THE SPACE OF LOVE:

PROBLEMATIZING THE INTERACTION OF CLOSE AND DISTANT READING

IN THE STUDY OF COURTLY LOVE

DAVIDE PAFUMI
Master of Arts, University of Padova, 2022

A thesis submitted
in partial fulfilment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

CULTURAL, SOCIAL AND POLITICAL THOUGHT

Department of English
University of Lethbridge
LETHBRIDGE, ALBERTA, CANADA

© Davide Pafumi, 2026

THE SPACE OF LOVE: PROBLEMATIZING THE INTERACTION OF CLOSE AND DISTANT READING IN THE STUDY OF COURTLY LOVE

DAVIDE PAFUMI

Date of Defence: August 12, 2026

Dr. D. P. O'Donnell Dr. B. Bordalejo Thesis Co-Supervisors	Professor Adjunct Professor	Ph.D. Ph.D.
Dr. M. Martin Thesis Examination Committee Member Department of English Princeton University	Professor	Ph.D.
Dr. A. Petrina Thesis Examination Committee Member Dipartimento di Studi Linguistici e Letterari Scuola di Scienze Umane University of Padova	Professor	Ph.D.
Dr. P. M. W. Robinson Thesis Examination Committee Member Department of English College of Arts and Sciences University of Saskatchewan	Professor	Ph.D.
Dr. M. A. Calabrese External Examiner California State University College of Arts and Letters Department of English	Professor Emeritus	Ph.D.
Jessica Copley Chair, Thesis Examination Committee	Assistant Professor	Ph.D.

DEDICATION

To my beloved ones. For the time I spent elsewhere, while you carried on your lives without me. Whatever is good in these pages was made possible by everything my absence asked of you.

ABSTRACT

This thesis asks whether computational methods can intervene productively in an established humanistic debate while remaining answerable to the interpretive traditions they seek to extend. It approaches this question as a methodological problem, using the longstanding debate over courtly love in later medieval English literature as a test case.

Reconstructing the Anglo-American debate from Gaston Paris through the decisive influence of C. S. Lewis, the thesis argues that its persistent impasse can be understood partly as a problem of archival restriction. The major interpretive traditions identified here as philosophico-theological, sociological, and psychological each recover genuine features of medieval representations of love, but do so from overlapping yet differently delimited bodies of evidence. Their internal coherence therefore derives partly from the evidentiary fields through which courtly love has been constituted as an object of inquiry. Drawing on Rüdiger Schnell, the thesis consequently reframes courtly love not as a unified doctrine or historical system but as a discourse: an ongoing and contested negotiation of love across texts, genres, authors, and linguistic traditions.

To make that discourse computationally analyzable, the thesis brings together two traditions of discourse analysis. A post-Foucauldian account provides the concept of the archive as a historically conditioned field of statements, while Zellig Harris's distributional hypothesis provides an empirical account of how patterned linguistic relations can be recovered from textual data. Spitzmüller and Warnke's DIMEAN supplies the multilevel framework through which these approaches can be synthesized, specifying how linguistic regularities at the level of words and texts may be interpreted as evidence of larger discursive formations. Vector-space modelling provides the computational means of representing those distributional relations.

The thesis also addresses the epistemological conditions under which such an intervention can count as humanistic interpretation. It argues that the opposition between close and distant reading rests on reciprocal misreadings in which each approach privileges a different object of authority: textual complexity and minutiae for close reading, and scale, archive, and evidentiary breadth for distant reading. Through a critique of the falsificatory logic exemplified by Matthew Jockers, this study argues that computational findings do not supersede or invalidate established interpretations, because humanistic interpretations need not succeed one another through falsification. Computational and non-computational approaches instead produce different, partial models of their objects. The consequential distinction is therefore not simply one of scale but of modelling: how evidence is bounded, represented, and made to stand for a larger field.

This framework is tested on a deliberately bounded corpus of Middle English texts. At the word level, distributional analysis shows that the theological, sociological, and psychological dimensions isolated by scholarship coexist within the semantic field of love. At the document level, the analysis uses C. S. Lewis as a benchmark because subsequent courtly-love scholarship has repeatedly developed from, revised, or defined itself against his synthesis. Testing the fourfold taxonomy of *The Four Loves* reveals that its categories do not map evenly onto the corpus and motivates a revised three-part description of literary discourse: courtly-romantic, secular-erotic, and theologico-philosophical. The purpose is not to falsify Lewis but to show how computational modelling can sharpen distinctions within an inherited interpretive framework.

Producing this analysis also required confronting the practical limits of computational work on pre-modern textual data. Corpus construction, periodization, textual selection,

orthographic variation, and cumulative processing noise make medieval evidence a problem of heterogeneous data integration rather than simply one of scale. The thesis develops an AI-assisted normalization procedure that maps dialectal and scribal variants onto shared Present-Day English lexical descendants while recording cases requiring semantic substitution for audit. These methodological solutions, together with the corpus-construction and modelling framework developed here, demonstrate how computational methods designed for large datasets can be adapted to small, historically mediated evidence. The thesis therefore offers a proof of concept for how computational modelling can complicate, refine, and extend established humanistic interpretation without claiming to replace it.

USE OF GENERATIVE AI

Beyond its role as an object of study, generative AI tools were used in a limited supporting capacity during the preparation of this thesis. Specifically, it was used to assist with proofreading, second-language correction, and editing of the text. Generative AI was also used to support code development, including debugging, annotating and improving its efficiency.

All results, interpretations, arguments, and citations presented in this thesis are the sole responsibility of the authors, who have verified the accuracy of all factual claims and references contained herein.

ACKNOWLEDGEMENTS

My deepest thanks go to my co-supervisors, Prof. Daniel P. O'Donnell and Barbara Bordalejo. Every page of this thesis carries the trace of the hours you spent reading, discussing, and redirecting work with remarkable patience. You guided me toward the right path more times than I can count and did so with a care I did not always make easy to extend. Your unwavering support has shaped both this thesis and the scholar I have become.

To Dr. Nathan D. Woods, my *shadow supervisor*. I am indebted to you for the countless hours at the whiteboard, working through ideas that were often only half-formed when I brought them to you. You had no formal obligation to give that time, and you gave it anyway, repeatedly, helping me sharpen arguments I could not yet see clearly on my own. I am deeply grateful for a generosity that went well beyond what anyone has a right to expect.

To my committee members, Prof. Meredith Martin, Alessandra Petrina, and Peter M. W. Robinson. Thank you for your availability and insight throughout this process. Your feedback shaped this work at every stage, and I found myself returning to it at exactly the moments I needed it most.

To Prof. David Kinney, prof. Simone Rebora, and Prof. David Wrisley. I am grateful you for the conversations, early in this project's conception, which set its methods in motion before I knew what shape the project would take.

To Prof. Warren Ginsberg, for advising on the courtly literature material, and to Prof. Brian J. Reilly, for his careful reading of my courtly love chapter. Thank you both for the attention you brought to work outside your own immediate obligations.

To the editorial board of the Middle English Text Series, in particular Dr. Anna Siebach-Larsen and Theresa Whitaker. Thank you for your generosity in waiving the copyright

restrictions on the data and for agreeing to publish the resulting work through the METS website. I hope that this work does justice to the extraordinary scholarly resource you have made available to the public.

To the colleagues and friends of my cohort, Stephanie Laine Hamilton, Janelle Marietta, Brendan Cummins, Morgan Pearce and Hanna Fantin, with whom I started this journey together. What began as stimulating conversation in our seminars grew over the years into something closer to a support group—the kind you can only build with people who are going through the same thing at the same time, and who show up for each other through it.

To the members of my lab goes my sincerest gratitude for the steady support along the way. To Jocelyn McKnight especially, for sharing the same passion for philology and for the conversations where, in the process of explaining my own work to you, I often ended up understanding it better myself.

I am also grateful to the professional communities that shaped this work: the Canadian Society for Digital Humanities/Société canadienne des humanités numériques, Digital Medievalist, and the Courtly Literature Society–North American Branch.

I owe a particular debt to the Medieval Association of the Pacific, and to Prof. Kimberly Klimek, who encouraged me to pursue this topic however expansive it seemed, reminding me that even the most explored terrain still holds lacunae worth finding. It is through connections made at MAP that I am now fortunate to have Prof. Michael Calabrese as the external reader of this thesis, whom I also thank heartfeltly.

Finally, to everyone else who supported this work in ways too numerous, or too easily taken for granted, to list here—thank you. This thesis carries more debts than any acknowledgements section can settle.

TABLE OF CONTENTS

DEDICATION iii
ABSTRACT iv
USE OF GENERATIVE AI vii
ACKNOWLEDGEMENTS viii
TABLE OF CONTENTS x
LIST OF FIGURES xii
LIST OF ABBREVIATIONS xiii
CHAPTER 1: INTRODUCTION 1
1.1 Preamble 1
1.1.1 About the thesis 7
1.1.2 A marginal note 13
CHAPTER 2: THE SCHOLARLY PROBLEM 17
2.1 About this chapter 17
2.2 If no love is 17
2.2.1 Boase: theories of meaning and theories of origin 18
2.3 The debate's precursors 21
2.3.1 Gaston Paris 22
2.3.2 C. S. Lewis 25
2.4 Love as a living anachronism 30
2.5 A different way of understanding the controversy 34
2.5.1 The philosophico-theological stake 35
2.5.2 The sociological stake 38
2.5.3 The psychological stake 40
2.6 Love as discourse 45
2.7 Next steps 52
CHAPTER 3: FRAMING THE METHODS 54
3.1 What is discourse, really? 54
3.2 Discourse après Foucault 56
3.3 The (post-)Foucauldian school 59
3.4 Discourse as language 64
3.4.1 Discourses and corpora 68
3.5 Harrisian distributional discourse 78
3.5.1 Meaning as space 84
3.6 DIMEAN as synthesis 87
CHAPTER 4: INTERPRETING WITH COMPUTERS 93
4.1 Recapitulation 93
4.2 The reading wars 95
4.3 Fetishizing complexity through close reading 98
4.4 The polemic against distant reading 108
4.5 The anxiety of distance 112
4.6 Distant reading's fetish for the *archive* 116
4.7 Melville, Milton, and Watts as examples 122
4.7.1 Jockers and the UMD English Department 123
4.7.2 Fish and Liberman 126

4.7.3 Watt and Jockers 129
4.8 Models Beyond Scales 134
CHAPTER 5: APPLICATIONS AND RESULTS 140
5.1 The heart of the matter 140
5.2 The corpus construction 143
5.2.1 Representativeness of the evidence base 144
5.2.2 Periodization and dating of texts 149
5.2.3 Special treatment of sources 152
5.2.4 Linguistic variation and etymological normalization 155
5.3 The space of meanings 163
5.3.1 Operationalizing the impasse 166
5.4 Why Lewis, and why *The Four Loves* 174
5.4.1 How the document-level analysis was constructed 179
5.5 Modelling the love discourse 183
5.5.1 Individual authors across the discursive space 189
5.5.2 What the revised map adds to the debate 192
CHAPTER 6: CONCLUSION 195
6.1 Implications for this study 195
REFERENCES 204
APPENDIX A: DATASET AND METADATA 223
APPENDIX B: ANNOTATED CODE 266

LIST OF FIGURES

Figure 1. Discursive continuum according to Jung..... 72
Figure 2. Hermanns's three corpus types..... 75
Figure 3. Levels of analysis in DIMEAN..... 83
Figure 4. Distribution of consonant classes in Milton's *Areopagitica*..... 127
Figure 5. Liberman's chart 1 showing the bilabial consonant distribution..... 128
Figure 6. UMAP projection of the love centroid's nearest neighbours..... 167
Figure 7. UMAP projections of anchor-based semantic neighbourhoods..... 171
Figure 8. PCA projection of Lewis's four-category scheme..... 177
Figure 9. PCA projection of the revised three-category scheme..... 185

LIST OF ABBREVIATIONS

AI	Artificial Intelligence
API	Application Programming Interface
BERT	Bidirectional Encoder Representations from Transformers
CBOW	Continuous Bag-of-Words
CSV	Comma-Separated Values
DH	Digital Humanities
DIMEAN	Diskurslinguistische Mehr-Ebenen-Analyse
GPU	Graphics Processing Unit
JSON	JavaScript Object Notation
LLM	Large Language Model
MED	Middle English Dictionary
METS	Middle English Texts Series
NLP	Natural Language Processing
OOM	Out of Memory
OOV	Out-of-Vocabulary
PCA	Principal Component Analysis
RAM	Random-Access Memory
TEAMS	Teaching Association for Medieval Studies
TTR	Type-Token Ratio
UMAP	Uniform Manifold Approximation and Projection
VRAM	Video Random-Access Memory
XML	Extensible Markup Language

CHAPTER 1: INTRODUCTION

L'errant, qui n'est plus le voyageur
ni le découvreur ni le conquérant,
cherche à connaître la totalité du monde
et sait déjà qu'il ne l'accomplira jamais.

—Édouard Glissant, *Poétique de la Relation*

1.1 Preamble

When Terry Eagleton published *After Theory* in 2003, declaring that "the golden age of cultural theory is long past," he was lamenting that "theory"—that is to say, explicit and self-critical reflection on the epistemologies used in humanities research—had been absorbed into academic routine, institutionalized to the point where its self-reflexive energy had dissipated (Eagleton 2003, 1). What he described was a field increasingly uncertain about its own epistemic objectives—not because it lacked methods, but because the animating questions those methods had once served were no longer generating the same productive friction. The irony, as it turned out, was that at almost the same moment that Eagleton was diagnosing this condition, DH was beginning to cohere as something recognizable, and with it came a new source of friction, and a new set of questions about what humanistic knowledge is and how it can be produced.

Stanley Fish gave one of the sharpest early formulations of the resulting anxiety in a 2012 *New York Times* column responding directly to the emergence of DH. Computation, he argued, does not extend humanistic inquiry but replaces its object, undoing the assumptions that once anchored literary interpretation—the unity of the author, the permanence of the text, the possibility of a reading whose significance exceeds its data (Fish 2011; Fish 2012b). What Fish was diagnosing—even if he did not quite name it—was a sort of categorical anxiety: the sense that something epistemologically alien had entered the field, and that the terms available for evaluating it—findings, results, patterns, evidence—belonged to a different intellectual culture

altogether. This is the culture of the intellectually fruitful yet unstructured exploration of texts without predetermined goals in what became known—approvingly or not—as "hermeneutics of screwing around" or "screwmeneutics" (terms coined respectively by Kirschenbaum 2010; Van Zundert 2015).

The question remains open whether computational approaches to literary and historical texts produce knowledge that is in a sense continuous with the humanistic inquiry these same approaches claim to extend—and if so, under what conditions, and of what kind. For Gavin, digital scholarship has consistently neglected this prior theoretical problem, substituting the question "What can digital methods contribute?"—a poor starting point of any such inquiry, in his view, because it privileges findings and results over reflection on the objects and conditions of computational knowledge (Gavin 2022, 5–6). However, this thesis argues, Gavin's position exemplifies the broader archive-centred tendency within digital scholarship examined here and, because of its limitations, provides the starting point for precisely the kind of inquiry he sets aside: whether a computational approach to literature via discourse analysis can engage with a traditional scholarly debate while each remains true to itself—intervening in such debate, refining it, and remaining answerable to it.

This thesis examines that question through one of medieval literary studies' most persistent debates and ultimately still unresolved: the discourse of love in medieval English literature. The study therefore approaches medieval love as a test case for a broader methodological problem rather than attempting to provide an exhaustive new account of courtly love itself. This is an object whose semantic variability, historical dispersion, and representational complexity make it unusually well suited to the kind of intervention proposed here. Medieval literature does not present a unified or stable conception of love; it offers instead

a multiplicity of often competing understandings: as devotional, didactic, and even parodic treatments of affective relationships or even a social code. Far from reflecting a single cultural consensus, these texts reveal a spectrum of meanings in which love is continually negotiated across time, genres, and authors—a complex discursive space that this thesis calls the *space of love*.

It is precisely in this space that a computational argument could prove most productive. Two recent studies establish the relevant precedents. Gabriel Viehhauser's work on *Minnesang* applies frequency analysis and topic modelling to test the long-debated hypothesis of a generic shift within the tradition, demonstrating that genres are defined by historically variable configurations of signals that gain or lose prominence over time (Viehhauser 2017). His study shows that even seemingly stable genre markers can be quantitatively tracked and shown to fluctuate diachronically, while new features emerge as alternative indicators of generic identity. Jean-Baptiste Camps et al. set out to test a literary-historical conjecture about the increased cultural importance of love in medieval French narrative by assembling a large corpus of digitized manuscripts and using embeddings to monitor the semantic environment of love across time (Camps et al. 2023). Their approach is especially relevant here because it shows that computational modelling can move beyond simple lexical counting and enriches existing literary knowledge. The results of Camps et al.'s study suggest a phase of expansion of love in medieval French fiction during the thirteenth and early fourteenth centuries, followed by a period of contraction associated with the crisis of the fourteenth century. What makes both these computational studies particularly instructive is not only their scale but the fact that they epitomize how computational methods can produce results that remain legible within existing literary-historical debates, rather than merely generating decontextualized quantitative patterns.

Building on these precedents for German and French, this study extends computational discourse analysis to Middle English while also taking the model one step further. It asks how the knowledge produced by such methods can enter into dialogue with the larger, predominantly non-digital tradition of scholarship surrounding a well-established literary problem. What is at stake is whether the medieval case can do what the epistemological argument requires of it: demonstrate, in practice rather than in principle, that close and distant reading are not competing but complementary, and that computational analysis can engage with, refine, and remain answerable to the interpretative debate it enters. The terms close and distant are retained here because they organize the disciplinary debate examined in this work, but they should not be understood as simple equivalents for non-computational and computational interpretation. The argument developed there is precisely that the opposition of scale obscures a more consequential distinction between different modes of modelling evidence.

As Andrew Piper argues, the emergence of distant reading has frequently been understood as a response to what he terms an "evidence gap"—the persistent difficulty of moving from part to whole in cultural criticism (2016, 3–4). Traditional literary scholarship has often relied on a small number of selectively chosen examples to support much broader claims, without adequately addressing the gap between the scale of the evidence and the scale of the generalization (Piper 2020, 4). This formulation identifies an important tension within literary-critical argument, but the model of knowledge implicit in it requires qualification. To define the problem principally as a gap in evidence risks reducing the distinction between computational and non-computational inquiry to one of quantity: traditional interpretations appear insufficiently supported because they derive from too few texts, while computational analysis promises to confirm, correct, or displace them by examining more—in other words, to falsify. Such a

formulation also moves toward the falsification framework through which computational studies have often sought disciplinary legitimacy. In both cases, prior interpretations are treated as hypotheses whose adequacy can be tested against an enlarged evidentiary field. Yet, humanistic interpretations do not ordinarily succeed one another through decisive falsification—they coexist. The epistemological task of computation is thus not to close an evidentiary deficit left by traditional criticism, but to bring the computational perspective into a productive relation with interpretations that operate according to different evidentiary and argumentative principles.

The richness and diversity of the surviving textual corpus—spanning romance, lyric, religious writing, and narrative poetry—provides an ideal ground for examining relations that become visible at a different level of observation from that of individual textual interpretation. It is precisely here that the significance of my research emerges, together with its broader methodological implications: namely, the possibility of questioning the applicability of computational linguistic methods to the interpretation of medieval texts, while also assessing their intellectual merit in enriching our understanding of historical literary discourse. Framed in this way, the research question is not entirely new, but rather rearticulates a foundational disciplinary concern of DH: how can computational and non-computational approaches produce different but complementary accounts of the same literary field without either claiming epistemological priority over the other?

This intervention is also shaped by the status of medieval textual evidence as Small Data. As Kitchin and Lauriault argue, the distinction between *Small* and *Big* data cannot be reduced to volume alone. Small data are generally limited in velocity, scope, and variety, are closely associated with particular research questions, and place especially strong demands on data quality, fidelity, and provenance (2015, 463–66). The Middle English corpus examined here is

small in this sense, but it is not therefore simple or internally homogeneous. Its surviving texts are generically broad while reaching the present through uneven and historically contingent processes of manuscript transmission, editorial selection, transcription, and digitization. The corpus is thus sufficiently diverse to disclose patterns unavailable through individual textual comparisons, but sufficiently unstable and mediated to make the construction of those patterns an interpretative problem in its own right. Nor is its breadth a claim to completeness: the corpus is a deliberately bounded and purpose-built approximation of a larger discursive field, and its exclusions therefore form part of the conditions under which its results must be interpreted. Chapter 5 consequently addresses not only what the corpus contains, but also identifiable absences and the limits those absences place upon the claims that can be made from it.

The significance of the research therefore lies not simply in applying computational linguistic methods to medieval texts, but in determining what adaptations are required before such methods can produce historically and philologically defensible knowledge that can be used by scholars whether or not they employ digital methods. Its broader methodological question concerns whether computational methods can contribute to the existing perspective of literary scholarship without treating scale as a substitute for interpretation, and whether the structures it identifies can enrich an existing account of historical literary discourse without claiming to supersede it. Framed in this way, the project rearticulates a foundational concern of DH: not merely how literary inquiry can examine more evidence, but how computationally constructed evidence differs from yet can be related to the objects and arguments of traditional criticism and how the two forms of knowledge can be placed in a productive relation.

1.1.1 About the thesis*

The methodological intervention developed in this thesis requires four related problems to be addressed. The first is theoretical: what is the established scholarly debate into which computation is being introduced, and why has it remained unresolved? Chapter 2 addresses this through the history of the courtly-love problem. The second is methodological: if, as that chapter argues, the debate is structured by differently restricted archives whose apparently shared terms designate different relations, what kind of analysis can recover discursive structure from a corpus without simply reproducing those restrictions? Chapter 3 answers this through discourse analysis. The third is practical: how can the fragmentary, heterogeneous, and linguistically unstable evidence of Middle English literature be constituted as a computational corpus, and what is gained or lost in doing so? Chapter 5 addresses this through corpus construction, normalization, and explicit consideration of the resulting evidentiary limits. The fourth is epistemological: even once such a method and corpus exist, how can computational research intervene in a traditional scholarly debate without treating existing interpretations as hypotheses awaiting confirmation or falsification? Chapter 4 addresses this through the close-vs-distant-reading debate. The chapters therefore do not constitute an extensive theoretical apparatus added before the analysis proper. Together they specify the conditions under which the analysis can count as the kind of intervention the thesis claims to make. Only once the scholarly problem, analytical framework, evidentiary model, and epistemological relation to previous interpretation have been established can Chapter 5 put that intervention into practice.

* Unless otherwise indicated, all translations from languages other than English are my own. Where a published English translation is used, this is identified in the relevant citation or note. Supporting materials are provided in the appendices: Appendix A lists the corpus and its metadata, while Appendix B contains the annotated Python code used in the data curation and processing. The shareable research data, metadata, trained model, and high-resolution figures are archived in Zenodo and are freely available at https://doi.org/10.5281/zenodo.21317382.

Chapter 2, then, traces the boundaries of the problem space that the computational argument will need to address. It asks the first necessary question: what is the debate this study wants to intervene in, and what relations has its existing organization left unresolved that a different mode of modelling might make visible? It reconstructs the historical development of the anglophone scholarly debate on courtly love from Gaston Paris's introduction of the concept in the nineteenth century through the field-defining synthesis of C. S. Lewis to the present, systematizing the extensive body of secondary literature around three interpretive stakes: philosophico-theological, sociological, and psychological. Each of these traditions identified something real in the primary texts; each was also built from an overlapping but differently delimited archive whose evidentiary priorities made some relations more visible than others and thereby contributed to the internal coherence of its interpretation. The chapter concludes, following Schnell (1989), that courtly love is best understood not as a static concept but as a discourse—an ongoing debate among medieval intellectuals about the meaning of love—and that this framing calls for discourse analysis as a method. Lewis occupies a distinctive position within this reconstruction because later scholarship repeatedly develops from, modifies, or defines itself against his synthesis. He therefore functions not simply as one participant in the debate but as an anchor through which its otherwise divergent interpretive traditions can be brought into relation.

Chapter 3 looks at two largely independent traditions of discourse analysis that become relevant to the methodological synthesis developed here and argues that what is often described as a confusion between them conceals a productive point of convergence. One tradition is historico-philosophical and centred on Foucault; the other is linguistic and centred on Harris. They do not share a direct genealogy or an identical definition of discourse, but they converge upon the problem of how larger structures of meaning can be reconstructed from patterned

relations among materially observable statements and linguistic forms. The chapter shows that post-Foucauldian scholarship expanded a restricted and methodologically bounded concept into a general theory of discourse, obscuring one possible compatibility between the two traditions. Its central claim is that the Foucauldian notion of archive and the linguistic corpus describe related but non-identical components of the analytical process: the archive designates the historically conditioned field in which statements become possible and intelligible, while the corpus provides the materially delimited means through which part of that field can be examined.

Jürgen Spitzmüller and Ingo H. Warnke's (2011) DIMEAN, short for *Diskurslinguistische Mehr-Ebenen-Analyse* (or Discourse-Linguistic Multi-Level Analysis), provides the methodological framework through which discourse can be examined across interconnected levels of analysis, from linguistic features within texts to broader discursive formations. DIMEAN is not a software package, nor does it prescribe computational procedures such as frequency analysis, distributional modelling, or vectorization. It is a model for the systematic human analysis of discourse, in which phenomena identified at different linguistic and textual levels can be brought into relation with one another. In this thesis, its role is therefore interpretive rather than computational. Harris's distributional hypothesis supplies the linguistic principle according to which patterns of distribution can be treated as evidence of semantic relation, while vector-space modelling provides the computational means of representing those relations. DIMEAN enters after, and alongside, those procedures as the framework through which findings made at the word, text, and document levels can be related to the larger discourse under investigation. The connection between DIMEAN and vector-space modelling is thus a methodological synthesis developed for the purposes of this thesis rather than something contained within DIMEAN itself.

Chapter 4 addresses a separate but equally necessary question raised by the nature of the proposed intervention itself: can computational methods engage with a traditional literary debate? The chapter engages with the intellectual history of close reading and distant reading, arguing that the apparent opposition between them is best understood as anxiety of influence (Bloom [1973] 1997) in which each approach fetishizes a different object of authority: minutiae and textual complexity for close reading, and archival breadth for distant reading. The latter is accompanied, in some influential formulations of computational literary study, by a further epistemological claim: that the enlarged evidentiary field made available by computation permits existing literary-historical interpretations to be confirmed, corrected, or falsified. The chapter develops its critique of this position particularly through Matthew Jockers, whose appeal to confirmation and Popperian falsification makes the epistemological consequences of the distant-reading argument unusually explicit. The falsification framework through which computational work has sometimes sought legitimacy in literary studies constitutes a misjudgement, since humanistic interpretations do not succeed one another but simply coexist. A computational representation of an interpretive claim may reproduce, qualify, complicate, or redirect it without thereby standing in a straightforward confirmatory or falsificatory relation to traditional interpretations. The chapter concludes that the real distinction between computational and traditional work lies not in scale but in modelling—i.e., the explicit articulation of how evidence stands in for a larger field under stated conditions. The archive, theorized across the preceding chapters, provides the point at which this resolution becomes concrete: corpus construction is not a preliminary technical step but the modelling decision that determines what the analysis can legitimately claim. Only once this is established can the analysis proceed without the spectre of replacement hanging over it.

Chapter 5 constitutes the thesis's analytical core. Given the methodological commitments established in the preceding chapters, how is the corpus to be constructed, how are the tools to be applied, and what does the analysis actually show? The chapter opens by confronting directly the practical problem that any computational engagement with medieval literary data must solve: how a materially incomplete and textually unstable surviving record can be transformed into a deliberately bounded corpus suitable for the particular analytical question being asked. Problems of textual variation, periodization, linguistic boundaries, composite works, editorial provenance, and orthographic normalization are therefore treated as modelling decisions rather than neutral preprocessing choices. The discussion also makes explicit what the resulting corpus does not contain. Its 155 documents are not offered as a comprehensive representation of Middle English literature or of medieval love discourse as a whole. The omissions identified in relation to the available textual record—including works that a close-reading account might reasonably expect to encounter—therefore delimit the claims made from the model. Some may be addressed in expanded iterations of the corpus; others arise from the larger problems of survival, accessibility, editorial mediation, and computational tractability that condition any retrospective analysis of medieval discourse. The purpose of the present corpus is consequently proof-of-concept representativeness: to provide an explicit and internally coherent evidentiary field within which the methodological intervention can be tested.

The analysis then proceeds at the word and intertextual level. At the word level, vector-space modelling reconfigures the existing debate by showing that the stakes identified in Chapter 2 are not competing explanations of separate phenomena: they describe overlapping semantic tendencies that coexist within the discourse of love itself. The analysis traces how these dimensions converge, separate, and recombine around the lexical field of love, revealing points

at which devotional, affective, social, moral, and courtly meanings become mutually constitutive rather than categorically distinct. At the document level, contextual embeddings extend this argument by examining how inherited typologies of love correspond to the distributional organization of whole texts. C. S. Testing Lewis is therefore a way of testing the organization of the field rather than an arbitrary choice among possible taxonomies. The *Four Loves* is used because its distinctions among Affection, Friendship, Eros, and Charity make explicit and operationalizable a broader conception of love that remains continuous with concerns developed across Lewis's earlier and later work. The analysis does not ask whether Lewis can be falsified, but how those categories behave when represented at the level of document distributions: which correspond to recurrent textual formations, which become difficult to sustain as autonomous categories, and which literary formations the inherited framework leaves comparatively obscure. The resulting profiles show that established distinctions retain explanatory value, but that their boundaries are neither stable nor exhaustive: different configurations of affect, devotion, social relation, desire, and moral evaluation overlap within texts and are further mediated by genre and textual tradition. The document-level analysis consequently motivates a revised three-coordinate description of the corpus into: courtly-romantic, secular-erotic, and theologico-philosophical. However, these categorization does not act as a replacement taxonomy of human loves but as an empirically motivated model of recurrent literary-discursive organization.

The conclusion returns from the medieval test case to the methodological question with which the thesis began: under what conditions can computational analysis intervene productively in an established humanistic debate? The answer offered by the study is that computation need not validate, falsify, or supersede previous interpretation in order to contribute to it. Its contribution may instead consist in representing the evidentiary field differently, making

recurrent relations, proximities, exclusions, and asymmetries visible at another descriptive level and returning those relations to interpretation. The refinement of Lewis and the revised three-coordinate model constitute the substantive demonstration of that possibility rather than a comprehensive new theory of medieval love. Reaching that result required a comparatively extensive apparatus developed to integrate unstable and scattered textual evidence, corpus construction, computational modelling, and established Humanities frameworks. That apparatus is itself part of the methodological contribution rather than merely the preliminary cost of obtaining the literary result. The difficulties encountered in moving from manuscript-derived texts to interpretable results reveal that computational methods cannot simply be applied to medieval evidence off-the-shelf. They must be adapted to the conditions of Small Data, that is: noisy, heterogeneous—and therefore difficult to process—and historically contingent datasets (Kitchin and Lauriault 2015, 463–66). The study's broader implication is that adapting techniques developed for large-scale analysis to Small Data exposes data integration as a central methodological problem and an important desideratum for future research.

1.1.2 A marginal note

The chapters of this work are also designed as both a series of responses to questions raised during the project's development—particularly in dialogue with my committee—and as a *vademecum*. This is to be understood not as a prescriptive model for all retrospective analysis of medieval discourse, but as a worked and reusable methodological framework: one possible sequence for moving from the identification of an established research problem, through the choice and justification of an analytical approach, to the evidentiary and computational decisions required to intervene in that problem. It is intended to be accessible not only to readers well-versed in digital medieval studies but also to those who may not identify as digital medievalists,

yet remain interested in the theoretical development of discourse analysis or in the study of medieval love as a discursive phenomenon. Its primary methodological aim is therefore pursued through a substantive intervention in a specialized field: the two are related as framework and test case rather than as independent projects of equal scope.

At the same time, the approach adopted here addresses a gap that, while not entirely new, has remained insufficiently explored. Even though the field of medieval literary studies has long engaged with the idea of love as a "discourse," it rarely incorporates the conceptual and methodological frameworks developed within discourse analysis. Conversely, discourse analysis as a field has produced a substantial body of work on how meaning is structured, transmitted, and transformed, but this work has not been fully integrated into the study of medieval literature. The result is a paradox: if discourse is often invoked in relation to medieval love, the criteria through which discourse is defined and analyzed are not consistently applied. In this respect, vector-space models provide one computational means of representing linguistic relations that can then be interpreted through the discourse-analytical framework established in Chapter 3. Overall, the project seeks to bring together strands of inquiry that have often remained separate, offering a framework in which literary analysis and discourse theory can be considered in relation to one another.

What began as a project of corpus-assisted discourse analysis gradually moved closer to Natural Language Processing—though without initially recognizing that NLP itself forms part of the broader AI endeavour to enable machines to understand language in ways that approximate human cognition. At the core of a vector space approach, the ambitions of early AI to model human cognitive design remain legible. In a different but related way, this thesis approaches the semantics of love as a structured system of associations between language and meaning,

represented through discourse by word embeddings: what artificial intelligence attempts at a broader level, this thesis attempts, in a more modest and historically situated way, at the level of literary discourse. The pace at which AI is evolving introduces an unavoidable problem of procedural obsolescence: the specific tools and methods employed here may become outdated within a relatively short period. The conceptual framework within which they are deployed, however—the argument that computational modelling can engage with traditional literary-historical debates on their own terms, rather than bypassing or replacing them—is intended to outlast any particular tool or implementation.

This thesis makes a deliberately bounded literary intervention, but the methodological apparatus required to produce that intervention is itself one of its principal contributions. Its extent reflects the problem the thesis set itself: not simply to apply a computational method to medieval texts, but to establish the conditions under which a computational representation can enter an existing and valid humanistic framework without bypassing the questions that framework is entitled to ask—and those questions, as the preceding chapters show, are neither few nor simple. The thesis does not aim to resolve the discourse of love in later medieval English literature in its entirety, nor to advance a *grand récit* where the existing scholarship has wisely resisted one.

The thesis is therefore more of a proof of concept, which operates at different levels. Substantively, it demonstrates that computational modelling can refine an inherited interpretive framework by making relations visible that the framework does not organize in the same way. Methodologically, it establishes a reusable sequence for moving from a contested scholarly object through discourse analysis, corpus construction, computational representation, and interpretive return. As per the results, the analysis offers instead a deliberately partial and

revisable account: a set of propositions intended to illuminate rather than exhaust a field whose historical dispersion and conceptual complexity exceed what any single study, computational or otherwise, can map. What it does claim is that the groundwork laid here is reusable. Others can take the same approach to other unresolved problems in the literature with considerably less preliminary labour than this thesis required.

In this sense, the study aligns with what Édouard Glissant ([1990] 1997) describes as a *poétique de la relation*: a mode of inquiry that resists totalization in favour of relational understanding, shaped not by metaphors of conquest or discovery but by the figure of the *errant*—the wanderer who seeks to know the totality of the world while recognizing in advance that such knowledge can never be complete. The exploration of the space of love proposed here remains necessarily incomplete. That incompleteness is a condition of the work's epistemological honesty, and it is precisely what keeps the study open to further interpretation, methodological refinement, and the continuing evolution of the discourse it seeks to explore.

CHAPTER 2: THE SCHOLARLY PROBLEM

S'amor non è, che dunque è quel ch'io sento?
Ma s'egli è amor, perdio, che cosa et quale?

—Francesco Petrarca, *Rerum Vulgarium Fragmenta*

2.1 About this chapter

The previous chapter advances the idea that the relationship between computational and traditional humanistic inquiry should not be framed as one of replacement, but as a question of whether the two can enter into a complementary and mutually answerable relation. This chapter turns from that general epistemological question to a specific scholarly debate in which such an intervention is both necessary and testable: the long-running debate over courtly love. Its purpose is not merely to recount the history of that debate, but to identify the structure of the impasse that has sustained it. In it, I argue that its various philosophico-theological, sociological, and psychological traditions do not simply offer competing interpretations of a common body of evidence: they construct different versions of medieval love from overlapping but differently delimited evidentiary bases—which we can also describe, anticipating subsequent chapters, as "archives." By showing how those archives shape the questions scholars ask, the relations they privilege, and the meanings they recover, the chapter establishes the problem to which the later computational analysis responds: how interpretations developed from separate evidentiary traditions might be placed within a common comparative framework without treating any one of them as simply correct or obsolete.

2.2 If no love is

For almost a century and a half since the introduction of the term in the scholarly debate, medievalists have discussed the nature of "courtly love" in the Middle Ages. The concept was developed to explain the heterogeneous range of love's literary instantiations during the Middle

Ages—a range that resists a singular definition and complicates scholarly synthesis due to its contradictory characteristics. John C. Moore synthesizes these contradictions by describing how courtly love has been variously defined as adulterous, yet also as virtuous and pure; as an expression of spiritual elevation yet as an erotic, sensual pursuit; as a choice freely made, yet equally as a force driven by fate or uncontrollable passion (Moore 1979). Theories on its origins are similarly diverse. As Roger Boase remarks in his book on *Origin and Meaning of Courtly Love*, courtly love's themes can be attributed to the influence of Ovidian ideas on romantic desire; they can be traced to the moral and spiritual teachings of orthodox Christianity; or their inspiration can be attributed to the cultural exchanges with the Arab world (1977). In his view, critiques of medieval courtly love suffer from an "intellectual malaise" fuelled by the fragmentation inherent to academic specialization (Boase 1977, 122). The resulting instability cannot, however, be attributed solely to the variety of medieval representations themselves; it also reflects the different bodies of evidence through which scholars have constituted love as an object of inquiry. Its definition changes not only according to period, place, and author, but also according to which texts are selected, which relations among them are foregrounded, and which explanatory priority the critic applies to the archive.

2.2.1 Boase: theories of meaning and theories of origin

In order to cope with this fragmentation, Boase provides a helpful framework for understanding the scholarly debate up to the 1970s, distinguishing between two main approaches: theories of meaning and theories of origin (1977, 3). Theories of meaning aim to interpret what courtly love represents or signifies within its socio-cultural and historical context, while theories of origin focus on tracing its beginnings and development over time. Boase attempts a synthesis of these theories developed to offer a comprehensive view of the modern erudite discussion as well as the

different meanings attached to the term up to the 1970s (1977, 122). Yet, even with this framework, Boase demonstrates that, despite the extensive body of work on courtly love, the scholarship remains fraught with unresolved interpretive challenges, reflecting what he calls a "bewildering variety of uses and definitions" (1977, 1).

The importance of Boase's survey lies therefore not in resolving the debate, but in revealing the reciprocal relation between origin, meaning, and evidence. As Boase explicitly states, theories of meaning and theories of origin cannot really be separated from one another (1977, 3). Indeed, the meaning of courtly love often depends on understanding its origins. If courtly love is assumed to derive from Ovid, its interpretation will tend to emphasize passion, eroticism, and literary convention; if it is located within Christian doctrine, it is more likely to appear as a spiritual, moral, or heterodox formation. The same relation holds for the archives from which those arguments emerge. Each interpretative tradition identifies genuine features of medieval writing, but each makes those features central by organizing a restricted set of texts and concepts around them.

Talbot E. Donaldson provides one of the clearest formulations of why such differences matter for what this chapter terms the archive, and what might more generally be described as the corpus available to interpretation. A definition based upon all medieval literature, he argues, becomes too broad to remain useful, whereas one "derived from only selected primary documents fits well only those documents from which it has been derived" (1970, 156). Moreover, "different documents yield divergent data," while even "a single document or a single body of related documents yields divergent data to different scholars" (157). Jill Mann extends the problem by observing that multiple forms of courtly love may coexist even within the work of a single author (2013, 93). These observations do not imply that evidence mechanically

determines interpretation. They suggest instead that the delimitation of an evidentiary base and the explanatory questions brought to it operate together: selection of certain sources makes some relations especially visible, while interpretative priority determines which of those relations are treated as fundamental.

Seen in these terms, the debate is not simply a collection of incompatible definitions. It is a series of differently constructed evidentiary objects. Boase's work clarifies not just points of agreement but, more importantly, where scholars diverge. His work demonstrates that, while no definitive solution has emerged, it is possible to map the broad spectrum of unresolved interpretive challenges within the scholarly debate. In fact, as Moore argues, the term "courtly" superimposes an artificial label on the cultural production from this time, implying that all literary expressions of love were tied to aristocratic settings (1979, 624). Representations produced outside that setting are consequently treated as exceptions or deviations from an assumed norm, rather than as evidence that the object itself may have been too narrowly constituted. "Courtly" love therefore imposes a restrictive and often artificial unity upon a much more varied body of medieval representations. The difficulty is that simply abandoning the term does not remove the recurrent relations among love, desire, devotion, marriage, service, morality, and social status that the term has imperfectly attempted to name. The category may be unstable, but the field of textual relations gathered beneath it remains available—and necessary—for analysis.

This chapter, then, aims to trace the history of the concept from its introduction and early interpretations by Paris, culminating in Lewis's influential but polarizing account, before examining the critiques that exposed its anachronism without dissolving the scholarly problem it had created**.** It then organizes the major historical and critical perspectives of the scholarship on

the semantics of courtly love, identifying three primary interpretive stakes in the debate: (a) philosophico-theological, (b) sociological, and (c) psychological. These stakes should not be understood as mutually exclusive schools, since each identifies one of the three recurrent explanatory configurations of evidence and explanation. One that interprets love through religious doctrine and spiritual conflict, one that situates it within institutions and social relations, and one that approaches it through desire, affect, and subjectivity. Dividing the schools in this fashion, however, does not eliminate the problem identified by Boase in his analysis of the theories of "Meaning" and "Origins." The impasse persists because philosophico-theological, social, and psychological meanings also coexist within medieval texts.

Drawing on Schnell (1989) and the lineage of research that—consciously or not—has developed from positions similar to his work, this chapter argues that courtly love is better understood as a discourse than as a unified doctrine, code, or singular historical phenomenon. Reframing the object in this way does not resolve its contradictions by selecting one interpretation over the others. It treats those contradictions as constitutive features of a discursive field in which religious, social, affective, moral, and erotic meanings are repeatedly articulated in relation to one another. The methodological question passed to the following chapter is therefore how that field can be reconstructed from a necessarily partial corpus without merely reproducing the categories through which scholars have previously divided it.

2.3 The debate's precursors

If the instability of courtly love arises partly from the evidentiary bases through which it has been interpreted, the first task is to determine how that evidence came to be organized as a scholarly object. The construct under discussion emerged through the selection of particular texts, genres, languages, and relations that established what would subsequently count as

evidence of courtly love. Gaston Paris and C. S. Lewis occupy foundational places in this process: Paris derived a general category from a highly specific romance tradition, while Lewis systematized the category into a generalizable framework, which gained popularity in Anglo-American academia. Together, they established the tradition within which much later scholarship would continue to operate through both the source and conceptual vocabulary of the debate.

2.3.1 Gaston Paris

In his two-part study "Études sur les romans de la Table Ronde," published in *Romania* in 1881 and 1883, Gaston Paris marks the moment at which love in medieval literature became a formal object of scholarly inquiry (Paris 1881; Paris 1883). Earlier writers had treated the topic. Stendhal's *De l'Amour* (1822), in particular, had introduced the crucial idea that love was historically and culturally variable rather than universal, and had drawn on Andreas Capellanus' *De amore* in proposing an Arabic provenance for the courtly tradition (O'Donoghue 2006, 9). But Stendhal's interest was in love itself, not in its medieval literary representations as objects requiring systematic historical explanation. It was Paris who made the historical turn—who took the relations dramatized in a particular romance (Lancelot and Guenevere's illicit love in Chrétien de Troyes' *Chevalier de la Charrette*) and subjected them to the kind of analysis that would produce a general scholarly category. For Paris and his contemporaries, the code of courtly love and the problem of the seriousness of Capellanus would prove generative topics that remain central to subsequent scholarly discussion. Paris did not invent the term *amour courtois* outright but re-coined it, drawing on prior instances—the troubadour Peire d'Auvergne in the twelfth century and, a century later, the *Roman de Flamenca* (O'Donoghue 2006, 9; although, as is rarely noted, the first mention in Paris's own work is found in the 1881 instalment rather than in 1883 as commonly reported; see Hult 1996, 199–200). What proved decisive was not the

lexical novelty of the term but the analytical authority Paris gave it. In the 1883 essay, he defines *amour courtois* as "an art, a science, a virtue, with rules just like chivalry or courtesy—rules that one masters and applies better as one makes more progress, and which one must not fail to follow under penalty of being deemed unworthy" (Paris 1883, 519).

As Mann observes, Paris identified its four main characteristics of this system: (1) Courtly Love is illicit and secretive; (2) the lady has a social advantage over her lover; (3) the male lover is in a subordinate, anxious position, while the lady remains distant and capricious; (4) to prove his worth, the lover must perform remarkable feats, as the lady imposes demands to improve him further (2013, 98–9). Courtly love consequently appears as a code of conduct centred upon an idealizing lover who seeks to prove himself worthy through courage, honour, service, and obedience. Sexual attraction is necessary to this structure, but sexual satisfaction cannot constitute its uncomplicated goal. The lady is frequently married and therefore unattainable, introducing adultery as both a constitutive condition and a persistent moral problem. Paris's definition thus converts the narrative organization of one especially influential romance into a general account of amorous behaviour. The particular relations among Lancelot and Guinevere, secrecy and service, hierarchy and ordeal become the criteria by which other literary representations may be recognized as instances of courtly love—or excluded as deviations from it. The evidentiary base matters because Paris's interpretation appears to derive naturally from the material only after the material has been delimited in accordance with the interpretation. A romance centred upon an adulterous and socially unequal relationship predictably makes certain traits appear foundational to the meaning of courtly love. This is not to say that the selection and organization of the evidence predispose the inquiry towards Paris's conclusions mechanically, but it makes some conclusions considerably more available than

others. His account moves beyond literary description by treating the relations found in *La Charrette* as evidence of a wider cultural system. The result is that the selected romance supplies the traits of courtly love, while those traits subsequently determine which other texts belong to the courtly-love tradition.

The historical circumstances of Paris's scholarship further shaped this movement from literary instance to cultural generalization. As Hult argues, Paris's work belongs to the intellectual and political environment of post-Franco-Prussian France and to a positivist scholarly culture invested in establishing the national importance of French medieval literature (1996, 195). Paris's notion of courtly love is therefore anachronistic not simply because it applies a modern label to medieval materials, but because it projects a form of conceptual coherence developed within nineteenth-century scholarship back upon heterogeneous medieval practices. The distortion lies in making the analytical category appear contemporaneous with the evidence it was created to organize. Modern understandings of medieval romance thus came to be shaped by Paris's conceptualization rather than solely by the more variable and context-specific representations preserved in the texts themselves. Between Paris and Lewis, whom we shall discuss next, the category underwent a consequential transformation: it ceased to be a literary-critical interpretation and began to appear as a historical social system. Lewis Freeman Mott (1896) and William Allan Neilson (1899) cast courtly love as a strict code governing social and amorous relationships, treating Paris's interpretation as a recoverable system that could be reconstructed across a wider body of literature. The same assumption is apparent in John Parry's translation of Capellanus' *Art of Courtly Love*, which promises to depict "that strange social system to which Gaston Paris has given the name of 'courtly love'" (1941, 3). By the time of Parry's formulation, a concept derived from modern interpretation had acquired the appearance

of a historical social system—or an institution—for which medieval texts supplied documentary evidence.

2.3.2 C. S. Lewis

It was in this quasi-institutionalized form that C. S. Lewis received the concept. In particular, his *Allegory of Love* (1936) gave it its most influential Anglophone formulation, describing courtly love as love "of a highly specialized sort," characterized by humility, courtesy, adultery, and the religion of love ([1936] 2013, 2). His intervention substantially expands Paris's evidentiary field. Rather than deriving the category primarily from a particular French romance, Lewis relates Provençal lyric, Ovidian convention, feudal service, Andreas Capellanus, Christian morality, allegory, and later French and English literary traditions. Yet this enlarged archive remains governed by the assumption that these materials participate in a coherent historical formation. The different sources do not merely offer contrasting representations of love; they become stages or components within an account of how courtly love emerged and developed. Lewis argues that a refined form of love first appeared in French troubadour poetry and subsequently shaped medieval literature in France and England. Although he rarely describes courtly love explicitly as a "code," his four defining traits align with Paris's representation of an intense passion governed by recognizable behavioural rules. Lewis further connects those conventions to "the actual practice of feudal society" ([1936] 2013, 16; see Mann 2013, 92). The language of amorous humility and service is thereby read in relation to feudal hierarchy, while its apparent religiosity is explained through the adaptation of Christian and Ovidian forms to an illicit object of devotion.

Marriage occupies a central place within this argument. Lewis contends that medieval marriage was ordinarily determined by practical considerations of property, alliance, and social

status rather than by romantic attachment. The idea of love within marriage could therefore be dismissed as "nonsense" ([1936] 2013, 16). Against this utilitarian institution, courtly love appeared as an idealized and emotionally refined relation founded upon tenderness, service, and desire. Lewis's most consequential claim is that, in a society where marriage is principally practical, the idealization of romantic love will tend toward the idealization of adultery ([1936] 2013, 16–17). At the same time, passionate attachment remained morally suspect even within marriage, leaving the medieval lover to choose between repentance and varying degrees of guilt. Lewis consequently makes the tensions among marriage, adultery, religious morality, erotic desire, and social hierarchy mutually explanatory. His courtly-love model is neither exclusively philosophico-theological nor sociological nor psychological; it is the foundational synthesis from which those later explanatory priorities can be separated. Lewis's importance therefore lies not merely in popularizing Paris's terminology, but in consolidating a broad archive around it. Under his lasting influence, generations of scholars continued to analyze medieval relationships as examples of, or deviations from, a courtly-love system whose defining characteristics were assumed to be already known. Its principal problems—idealization, the lover-beloved hierarchy, adultery, marriage, moral ambiguity, sexual restraint, and social service—became recurring points of reference even for scholars who rejected Lewis's conclusions. Medievalists progressively treated courtly love as a *datum*: a formalized and apparently self-evident phenomenon against which individual texts could be measured.

Lewis returned to the classification of love in *The Four Loves* (1960), although this later book did not acquire a comparable role in the scholarly discussion of courtly love. The reasons are partly circumstantial. In 1958, the Episcopal Radio-TV Foundation of Atlanta invited Lewis to prepare a series of short talks for broadcast in the United States, leaving him free to select the

subject. Lewis proposed the four loves—Affection (storge), Friendship (philia), Eros (romantic love), and Charity (agape)—as a theme that would encompass much of Christian ethics, and recorded ten talks in London in August 1958. The recordings and their separately issued transcripts supplied the basis for the substantially expanded book published two years later (see Lepojärvi 2022). The *Four Loves* was therefore conceived for a general broadcast audience rather than as a contribution to medieval literary history, and its conversational, theological, and pastoral orientation separates it from the historical argument of *The Allegory of Love,* even where Lewis returns to medieval material, but the reasons are also professional.

By 1960, Lewis also had spent two decades producing popular theological works—from *The Problem of Pain* (1940) and *The Screwtape Letters* (1942) through the BBC talks that became *Mere Christianity* (1952) to the Narnia chronicles (1950–56)—that made him one of the most widely read Christian writers in the English-speaking world and one of the most resented within his own faculty. When the Merton Chair of English Literature became vacant in 1947, as Wilson puts it, Lewis was passed over in favour of a candidate who

> Could be relied upon not to cause trouble, and not to embarrass his colleagues by writing books about Christianity. Even colleagues who were Christians found Lewis's career as a popularizer embarrassing, 'not quite the thing.' (Wilson 1990, 208)

Lewis did not receive a full professorship until Cambridge offered him a newly established chair in 1954, by which time the pattern was entrenched. *The Four Loves*, begun as a series of radio talks and published six years after his departure from Oxford, belongs squarely within the body of work that provoked this professional cost. Later medievalists treated it accordingly: they continued to engage with the historical model of *The Allegory of Love* rather than with the transhistorical taxonomy of a book their discipline had no reason to take as scholarship. *The Four Loves*, however, is relevant to the courtly love debate. It distinguishes Affection, which

Lewis associates with *storge*; Friendship, related to the classical traditions of *philia* and *amicitia*; Eros, understood as the condition of being "in love"; and Charity, the divine love through which the others may be transformed (Lewis 1960, 42, 69–71, 106, 133). These are not four equivalent historical categories: affection, friendship, and eros are natural loves, whereas charity occupies a distinct theological position, since the natural loves are "summoned to become modes of Charity" while remaining the kinds of attachment they already are (151–52).

In *The Discarded Image* (1964), Lewis returns to several of the relations that had structured his earlier accounts of love. Describing the medieval tendency to organize experience into formal systems, he observes that just as war was regulated by chivalry, "sexual passion" was formalized by "an elaborate code of love" (Lewis 1964, 10). In his discussion of the medieval cosmos, Lewis explains Aristotle's Prime Mover as one who "moves as beloved," so that the universe is set in motion by the desire of creatures for God (1964, 113). He then distinguishes this ascending love from the "provident and descending love" of God characteristic of Christianity, while insisting that the two senses of the "love of God" can coexist within the same universe (113–14).

These works concur in their treatment of love and its correlates, revealing a notable consistency across Lewis's career in his handling of adultery, marriage, divine love, sin, and the ordering of desire. The same pattern extends beyond the works most readily accepted within medieval literary scholarship to the theological and more popular writings that have generally remained peripheral to the courtly-love debate. A line can be traced from the historical and literary account developed in *The Allegory* in the 1930s, through the treatment of desire and sin in *The Screwtape Letters* in the 1940s, to the more systematic classification of love in *The Four Loves* and the broader medieval cosmological synthesis of *The Discarded Image* in the 1960s.

Throughout this trajectory, Lewis returns to questions concerning human attachment, erotic desire, moral order, and divine love.

The later taxonomy develops concerns already present in *The Allegory* twenty-four years earlier. *The Allegory*'s courtly love is recognizably a configuration of Eros—defined by longing, service, and idealization. Its "religion of love" names the tension between Eros and Charity: in *The Allegory*, that religion had appeared principally as a historically specific complex of literary conventions; in *The Four Loves*, Lewis treats it instead as one manifestation of Eros's more general tendency to demand absolute allegiance and become "a sort of religion" (127–28). *The Allegory*'s emphasis on feudal hierarchy, courtesy, and companionship addresses the social relations that Friendship and Affection organize differently. *The Four Loves,* therefore, does not abandon the material of *The Allegory of Love*; it abstracts that material into a transhistorical framework whose categories—natural loves in need of divine correction—are no longer tied to any single literary period.

Yet this reinterpretation did not substantially redirect the specialized debate over courtly love. Later medievalists continued to engage primarily with the historical model established in *The Allegory of Love*—its account of troubadour innovation, feudal service, adultery, and the religion of love—rather than with the transhistorical taxonomy of the later book. *The Four Loves* became influential chiefly as a work of Christian ethics and popular theology, while its revision of Lewis's earlier interpretation remained largely peripheral to arguments about the existence, origins, and defining characteristics of courtly love. The contrast in reception matters here: the historical generalisations of *The Allegory of Love* entered medieval scholarship as claims to be adopted, modified, or rejected, whereas the categories of *The Four Loves* did not become the principal conceptual vocabulary through which medieval amorous relations were debated.

2.4 Love as a living anachronism

Almost as soon as Lewis's synthesis was published, scholars began to critique it—and over three decades their objections accumulated into a kind of criticism with a central claim that was difficult to dismiss: i.e., that courtly love, in the way Paris had conceived it and Lewis had codified it, was a modern construction projected backwards onto medieval evidence, an anachronism whose coherence belonged to the scholarship that produced it rather than to the medieval culture it purported to describe. A 1967 symposium held at the State University of New York at Binghamton—under the auspices of the International Courtly Literature Society—brought previously dispersed positions—both supportive and skeptical—into direct confrontation; its proceedings appeared the following year as Francis X. Newman's *The Meaning of Courtly Love* (1968). Boase's later "chronological survey" of European scholarship, extending from the sixteenth century to the twentieth, retrospectively organized the debate into a historiographical narrative that clarified its major questions without claiming to resolve them (1977, 5–26). Courtly love was no longer only a proposed feature of medieval culture; it was also an inherited scholarly problem. As David Hult remarks, the assumption that this modern historiographical concept reflects a coherent medieval attitude amounts to "one of the great critical anachronisms of all time" (1996, 194), or, in Howard R. Bloch's formulation, "of the great scholarly errors of all time" (1991, 69). The critique formulated by these scholars exposed the category as artificial, but it did not dissolve the literary relations that had made it persuasive, nor did it alert subsequent scholars to the implications of the archive through which it had been constituted.

John Tatlock was the first to say so directly, calling the concept "a fundamental modern error" as early as 1941 (88). By the 1960s and 1970s, the critique had sharpened and differentiated, with scholars arriving at the same conclusion from distinguishable starting points.

For Durant W. Robertson, the problem was hermeneutic: courtly love distorted the interpretation of medieval texts by obscuring the Christian allegorical framework through which they should properly be read. Following William Matthews, Robertson described the concept as an "impediment to the understanding of medieval texts" (Robertson 1968, 1; see Matthews 1965). For John F. Benton, the objection was historical: the literary conventions attributed to courtly love could not be shown to correspond to documented social practice, and the term offered "no advantage" since it "confuses the issue" by conflating convention with lived experience (1968, 37). For Talbot E. Donaldson, the difficulty was one of specificity: what was called courtly love was not distinctively medieval but a set of attitudes toward erotic experience that were "no more real in the Middle Ages than they had been before and have been since" (1970, 163). James J. Wilhelm went further still, recommending that the category of "courtly love poetry" be rejected *ex toto* (1965). John C. Moore subsequently turned this critique into an explicit terminological prescription, dismissing courtly love as "the creature of every wind stirred up by scholars" and "a mischievous term which should be abandoned" (1979, 630–31). Conor McCarthy later observed that C. S. Lewis's claims about medieval marriage—that it was governed principally by material interests, that passionate attachment remained morally suspect even between spouses—substantially oversimplified the social and emotional possibilities available in the period (2022, 7). The difficulty was not merely that individual claims were inaccurate. It was that by presenting one configuration of love as representative of medieval culture, the inherited framework made texts intelligible according to a conceptual structure whose historical existence had never been independently established.

Each of these critiques—theological distortion, historical unverifiability, lack of historical specificity, categorical overreach—identified a different weakness in the inherited

framework. But they converged on a common suspicion: that the circularity identified in the preceding section had become the category's defining feature. Texts were classified as courtly because they displayed traits derived from the definition, and the recurrence of those traits was taken to demonstrate that the definition corresponded to a historical reality. What Paris and Lewis had presented as a coherent structure could instead be understood as the retrospective organization of heterogeneous literary phenomena under a modern analytical term.

Yet these calls to reject or abandon courtly love did not produce a scholarly consensus. Francis Utley (1972), Jean Frappier (1973), and William C. Calin (1980), among others, defended the concept as important—perhaps indispensable—to the interpretation of medieval literary and social ideals. For these scholars, the instability of the term did not necessarily prove that the relations it named were imaginary. Secrecy, service, hierarchy, refinement, adultery, and idealization recur across lyric and romance with a regularity that still appeared to require some form of collective explanation, even if the explanation inherited from Paris and Lewis was no longer satisfactory. The terminological problem identified by Moore (1979)—and which also provides the title of his essay—proved especially revealing in this regard. The French expression *amour courtois* was not the ordinary term through which medieval poets described what they were doing. Occitan writers employed *fin'amor*; German traditions developed the vocabulary of *Minne* and *Minnesang*.

These expressions do not map neatly onto one another, nor do they necessarily designate the complete set of characteristics later attached to courtly love. The fundamental error, in other words, was not anachronism but the assumption of shared reference. Scholars discuss courtly love, Moore observed, "as though they had the same idea at the back of their minds" (1979, 632), when in fact their definitions designate incompatible phenomena. What one critic understands by

courtly love as adulterous passion governed by a code of secrecy and service, another may understand as spiritualized devotion modelled on religious worship, a social practice shaped by feudal hierarchy, or an affective disposition analyzable through the categories of desire and sublimation. The term functions as a false cognate within the scholarly conversation itself: it creates the appearance of a shared object where none exists. Moore consequently dismissed courtly love as "the creature of every wind stirred up by scholars" and "a mischievous term which should be abandoned" (1979, 630–31; see also Boase 1977, 123, 129; Mann 2013, 92). Moore's objection concerns the instability of the term's scholarly usage. This thesis's argument, however, is that the problem is not merely terminological. The phrase gathers under one heading phenomena that may belong to different genres, languages, institutions, moral traditions, and affective systems. The adjective "courtly" privileges aristocratic and romance-centred evidence; the noun "love" implies a semantic unity—among devotion, desire, friendship, charity, marriage, and social service—that medieval texts do not consistently sustain.

Bernard O'Donoghue proposed one way of making sense of this proliferation by casting the disagreement in geographical terms: a division between an American critical tradition, which tends to regard courtly love as an exaggerated modern construction, and a European tradition more inclined to defend its historical and literary significance (2006, 7–9)—a distinction he relates to supposedly "practical" and "theoretical" tendencies within the two academic cultures (2006, 8). The idea that the debate is structured by different scholarly traditions is, I think, a useful one—more useful, in fact, than the anachronism critique on its own, because it shifts attention from whether the category is historically valid to how different communities of scholars have prioritized different explanations of the recurrent but heterogeneous field of texts, concepts, and relations gathered under the name of courtly love. But the distinction is not geographic:

affirmative and sceptical positions do not divide consistently along national lines; scholars placed on opposing sides often rely upon the same canonical texts, the same inherited vocabulary, and the same historical problems.

2.5 A different way of understanding the controversy

The preceding sections have shown that the critique of anachronism, while persuasive in its central claim, did not resolve the debate. Moore's diagnosis of false shared reference helps explain why: scholars employ the same term while assigning it different meanings, and the resulting divisions cannot be adequately captured by a geographical distinction between American and European schools. If geography does not adequately account for the structure of the disagreement, the positions must be organized according to another principle. This chapter therefore proposes a provisional classification based not on where scholars were trained, or on whether they affirm or reject the category of courtly love, but on the kinds of explanation their arguments repeatedly privilege. Organizing the debate in this way reveals three recurring explanatory tendencies. One group of interpretations foregrounds religious doctrine, heterodoxy, and moral conflict. Another gives explanatory priority to social institutions and relations, including feudal hierarchy, court culture, marriage, gender, and class. A third approaches the material principally through desire, affect, fantasy, and subject formation. For the purposes of the following analysis, these tendencies will be described as philosophico-theological, sociological, and psychological.

These labels are products of this thesis's reorganization rather than autochthonous categories within the debate itself. The scholars discussed below do not necessarily identify themselves with one of these categories, nor do their arguments always remain confined to a single explanatory framework. A scholar placed within one tradition—for example, Wilcox—can

be situated differently depending on which aspect of his argument is emphasized. The same canonical texts—Provençal lyric, Chrétien de Troyes, Andreas Capellanus, and the Tristan tradition—also recur across all three tendencies. The claim, then, is not that every medievalist belongs neatly to one of three schools, but that these recurring explanatory priorities may provide a more useful account of the debate than the alternatives considered above. The proposed distinction concerns not only explanatory emphasis but also the evidentiary bases from which those explanations are developed. The three tendencies draw upon overlapping but differently delimited evidentiary bases. Provençal troubadour lyric, Old French romance, the Tristan and Lancelot traditions, Capellanus's *De amore*, and Ovid recur throughout the debate, but they are placed within different interpretative contexts. Accounts concerned with religious meaning situate these texts alongside patristic, scholastic, heresiological, and ecclesiastical materials.

The following sections test whether this tripartite scheme clarifies the structure of the debate without reducing its participants to fixed or mutually exclusive schools. Each examines a recurring explanatory tendency, the scholars associated with it, the evidence on which it principally relies, and the dimensions of medieval love that it brings into focus or leaves comparatively peripheral.

2.5.1 The philosophico-theological stake

According to a philosophico-theological reading of courtly love, religious doctrine—its conflicts, its heterodoxies, its moral anxieties—is taken as the interpretation principle through which the literary phenomenon becomes intelligible. Even more clearly than the other stakes, this line of interpretation stems from both C. S. Lewis's work and a longer engagement with the medieval body of religious writing on the nature of love, desire, and sin. The biblical notion of love had long been a rich ground for intellectual exploration within the Christian tradition, with

figures such as Clement of Alexandria, Origen, St. Augustine, St. John Chrysostom, St. Gregory of Nyssa, and St. Maximus the Confessor, each examining different facets of love's nature and purpose—connecting it to knowledge, wisdom, grace, charity, friendship, and ultimate beatitude (e.g., see Vandevelde 2001, 23–24). Their works reflect a complex landscape of tensions: between self-love and pure love, between natural and supernatural love, between physical desire and spiritual devotion. These tensions laid the groundwork for the medieval inquiries that the theological strand of the courtly love debate would later foreground as the tradition's defining feature.

The Lewisian synthesis described in the preceding section already contained the theological argument in embryo—i.e., the claim that Provençal courtly poets, misreading Ovidian erotic literature as fundamentally serious rather than comic, produced feudalized love religion that stood in overt polemic in tension with Christian orthodoxy (Lewis [1936] 2013, 4–27). The scholars who developed this line most influentially—de Rougemont, Denomy, Russell, and, more recently, William M. Reddy—arrived at their interpretations through different philosophical and historical routes, but their evidence converges on a remarkably small cluster of canonical texts: Provençal troubadour lyric, the Tristan legend, and Andreas Capellanus' *De amore*, read against a body of patristic, scholastic, and heterodox writing. Denis De Rougemont radicalized Lewis's claim by reading the troubadour lyric and the Tristan legend through the lens of Cathar and Manichaean dualism, producing courtly love as a spiritualized passion defined by distance, suffering, and the attraction of death—a crypto-heretical formation in which individual sexual desire was always a potential threat to societal order and proper conduct (de Rougemont 1939). Alexandre Denomy moved from Avicenna and Neoplatonic philosophy of the soul through the problem of heterodox thought, arriving at Capellanus' *De amore* as evidence of an

Averroist "double truth"—the controversial philosophical stance, attributed to Averroes, which suggested that two contradictory ideas could both hold truth, one from the standpoint of reason, the other from faith—that allowed Christian restraint and worldly passion to coexist (Denomy 1965). Russell evaluated the role of Catharism further, reading Cathar documents and anti-heresy chronicles against the same *De amore* and hypothesizing that courtly love might have reflected the tensions then present in French society, possibly signifying a reaction against ecclesiastical hierarchies (Russell 1965). More recently, Reddy proposed—drawing on the same body of troubadour lyric and clerical writing—that courtly love may be understood as a form of resistance against the sweeping condemnation of sexuality that emerged with the Gregorian Reform: during the twelfth century, the Church developed extensive theological doctrines and legal structures aimed at controlling sexual behaviour and marriage, and courtly love would have provided a way to assert the primacy of love over ecclesiastical constraints (Reddy 2012, 26–27).

What unites these otherwise diverse positions is the explanatory priority they assign to the theological dimension of the evidence. Whether courtly love is read as crypto-Catharism, as Averroist double truth, as anti-clerical resistance, or as the misapplication of Ovidian convention, the interpretive move is the same: love becomes intelligible through its relationship to doctrine, heresy, orthodoxy, and the moral economy of salvation. And in each case, the interpretation is built from a restricted evidentiary base—Provençal lyric, the Tristan legend, Capellanus, and a small body of patristic and heterodox texts. When these texts are read alongside patristic, scholastic, and heterodox materials, their religious and philosophical relations become especially visible.

This priority makes visible a real dimension of medieval love literature—the persistent entanglement of erotic desire with spiritual aspiration, religious guilt, and the institutional

authority of the Church. What it tends to leave at the periphery is everything in the medieval treatment of love that is not organized around the theological problem: the social logic of feudal service and courtly conduct, the psychological dynamics of desire and fantasy, and—most consequentially for the present study—the bawdy, the appetitive, the transactional, and the comic, which occupy no place within a framework whose fundamental assumption is that love is oriented toward some form of the sacred.

2.5.2 The sociological stake

The sociological interpretation assigns explanatory priority not to religious doctrine but to social structure—to the institutions, hierarchies, and cultural transformations within which love is produced, regulated, and represented. Its archive overlaps substantially with that of the theological strand but the questions it puts to that archive are different: not what love means in relation to God or salvation, but what love does in relation to feudal obligation.

Of the three strands, the sociological has the earliest roots—and the deepest entanglement with the category's original construction. Paris's own definition, as the preceding section showed, already contained the sociological reading in germ: the four characteristics he identified are fundamentally social in nature, describing a code of conduct anchored in relations of rank, gender, and obligation. The scholars who built most consequentially on this foundation—Ker, Wilcox, Dronke, Köhler, Bloch, and Kay—drew from a body of evidence centred on French romance, Provençal lyric, chivalric culture, and a small number of Latin theoretical texts, above all Capellanus' *De amore*. Ker, writing in 1897, argued that the shift from epic to romance—moving from the *Chanson de Roland* through Chrétien, Benoît de Sainte-Maure, and Ovid's *Heroides* and *Ars amatoria*—signalled a "twelfth-century revolution" in which broader societal transformations were mirrored in literary form ([1897] 1957, 349). The rise of romance was, for

William P. Ker, emblematic of the development of Western medieval civilization itself: a move from martial heroism to personal devotion and emotional depth, in which romances like those of Chrétien cultivated the "courteous science of love" (345). Wilcox narrowed this line through the court of Marie de Champagne and the later romances of Chrétien, defining courtly love through the worship of women, doctrinaire free love, and sublimation through chivalric deeds—and suggesting that it likely originated as "an idealization of grosser social facts," refining everyday social interactions into an elevated code of conduct (Wilcox 1930, 322–23).

The tradition's strength lay in its capacity to connect literary representation to the material and political conditions of its production. But it also generated its own internal challenges. Peter Dronke challenged the sociological strand from within by arguing that the themes attributed to courtly love appear not only in medieval Western aristocratic literature but also in non-Western love poetry, suggesting that such ideals might reflect more universal human experiences than any specific socio-cultural milieu (Dronke 1968). If the same patterns appear in cultures without feudal courts, the sociological explanation cannot be sufficient. Erich Köhler, working from the troubadour lyric and reading *fin'amor* against feudal and class relations, offered a more explicitly class-based response: courtly love functioned as a unifying ideology, bridging the gap between upper and lower nobility by fostering a shared set of ideals—a form of social cohesion rather than mere literary convention (Köhler 1976). Bloch proposed that courtly love required the individual to submit to an abstract authority—love itself—in a structure that foreshadowed the modern relationship between citizen and state (1977, 215–48). Sarah Kay pushed furthest in this direction, arguing—on the basis of the same troubadour and romance archive—that courtly love eludes univocal definition precisely because it is structured around internal tensions that it simultaneously seeks to obscure: the questions of whether love is foolish

or wise, ecstatic or rational, socially beneficial or disruptive, spiritual or sensual, are evidence of existing tensions within the courtly environment—a setting marked by an intricate blend of decorum, elegance, display, and wealth, invested above all in containing potential divisions and mediating between lay and clerical interests (Kay 2000, 85–92).

The sociological strand makes visible the institutional and relational dimensions of medieval love—the ways in which desire is shaped by service, hierarchy, gender, and social aspiration. Its trajectory, from Ker's civilizational narrative through Köhler's class analysis to Kay's structural reading of courtly tensions, traces a progressive internalization of the social question: from what love reflects about society to how love is itself a social mechanism. But the archive from which these conclusions are drawn remains centred on Provençal lyric, Old French romance, and especially Capellanus. When the same literary materials are placed in relation to feudal institutions, courtly conduct, and class structure, their social functions become the principal object of explanation. What it tends to leave at the periphery is the theological intensity of the devotional tradition—the mystical and penitential treatments of love that are not easily reducible to social function—and, once again, the appetitive and comic registers that do not participate in the idealizing logic of courtly conduct. While the sociological has less to say about representations governed by private compulsion or by registers that resist courtly refinement altogether; it surely explains especially well how love mediates hierarchy, conduct, and societal tension.

2.5.3 The psychological stake

The psychological strand is distinguished from the sociological not by subject matter but by explanatory priority: where the sociological strand treats desire as a product of social structure, the psychological strand treats social structure as the vehicle through which desire

expresses itself. Its archive is, once again, substantially the same—troubadour lyric, Chrétien, Capellanus, the Tristan material—but the questions it brings to that archive concern interiority, affect, the dynamics of longing, and the construction of the beloved as an object of fantasy.

The psychological reading emerged early but found its most influential formulation late. The scholars who shaped this tradition most decisively—Wilcox, Moller, Lacan, Gaunt, Žižek, and O'Donoghue—drew from the same Provençal and romance material as the other two strands, supplemented by the Minnesang lyric and the Tristan tradition, but read that material through the categories of psychoanalysis through categories of desire, affect, and subject formation. Wilcox, already discussed in the sociological context for his emphasis on courtly conduct, also introduced one of the earliest psychological readings: the repetitive themes of courtly love as manifestations of unconsummated desire, in which the lover's devotion becomes a form of sublimation (1930, 322–23). Herbert Moller combined *Provençal* and *Minnesang* lyric with the social-demographic record of southern France to explain courtly love as frustration, sublimation, and social training: the male lover, captivated by the woman's aloofness and idealized perfection, remains devoted, performing acts of "love service" for an unattainable lady, and through this process learns to suppress coarseness, tolerate frustration, and maintain a refined disposition—biological impulses redirected into culturally valued expressions (Moller 1959, 48–49). What Wilcox and Moller shared was a fundamentally Freudian assumption: that the literary representation of love could be read as evidence of psychic processes operating beneath the social surface. Their disagreement concerned which processes—regression, sublimation, socialization—but not whether psychoanalytic categories were the appropriate lens.

The Lacanian school radicalized this line by shifting the question from the lover's psychic development to the structural logic of desire itself. For Jacques Lacan, courtly love—

read through the troubadour lyric of the twelfth century—was "a poetic exercise," an activity of poetic creation that significantly influenced social behaviour, in which the idealized beloved functioned not as a social person but as an object produced by desire, and in which historical variation was secondary to the transhistorical logic of that structure ([1959–60] 1986, 148–49). Gaunt argued, on the basis of the same troubadour corpus, that the love song is inherently a "monologue" whose absence of true dialectic between poet and lady highlights "the phallocentric nature" of courtly love lyrics—the poet engages not with real women but with an idealized creation upon which he exerts control (1995, 142–43). Slavoj Žižek pushed further, arguing that courtly love is not sublimation but perversion: the lady's "arbitrary dictatorship" reduces the lover to a subservient role in which suffering becomes the ultimate expression of desire (2005, 89–94). O'Donoghue identified "the central plank in the construction of the courtly lover's situation" as the individual's drive to pursue desire even when it defies social conventions—a "transgressive but individually imperative passion" reflecting the Freudian concept of desire as developed by Lacan (2006, 11–12).

The psychological stake, particularly in its Lacanian form, thus reverses the direction of the sociological explanation. Where sociological accounts treat desire as historically shaped by social organization, Lacanian accounts treat social convention as one of the forms through which a more general structure of desire becomes visible. This opposition should not, however, be made absolute. Both traditions reason from overlapping evidentiary bases and distinguish themselves principally through the explanatory priority they assign to the relations found within them: social organization in one case, psychic structure in the other. Within the psychological tradition, troubadour lyric, Old French romance, the Tristan material, and a small number of Latin theoretical texts are organized around distance, prohibition, suffering, deferred fulfilment,

and the inaccessible beloved. This arrangement makes the dynamics of fantasy and desire especially visible, but it also risks treating historical and social variation as secondary to a transhistorical psychic structure. A differently delimited evidentiary base—including *fabliaux*, devotional prose, penitential writing, or representations of marital and transactional love—might foreground affective relations less readily accommodated by the Lacanian model. It is precisely this tendency to detach desire from its historically specific forms that recent work in the history of emotions challenges.

More recently, Don A. Monson has reformulated the psychological interpretation through cognitive science and evolutionary psychology, explicitly rejecting the Lacanian framework influential in courtly-love criticism since the late twentieth century. For Monson, Lacan's theories "cannot serve as the basis for understanding medieval love psychology" because they are "predicated on a false theory of human nature" (2022, 37). His alternative draws on cognitive and evolutionary research to interpret courtly themes as culturally specific elaborations of biologically universal patterns (43–53). Although this framework abandons psychoanalysis, it does not abandon the psychological tradition's explanatory priority: love remains intelligible primarily through the structure of the mind, even if that structure is now evolutionary-cognitive rather than psychoanalytic. Monson's corpus is also generically broader than those of many earlier psychological accounts, incorporating lyric, narrative, didactic, and satirical material; it nevertheless remains concentrated on twelfth- and early-thirteenth-century France and on Old Occitan, Old French, and Medieval Latin sources (9–10). It therefore confirms the archival pattern traced in this chapter: even an account that deliberately broadens the generic range of its evidence continues to reason primarily from the French, Occitan, and Latin materials that Paris and Lewis helped establish as the debate's canonical archive.

The history of emotions therefore represents both the most recent development of the psychological strand and a partial correction of its premises. It challenges Lacanian assumptions by questioning whether desire can be treated as a universal psychic structure operating independently of historical context. The methodological foundation for this historicizing challenge was established by Peter N. Stearns and Carol Z. Stearns, who introduced the term "emotionology" to describe the attitudes and standards through which a society, or a definable group within it, evaluates emotions and regulates their appropriate expression (Stearns and Stearns 1985). Their distinction between emotional experience and historically variable emotional standards redirected attention from supposedly universal psychic mechanisms toward the social norms governing how emotions are understood, expressed, and judged. Barbara H. Rosenwein identifies a dominant scholarly narrative in which medieval emotional life is cast as excessive, childlike, and insufficiently regulated—a view shaped by the combined influence of Huizinga and Freudian social psychology (Rosenwein 2002). Psychological readings of courtly love that interpret longing as sublimation, regression, or fetishization risk reproducing this framework by projecting modern categories of psychic development onto medieval texts. Rosenwein argues instead that emotions must be understood as historically constituted practices embedded in social, linguistic, and institutional systems. Her concept of "emotional communities"—groups that share norms governing the evaluation, expression, and meaning of feeling—relocates emotions from the realm of universal psychology to that of social practice (Rosenwein 2002). Damien Boquet and Piroska Nagy extend this position by foregrounding the political and socially embedded character of medieval emotions. Courtly literature, they argue, was not meant "to resound quietly within the heart" but to be sung publicly, in settings where emotional bonds had "to be seen and felt in order to be recognized," giving them real social

efficacy (Boquet and Nagy 2018, 105–6). Beyond defining courtly love as an "erotica of mastering desire," they show that it articulated a shared affective logic, intelligible within a courtly emotional community whose members recognized, evaluated, and rewarded specific forms of emotional restraint, endurance, and display through established literary and moral codes (110).

The history of emotions therefore emerges as the most recent phase of the psychological strand—and also, in a sense, as its self-correction. It inherits the strand's questions about affect and interiority but repositions them within a historically grounded analysis that reconnects spiritual frameworks, social structures, and subjective experience without collapsing any of them into the others. Its intervention is significant for the present study because it demonstrates, from within the psychological tradition itself, the limitation that this chapter's argument has been tracking across all three strands: the tendency to treat one explanatory dimension as foundational while the others remain peripheral.

2.6 Love as discourse

The three strands examined above provide a more detailed and productive analysis of the different schools in the historical scholarly love debate. But it does not help these schools resolve their differences. Each identifies a real dimension of the medieval treatment of love but each treats its own dimension as foundational while the others remain peripheral. The impasse between the stakes is therefore a structural consequence of how the scholarly object has been constituted. The three stakes arise from overlapping but restricted archives, each producing a version of medieval love whose internal coherence derives from the evidentiary base through which it was constituted. This thesis proposes that the impasse can be reframed by changing the status of the object under discussion. Rather than asking which restricted account supplies the

correct definition of courtly love, it asks what becomes visible when those accounts are treated as partial descriptions of a larger discursive field. As long as courtly love is treated as a doctrine, a code, or a singular historical phenomenon, the debate will remain organized around the question of which account is correct—and each tradition will continue to confirm itself from its own archive. But if the contradictions among the three strands are treated not as evidence of scholarly failure but as evidence of the object's own discursive character, the question changes. It is no longer which tradition identifies the true meaning of love, but how the philosophico-theological, sociological, and psychological dimensions each foreground coexist, cross, and contest one another within a shared discursive field of textual production. On this account, the medieval love discourse is not a single meaning that scholars must recover but an ongoing negotiation whose structure is inscribed in the language of the texts themselves. If this is right, then the impasse among the three strands is not a problem to be solved by choosing one tradition over the others. It is a consequence of each tradition having recovered one region of a larger discursive field—a field that, in principle, can be reconstructed from a corpus broad enough to span the full range of the literary production rather than from any single strand's preferred texts. The question that follows is whether the existing scholarship already contains evidence for this reframing—which it does.

Attempts to preserve the category by allowing for multiple forms of courtly love did not eliminate this problem of shared reference. Utley, for example, defines courtly love as "the formalized codification of love" while conceding that the category may contain "twenty or thirty" variations (1972). Mann similarly observes that it is often difficult to determine what scholars mean by courtly love because the term is treated as though its meaning were self-evident and is consequently left undefined (2013, 92). The instability operates at two related

levels: medieval texts employ diverse and sometimes incompatible vocabularies for love, while modern scholars organize those vocabularies through categories that are themselves inconsistently defined. Since medieval representations range from spiritual and virtuous love to erotic, profane, and transgressive desire, neither medieval authors nor modern critics can be assumed to occupy a single conceptual ground. The anachronism of courtly love therefore arises not only from the retrospective application of a modern term, but from the assumption that the term establishes a shared object across medieval texts and modern interpretations.

Patricia Kean, writing in 1972, had introduced quite early a notion that anticipates the discursive reframing from a different direction entirely. Rather than invoking discourse theory, she identified what she called a "common narrative *koinê*"—adopting the term referring to the common language existing to communicate among different ancient Greek dialects to indicate a shared stylistic, narrative, and thematic elements connecting Chaucer's work to a larger tradition of romance literature (1972, 6). The concept of a *koinê* points at precisely the kind of patterned recurrence across texts that linguistic discourse analysis is designed to detect: not a fixed system of beliefs but a set of recurring configurations—stylistic, thematic, and lexical— the distribution of which across a corpus provides one observable point of access to a discourse. Kean did not pursue this line into formal discourse analysis, but the intuition is the same: if love in medieval literature is an ongoing debate rather than a settled doctrine, its structure should be visible in the patterns of language, form, and argument that recur across texts—and those patterns are, in principle, recoverable.

Several scholars had already gestured, in different ways, toward a conception of love that exceeded the explanatory framework their own tradition supplied, and without the formalization Kean's *koinê* anticipates. Similarly, Nancy Thompson explores the idea of a love debate in

Chaucer in relation to broader literary traditions of disputation about love (Thompson 1996). Kay—discussed above as a sociological critic—argued that courtly love is "debated across a large number of texts" (2000, 81), a formulation that treats the phenomenon as an ongoing argument rather than a stable doctrine. Something similar to what Paxson and Gravlee, working under the heading of "discourse of desire," went further in foregrounding as the discursive character of the phenomenon:

> People in a mimetic narrative produce discourse; yet discourse, like a sentient agent, generates a semiotic or rhetorical field wherein human agents are produced, constructed, arranged. Discourse determines and disposes the social world and formal structure of the literary text invested in desire and love. (1998, 15)

One further connection characterizing love in linguistic terms deserves notice. O'Donoghue—discussed above as a psychological critic—similarly insisted on the presence of a "body of literature which can be conveniently taken together with common conceptual elements in it" (2006, 13). This focus on the corpus draws attention to the pan-European circulation of love-situations and language across vernacular lyrics, and, in particular, on role of Anglo-Norman bilingualism in facilitating lexical connections that in German contexts required translation (2006, 16–21; see Jaeger 1985, 127–35). Stephen C. Jaeger himself described the result as "a large grid crisscrossed with conflicting opinions" (1999, 120)—a formulation that points toward a textual nature of discourse without yet naming it as such. McCarthy, in his anthology of legal, medical, devotional, and literary multilingual sources from the fourth to the fifteenth century, presenting a "visible literary dialogue about the nature of love taking place across a range of medieval texts, particularly lyric and romance, from the troubadours onwards" (2022, 7, 25).

At different levels of understanding, what these scholars share—despite belonging to different strands of the debate and working from different theoretical commitments—is the

recognition that medieval love is not a fixed entity whose meaning must be recovered but an ongoing negotiation whose structure is visible in the patterns of recurrence, tension, and contestation distributed across a body of texts. It was Schnell who gave this recognition its most decisive formulation—and who did so independently of Kean. In his 1989 study, he drew upon the terminological and definitional problems identified by Utley (1972) and Moore (1979), changing the status of the contradictions that had frustrated earlier scholarship. In fact, if for Utley and Moore the proliferation of definitions attached to the label *courtly love* exposed scholarship's failure to establish a stable and commonly shared object, Schnell, by contrast, sees that same plurality not as evidence of a failed synthesis, but as a constitutive feature of an ongoing discourse about amorous conduct.

> [T]he vague and inconsistent terminology used to depict romantic relationships in vernacular poetic works suggests a decisive element: we are not dealing with a clearly defined theory of love, but rather with a courteous discussion on courtly forms of amorous conduct. Yet, different perspectives among most poets converge with the aim of portraying exemplary behaviour. As will be demonstrated, the literary phenomenon of courtly love is better understood as a discourse on true, virtuous love, rather than as the reproduction of a fixed, clearly established convention of love. (1989, 78–9)

According to Schnell's definition, love's ambiguity and incompatible features—the contradictions that had sustained and frustrated the debate since Paris—point to its inherent polemic nature. They are constitutive features of a discursive field in which religious, social, affective, moral, and erotic meanings are repeatedly articulated in relation to one another. The terminological instability that Moore diagnosed as a problem becomes, on Schnell's account, evidence of precisely the kind of ongoing contestation that discourse analysis is designed to recover.

The growing influence of this reframing became visible in 2003, when the first International Symposium on Medieval and Early Modern Studies, held at the University of

Arizona, brought together scholars to examine love, marriage, and transgression across a wide range of historical, linguistic, and generic contexts. The publication of its proceedings the following year explicitly foregrounded these themes as discursive categories rather than stable concepts (Classen 2004). To my knowledge, Albrecht Classen is the scholar who pushed furthest toward operationalizing the discursive reframing (2002; 2008) conceptualizes as the idea of *kommunikative Gemeinschaft* (i.e., a "discourse community"), not as a group defined by shared beliefs or fixed doctrines, but as one constituted through participation in communicative processes themselves. In this sense, membership is not determined by adherence to a stable ideology, but by the capacity to engage in, respond to, and reproduce the discursive practices that sustain the community over time. Classen explicitly situates his interpretative model within the work of the sociologists Jürgen Habermas and Niklas Luhmann, to argue that the central problem in (German) medieval literature is discourse itself (*Kommunikation* in the original; 2002, xviii). Classen explicitly frames this as the foundation of social existence, asking to what extent "functioning communication, or communicative action, contributes to the formation of society" (xxiii). In this sense, Habermas' notion of discourse—where social cohesion depends on mutual understanding achieved through language—provides a model for interpreting literary interactions as attempts—often unsuccessful—to establish shared meaning; while Luhmann's systems theory allows us to conceive of discourse as a relational network in which meaning emerges through difference and interaction rather than as a fixed entity (xxii). Despite this quite divergent treatment, with respect to other scholars, what matters is that Classen explicitly pushed the boundaries of literary debate towards positions integrating discourse theories into existing frameworks by affirming how much literary texts are always "embedded in their respective society in socio-historical and discursive terms" (Classen 2002, xviii). Thus, authors and texts

become part of this Gemeinschaft insofar as they participate in the ongoing negotiation of meaning—whether successfully or through communicative failure—thereby contributing to the maintenance or disruption of the social fabric articulated through discourse.

Classen's contribution, however, does not stop there, but extends to the application of this interdisciplinary method to the understanding of love. In Classen's view, in fact, medieval literature can be understood as staging a form of discourse on discourse, in which communication itself becomes both the medium and the object of reflection (Classen 2008). Classen affirms that medieval literature can be understood as being embedded within a "larger discourse community" that actively seeks to establish "rational principles, logical strategies, and communicative operatives" for the maintenance and improvement of society (4–5). The aim is "to uncover the true nature of love poetry and love discourse as the critical vehicle for all human efforts in the interpretation of our physical and metaphysical environment" (5–6). What is at stake, therefore, is not only the representation of love, but the capacity of language to function as a reliable medium of sociocultural exploration, so that the prominence of love discourse must be understood in this light: it constitutes a privileged domain in which the limits, possibilities, and failures of communication are explored through the ambivalent tensions of the discourse on love itself.

If, as Classen suggests, medieval literature continuously negotiates meaning through discourse rather than presupposing it, then the question becomes methodological: how can such a discourse be identified, described, and analysed without reifying it into a fixed system? It is precisely at this juncture that returning to Kean's notion of a *koinê* becomes essential. Her emphasis on recurrent textual features offers a way of connecting the abstract idea of a discourse community to relations that can be observed within texts. These patterns represent not only the

constitutive elements of medieval literary discourse on love—emerging directly from the texts—but also the observable traces of what might otherwise remain an abstract discourse community. While modern scholarship has questioned the historical reality or theoretical coherence of courtly love as an established system for medieval audiences, discourse analysis shifts the inquiry: it invites us to examine the textual patterns and clusters that emerge naturally, fostering a fresh understanding of medieval love and its literary expression. The goal is not to disprove existing assumptions but to complicate established claims, starting at the periphery of scholarly interpretations. By computationally analysing a corpus of texts representing more or less clear interventions to the debate, this study pivots from the inductive interpretations of traditional scholarship to a more deductive, data-driven approach, as elaborated in the following chapters. This approach builds an understanding of the genre from the "inside out," resisting assumptions of discourse without first identifying the nodes of the debate. Without this, the relationships within medieval love literature would remain obscured, leaving us with a reconstruction that risks circling rather than advancing our understanding of love's representation in the Middle Ages.

2.7 Next steps

The question that follows on the above discussion is methodological. If courtly love is a discourse—an ongoing, contested, and historically distributed negotiation about the meaning of love—then the method required to examine it must be one that can represent relations distributed across a corpus while making explicit the status of the categories through which those relations are examined. The distinction matters for the analysis that follows: one stage examines distributional relations without beginning from the three scholarly stakes as predetermined classes, while another deliberately introduces Lewis's taxonomy as an inherited benchmark to be

tested against document-level organization. But the term discourse analysis itself carries a burden of confusion that must first be addressed. The following chapter therefore clarifies what is meant by discourse analysis, explains the persistent ambiguities surrounding the term, and defines the form adopted by this thesis. The later computational analysis will place these stakes within a common representational field, examining how their associated meanings are distributed, connected, and separated across an explicitly delimited Middle English corpus. Its purpose is neither to replace the earlier scholarship nor to adjudicate among its traditions as competing hypotheses, but to make their partial accounts mutually comparable and to identify relations that their separate explanatory priorities may have left comparatively peripheral.

CHAPTER 3: FRAMING THE METHODS

> Je n'écris pas un livre pour qu'il soit le dernier;
> j'écris un livre pour que d'autres livres soient
> possibles, pas forcément écrits par moi.
>
> —Michel Foucault, "Entretien avec Foucault"

3.1 What is discourse, really?

Discourse and discourse analysis are terms representing divergent theories, methods, and approaches. From their early uptake across the Humanities and Social Sciences, critics have noted the term's conceptual looseness. In cultural studies, Julian Henriques and colleagues famously dismissed the proliferation of discourse talk as mere "discourse babble" (1984, 105). Similar reservations emerged in anthropology, where Lila Abu-Lughod and Catherine A. Lutz observed that "'discourse' has become [...] one of the most popular and least defined terms in the vocabulary of Anglo-American academics […] defining discourse precisely is impossible because of the wide variety of ways it is used" (1990, 7). Some scholars went further, treating discourse analysis as a scientific fad—one that would disappear as swiftly as it emerged. Henry G. Widdowson (1995, 158) also describes the term discourse "as vague as it is fashionable," adding that this is a concept "everybody is talking about but without knowing with any certainty just what it is: in vogue and vague." This view characterizes a perspective on discourse, shared by its detractors, that sees it—*horribile dictu*—as an "inflationary buzzword" (Kerchner and Schneider 2006, 6) to join the latest academic trend with the highest possible chance of attracting attention; or further, "a trendy concept [...] whose arbitrary use in the most diverse fields makes a general understanding seem impossible" (Ruoff 2007, 92). Moreover, concerns have been articulated within the broader field of linguistics itself, where the concept of discourse has been described as inherently vague, ambiguous, and associated with sometimes incompatible theoretical traditions (e.g., see Keller et al. 2006). Such critiques are closely tied to the persistent

difficulty of determining what discourse analysis—which at times seems to resist stable disciplinary classification and has been described as a "methodology," a "method," a "theory," an "attitude," or even an "art" (e.g., see Gardt 2007, 28–38; Niehr 2014, 44–50)—actually is. Spitzmüller and Warnke (2011) argue that this conceptual indeterminacy is not unique to discourse analysis nor necessarily a sign of analytical weakness; rather, it reflects the historically layered development of the field and the absence of a single, fully systematized theoretical foundation.

But the critics quoted above are not complaining about definitional looseness: several of them deny that discourse analysis constitutes a coherent intellectual enterprise at all. Against that scepticism, the record of the field's applications tells a different story. Discourse analysis has proven to be among the most productive frameworks for exploring the entanglement of power, identity, social practices, and language. Its versatility is evident in a range of research areas, including more traditional fields such as literature (e.g., see Johnstone 2005; Maingueneau 2010), politics (e.g., see Wodak 2009; Van Dijk 2008), and techno-cultural studies (e.g., see Akkaya 2014; Roderick 2018; Brock 2018; Oprea 2019). It has also been effectively combined with critical theories, such as feminism (e.g., see Lazar 2005; Wagner and Wodak 2006; Baxter 2008), post-colonialism (e.g., see Jackson 2015; Guilherme 2019; De Melo Resende 2021), and environmental studies (e.g., see Stibbe 2014; Fernández-Vázquez and Sancho-Rodríguez 2020). This breadth of application reflects not a unified theoretical foundation but the historical dissemination of discourse analysis across multiple disciplines, each of which has adapted the concept to its own analytical concerns.

3.2 Discourse après Foucault

In nearly every account of discourse analysis, the name that emerges—more or less directly—is that of the French philosopher Michel Foucault. But while Foucauldian concepts have become highly influential across disciplines (e.g., Mills 2004), Foucault did not provide a unified theory or systematic procedure for discourse analysis.

What is now commonly termed *Foucauldian* discourse analysis is therefore best understood as the result of a post-Foucauldian line of reception in which different scholars adapted selected elements of his work—historicity, power, knowledge, subject positions, exclusion, and discursive formation—to distinct disciplinary and methodological purposes. This post-Foucauldian expansion is crucial for the argument of this thesis because it moves discourse beyond the narrow sense of turn-taking in conversation. Discourse comes to name the broader circumstances in which conversations take place: relations of power, social positions, institutional constraints, gendered and racialized standpoints, epistemic assumptions, and rules of formation that make some statements possible, authoritative, marginal, or excluded. The diversity of contemporary discourse-analytical approaches stems partly from repeated attempts to operationalize the archaeological programme that Foucault himself never conceived as a single method and never fully systematized as a transferable procedure. His work, by his own admission, is defined more as "*boîtes à outils*" (1994, 720)—that is, as toolboxes offering an open-ended programme rather than a rigid set of procedures for discourse analysis.

This is also why the post-Foucauldian tradition provides the conceptual base for constructing a heterogeneous archive. A discourse such as courtly love is not confined to a single text, genre, school, or scholarly tradition; it is constituted across multiple conversations and bodies of evidence. A romance, a devotional text, a didactic work, and a fabliau may therefore enter the same analytical field without being treated as instances of the same genre or as

expressions of the same beliefs. Their relevance derives from the different positions they occupy within the discourse of love and from the relations that become visible when they are examined together. The theological, sociological, and psychological explanations of love identified earlier may each be understood as partial archives, or historically organized conversations, in their own right. The post-Foucauldian understanding of discourse makes it possible to bring these separate archives into relation without assuming that they are identical, continuous, or reducible to one another. What unites them is not a single doctrine of love but a field of historically conditioned relations among statements, concepts, subject positions, interpretive authorities, and forms of knowledge.

However, Foucault also famously describes knowledge as "the space in which the subject may take up a position and speak of the objects with which he deals in his discourse" (Foucault 1972, 182; he also mentions the "trajectory of meaning", 150; and the "space of dissensions", 152), thereby situating discourse within a spatial framework. But as Lefebvre observed, Foucault "never explains what space it is that he is referring to, nor how it bridges the gap between the theoretical (epistemological) realm and the practical one, between mental and social" (1991, 14). The metaphor of space remains ungrounded—a suggestive figure without an operational method for recovering the structure it names.

If corpora provide the epistemological means of accessing discourse—the material from which its patterns can be recovered—the question that follows is methodological: what kind of analysis can reconstruct discursive structure from a body of texts without presupposing the categories it is looking for? The answer this thesis adopts originates in a largely independent programme within linguistics itself—one that predates Foucault's *Archaeology* by nearly two decades and that arrived at a compatible but more tractable conception of discourse through a

different route entirely: this is the distributional hypothesis by Zellig Harris. Harris coined the term *discourse analysis*, not to describe the analysis of power or historically authorized knowledge, but to extend linguistic analysis beyond the sentence through the study of recurrent distributions in "connected speech (or writing)" (1952, 1–2). Whereas the post-Foucauldian tradition asks how historically situated conditions regulate statements, subjects, and forms of knowledge, Harris asks how discourse can be reconstructed from observable patterns of recurrence, co-occurrence, and linguistic distribution within textual corpora. His programme is rarely placed at the centre of histories of discourse analysis, despite providing the distributional principle from which modern computational semantics and vector-space modelling conceptually descend: linguistic elements that occur in similar environments acquire related functions and, by extension, comparable semantic representations.

The chapter does not adopt either genealogy in isolation; instead, it draws upon their convergence in DIMEAN. This multi-level interpretative framework offers a way to hold together the top-down and bottom-up traditions of discourse analysis, allowing both movements without prescribing any. It supplies the framework within which the patterns these traditions identify can be interpreted as components of a discourse. The post-Foucauldian tradition licenses why a discourse may be reconstructed from heterogeneous materials, while the linguistic tradition explains how recurrent textual patterns within those materials can be identified and modelled. Clarifying the two routes into discourse analysis therefore establishes the scope and mandate of discourse analysis as it is employed in Chapter 5, while also preparing the discussion of the archive in Chapter 4: the archive is not a neutral container of sources but a constructed field in which heterogeneous conversations can be brought into relation through explicit criteria of selection, modelling, and interpretation.

3.3 The (post-)Foucauldian school

Foucault's intellectual endeavour is commonly situated within the broader structuralist agenda concerned with contingency, historicity, and the limits of interpretation. This trajectory unfolds across three main phases: archaeological, genealogical, and ethical (Kögler 2004, 9–10). Each phase, in fact, corresponds in quick succession to a distinct area of investigation centred on the concepts of knowledge, power, and the subject. The archaeological phase is represented by *L'archéologie du savoir* (Foucault 1969), which outlines his early formulations up to the 1970s. The genealogical phase follows with *Surveiller et punir* (Foucault 1975; Foucault 1977), focusing on power relations within society and synthesizing his work throughout that decade. Finally, from the late 1970s until his death, the ethical phase is captured in the three-volume *L'histoire de la sexualité* (Foucault 1976; 1984a; 1984b), wherein the more mature developments of his thought take shape.

The precursors of what would later be recognized as the foundation of discourse analysis can be discerned in Foucault's early literary-critical work. As Philipp Sarasin observes, the genesis of discourse analysis is already visible in Foucault's 1963 study of Raymond Roussel, where he examines literary language through the identification of "two opposing propositions" constitutive of the text, one of which remains unspoken yet nonetheless determines "the logic of the story" (Sarasin 2005, 41, 64–65). Foucault's early sensitivity to linguistic and stylistic analysis is further evidenced by his engagement with Leo Spitzer, most notably through his translation of *Linguistics and Literary History*. The precise date of this translation remains contested according to Stuart Elden (2023). While the Gallimard edition is clearly dated to 1970, several early bibliographies attribute the translation to 1962—a dating that would align more closely with Foucault's intellectual trajectory around the time of his thesis work (see Elden 2023, n. 21). Regardless of its exact timing, Foucault's engagement with Spitzer situates him within a

tradition attentive to the systematic analysis of language—amongst other elements of historical development—providing an additional, often overlooked context for his later reflections on discourse. This attention to implicit regularities, structural constraints, and conditions of intelligibility anticipates later concerns with discourse, even though it remains embedded in literary analysis rather than articulated as a general analytical framework.

It is only with the 1970 inaugural lecture at the Collège de France, *"L'Ordre du discours"* (Foucault 1981, 51–78), that Foucault explicitly foregrounds discourse as a central analytical category. Here, discourse is no longer approached primarily through literary structure or stylistic analysis but is articulated in relation to institutional constraints, exclusionary procedures, and power relations. This lecture inaugurates a phase of Foucault's work in which discourse becomes inseparable from questions of power, knowledge, and social regulation, marking a decisive shift toward the analysis of discursive formations and their mechanisms of control. In fact, if Michel Foucault is to be associated with a theory of discourse at all, this association must be grounded in *The Archaeology* and understood within its explicitly limited scope.

In Sawyer's reconstruction, this work represents the only moment in Foucault's *oeuvre* in which *discourse* is employed as a formal analytic category (Sawyer 2002, 436–437). Crucially, Foucault's use of the term in *The Archaeology* remains consistent with its established meaning in French structuralists—particularly in the work of Émile Benveniste—where discourse denotes concrete events of language use rather than abstract social forces or a general explanatory principle (Sawyer 2002, 436). Throughout the book, discourse is defined neither independently nor foundationally, but relationally and derivatively, through its connection to adjacent concepts developed within the archaeological project. As Sawyer emphasizes, Foucault understands

discourse as bounded ensembles of statements that exhibit regularities of use, organization, and transformation over time (Sawyer 2002, 436–437). These ensembles are not unified by authorship, intention, or thematic coherence, but by historically specific rules governing what can be said and how it can be said. Within this framework, discursive formations function as the primary units of analysis. When such formations achieve autonomy and temporal coherence, they cross what Foucault calls the threshold of positivity, marking the loosest and least regulated form of discursive unity (Sawyer 2002, 437). The *archive*, in turn, designates the historically specific system governing the formation and transformation of statements and thus constrains what can be thought and articulated within a given period. As Sawyer notes, this concept displaces the earlier notion of the *épistémè*, allowing Foucault to emphasize historical transformation (Sawyer 2002, 437).

After *The Archaeology of Knowledge*, Foucault no longer employs discourse as a technical analytic category, and his subsequent avoidance of the term is striking (Sawyer 2002, 436–437). Sawyer explains this shift by situating it within the intellectual and political context of the late 1960s. In the wake of May 1968, *The Archaeology* was sharply criticized—particularly by Althusserian Marxists—for bracketing material, institutional, and ideological conditions and for excluding non-discursive practices from analysis (Sawyer 2002, 440–442). In response to these critiques, Foucault decisively redirected his attention toward power relations, institutional mechanisms, and non-discursive practices—inaugurating his genealogic phase. As Sawyer (2002) argues throughout his paper, however, post-Foucauldian scholarship has frequently misunderstood the change in Foucault's research interests as evidence of a single, internally coherent intellectual endeavour, rather than as a sequence of analytically distinct projects. In this way, by treating *The Archaeology* as the foundation of a general theory of discourse—rather than

as a methodologically constrained intervention tied to a specific moment—later scholars abstracted discourse from the conditions that originally limited its analytic function to create something else. Sawyer defines this process as a retrospective broadening of the original concept into a general theory of discourse that was never intentionally considered by Foucault himself (Sawyer 2002, 434). In order to justify the expansion of discourse, many later attributions rely on selective quotation, most often from a single passage of *The Archaeology*, while ignoring the surrounding context and Foucault's explicit definitions later in the book (Sawyer 2002, 451–452). As a result, secondary literature routinely conflates discourse with other concepts such as *épistémè*, discursive formation, discursive practice, or dispositif, despite Foucault's insistence that these are distinct analytical terms (Sawyer 2002, 446–448).

Sawyer's account shows that the widespread attribution of a general theory of discourse to Foucault is best understood as the product of retrospective convergence rather than original intention. The general, power-oriented concept—and thus theory—of discourse analysis that dominates much contemporary theory did not emerge from Foucault's intentional effort, but by later Althusserian- and Lacanian-inspired scholarship. In fact, the two main ideas associated with Foucault come from other sources: while his Marxism *à la* Althusser had provided the basis for ideological analysis of power mechanisms, Lacan had offered a viable framework for the socially constituted subject within these mechanisms. Despite the evident misreading of primary sources, however, the conflation was also encouraged by Foucault's rising popularity in the Anglo-American academia during the 1980s as well as the practical appeal of the term discourse, which proved capable of replacing the increasingly contested concept of culture by absorbing explanatory roles once assigned to adjacent concepts (Sawyer 2002, 448–451).

This misunderstanding is well encapsulated by Constanze Spieß' acknowledgment of the "changes and discontinuities with regard to [Foucault's] conception of discourse and, thus, of thought" (Spieß 2011, 78; see also Bendel Larcher 2015, 19). Spitzmüller and Warnke similarly identify a growing association between Foucault's analyses of discourse and questions of power over the course of the 1970s; while Axel Honneth extends this timeline, arguing that Foucault consistently explored the interrelations among power, knowledge, and discourse throughout his entire life and career (Honneth 2003, 21). What Sawyer's reconstruction makes clear, however, is that this association reflects a later phase of Foucault's work and its reception, and that it is often retrospectively imposed on his earlier archaeological writings, thereby reinforcing precisely the interpretive superimpositions that contribute to the conceptual instability of discourse analysis. The expanded theory that resulted from this retrospective broadening treats discourse as a vehicle for power relations—a well-regulated system responsible for the formation of social knowledge, namely "the whole framework of knowledge" through which interpretation is possible (Foucault 1981, 53). Within this perspective, discourse shapes orders of thought, social knowledge systems, and the thought structures of individual epochs (Spieß 2011, 79), and power structures control discourse through three mechanisms: "procedures of exclusion" (1981, 52), "procedures of controlling and delimiting" (1981, 56), and a "system of subjection" (1981, 64). The first concerns the opposition between "the true and the false" (1981, 54) along with what Foucault defines as "forbidden speech[es]" (1981, 55). The second refers to the principles of authorship, commentary, and disciplines (1981, 56–8), which maintain coherence by defining the boundaries of a "society of discourse" (1981, 63)—a concept closely related to what later scholarship would formalize as *discourse community* (see section below). For later scholars, this line of thought culminated in the concept of the dispositive (or apparatus, as it is often rendered

in English), designating “the system of relations that can be established between these elements” (Foucault 1980, 194). Through it, Foucault implies that no form of knowledge—including law, philosophy, morals, science, and even architecture—can exist without presupposing an overarching power relationship (1980, 194–228).

Another consequence of the widespread assumption that Foucault provided a general theory of discourse was a fundamental debate over the purpose and mandate of discourse analysis—one that affected its scholarly legitimacy directly. While scholars broadly agree that Foucault’s engagement with discourse was schematic rather than systematic, and that it never crystallized into a unified theoretical framework, this very openness generated a theoretical gap that post-Foucauldian scholarship sought to address in order to operationalize the approach. Rather than remaining a productive indeterminacy, the gap proliferated into competing interpretations of what discourse analysis is meant to do in practice. As a result, discourse analysis came to be pulled in divergent directions, with approaches often operating under the same label while pursuing markedly different analytical goals. Spitzmüller and Warnke characterize the resulting situation as one of “discourse bewilderment” (*Diskursverwirrung*), a condition in which the need to resolve conceptual ambiguities left open by Foucault became a central task for post-Foucauldian scholarship (Spitzmüller and Warnke 2011, 65).

3.4 Discourse as language

The preceding sections have argued that the widespread attribution of a general theory of discourse to Foucault is best understood as a retrospective construction—one that expanded a restricted and methodologically bounded concept into an all-purpose explanatory category, and that this expansion produced the “bewilderment”: a field split between competing interpretations of what discourse analysis is meant to do, with no shared agreement on its objects, methods, or

analytical scope. But the bewilderment forced scholars to ask a question that Foucault's own work had left largely unaddressed: if discourse is an order of knowledge, how is that order materially realized—and through what observable evidence can it be recovered?

The response that matters most for this thesis began with a decisive reorientation of what "discourse analysis" was taken to mean. Where the broader post-Foucauldian tradition had treated discourse as a theory of power, knowledge, and social regulation, a more linguistically oriented strand of that same tradition recovered from Foucault's work something more specific and more tractable: the analysis of the functional elements of the epistemic system—statements—along with the relationships among them. This reorientation shifted attention from discourse as an abstract order of knowledge to the linguistic forms through which that order becomes materially observable. The question became: how are statements realized in language, and how can they therefore be identified, compared, and analysed?

Although Foucault did not provide a procedure, he supplied a basic nomenclature for this linguistic approach to discourse. The French philosopher defined discourse as "the general domain of all statements, sometimes as an individualizable group of statements, and sometimes as a regulated practice that accounts for a number of statements" (Foucault 1972, 80). With this, he means a system loosely modelled on the example of Saussurian structuralism, formed by the relationship between the individual constitutive elements of a system and the system itself. These constitutive elements are the *énoncés* (statements) and the *énonciations* (enunciations), which represent a certain discourse at different degrees (Warnke 2007, 8). They are distinguished from each other because statements are functions defining conditions under which certain expressions are possible within a discourse, while enunciations are the linguistic materializations of said conditions. Statements can be produced at a given moment in specific social and cultural

contexts, whereas enunciations are realized sequences—either in writing or orally—in relation to a certain statement. The linguistic encodings of statements are the starting point from which the discourse analyst can begin to reconstruct the discourse from a linguistic perspective (2007, 8).

This approach to discourse considers not the single element—the individual *énoncé*—alone as the object of research interest, but rather always in relation to other elements, in its position within systems of formation, webs of rules, regulatory mechanisms, and the conditions constituting discourses. Both statements and enunciations represent discourses at varying levels; however, enunciations serve as the linguistic encoding of statements and thus act as the primary "units of analysis" in discourse linguistics (Wodak and Meyer 2015, 20). This implies that language itself provides the starting point for discourse analysis from a linguistic perspective. Such an approach is focused on the study of social processes in different analytical categories, namely "genre, discourse and style," to explore how semiosis mediates other elements of social practices within a specific network called "*orders of discourse*" (2016, 88–9).

These orders operate across multiple scales, linking local, situated interactions to broader institutional and societal formations. At a more delimited scale, discourse can be analyzed through the notion of *discourse communities*. John Swales offers a useful formalization of this concept. In his account, a discourse community is defined by six criteria: shared goals, established channels of intercommunication, mechanisms for information exchange, dedicated genres, specialized lexis, and thresholds of membership (Swales 1990, 24–27). Of these, it is the latter three that bear most directly on the present argument. The genres a community develops regulate its communicative activity; its specialized lexis demarcates insider from outsider knowledge; and its membership thresholds govern entry into the community itself. These features are especially important for the thesis because they explain how a discourse such as

medieval love can be recognized without being reduced to a single doctrine, code, or school. If, as Schnell (1989) argues, courtly love is better understood as a discourse on amorous conduct than as the repetition of a fixed convention, then its coherence must be sought not in a stable definition but in the recurrent forms through which that debate is conducted: the genres that stage love, the terms that organize it, and the conventions that determine who may speak authoritatively about it.

It is this latter cluster of features—genre conventions, terminological boundaries, and gatekeeping mechanisms—that connects the discourse community to the broader Foucauldian edifice discussed above, in which each operates as a site where what counts as legitimate speech is produced, constrained, and policed. At this level, discourses—together with the affects, symbols, and narratives they mobilize—contribute to ideological cohesion within a discourse community through repeated interaction. These patterned interactions give rise to distinct linguistic norms and practices, forming what Patricia Bizzell describes as a community's "language-using practices" (Bizzell 1992, 222). Classen's notion of a *kommunikative Gemeinschaft* clarifies the same point from the side of medieval literary practice (Classen 2002). A discourse community is not simply a group of people who share beliefs in advance; it is a community sustained through participation in communicative processes. Medieval love discourse, on this account, is constituted by authors, texts, genres, and readers repeatedly negotiating the meanings of love, marriage, desire, virtue, transgression, and social order. Swales' categories make this participation analytically concrete. Genre identifies the literary and institutional forms through which the discourse circulates; specialized lexis identifies the vocabulary through which love becomes recognizable as an object of debate; and thresholds of

membership identify the cultural, textual, and interpretive competencies required to participate in that debate.

As a result, discourse linguistics as a field originates from the turn that occurred during the 1980s, leading to a general reconsideration of "[t]he limits of traditional linguistics" regarding textuality (Spieß 2011, 12). This development first concerned scholars who treated the boundaries of the single sentence as the fundamental unit of analysis without conceiving something beyond "the semantic analysis of linguistic units at the level of sentence statements and propositional linkages" (2011, 12), involving a looser definition of text as well as "a limited sequence of linguistic signs which is internally coherent and which in its entirety signals a recognizable communicative function" (Brinker 2005, 17). According to Kai Vogelsang (2012), the turning point came with the birth of Busse's historical semantics, which elaborates Reinhart Koselleck and Rolf Reichardt's work (e.g., Koselleck 1972; Reichardt 1985). They pioneered a conceptual history framework where *Grundbegriffe* ("fundamental concept") act as both indicators and factors of historical change. Rather than analyzing terms in isolation, their methodology examines how these concepts circulate within the broader socio-political sphere, mapping the fluid *semantische Felder* ("semantic fields") where clusters of interconnected meanings coexist and evolve. Both Koselleck and Reichardt draw their approach from the work of scholars like Erich Rothacker and Otto Brunner, who integrated social history and conceptual analysis to address gaps in historical studies, establishing conceptual history as a field dedicated to understanding past eras through their own language and concepts.

3.4.1 Discourses and corpora

The convergence between the post-Foucauldian and linguistic traditions becomes most concrete in their treatment of the relationship between discourse and its textual evidence. From

the post-Foucauldian perspective, discourse exceeds any individual utterance or document: it consists of historically conditioned relations among statements, forms of knowledge, subject positions, and rules of formation. Foucault's idea of the *archive* therefore does not designate primarily a repository of surviving documents, but the historically specific system that governs the conditions for which statements can appear, recur, and be intelligible. This is the point developed at the beginning of the chapter: the post-Foucauldian expansion of discourse provides the theoretical warrant for treating courtly love as a heterogeneous archive. The theological, sociological, and psychological strands reconstructed in the previous chapter may each be understood as partial archives—historically organized conversations with their own evidentiary habits, privileged texts, and explanatory vocabularies. Discourse analysis makes it possible to bring these partial archives into relation because it defines discourse not by the homogeneity of its sources, but by the historically conditioned relations among statements that circulate across them. Yet, precisely because this archive cannot be observed directly or recovered in its entirety, discourse analysis requires a materially delimited means of access to it. In linguistic and computational research, this is done through the corpus.

For linguistic discourse analysis, a version of the Foucauldian concept of the archive—in other words, the evidentiary base for the analysis—becomes operational as a corpus: a selected and materially bounded collection of texts whose internal relations can be examined systematically. The corpus is therefore not the archive in its totality, nor a neutral sample of it, but the constructed point of access through which any discursive field can be made analyzable (see Scharloth, Eugster, and Bubenhofer 2013). Gavin provides a useful way of conceptualizing this relational function. In his terms, a corpus constitutes "a topological space with an underlying set of elements (tokens)" described by metadata (2022, 7). The corpus is therefore not simply a

container holding texts, but a structured field in which linguistic units, documents, authors, genres, dates, and other bibliographical attributes can be related to one another. Its elements become analytically meaningful not in isolation, but through the patterns produced by their organization within the corpus. The corpus thus provides a bounded space in which their semantic relations can be compared.

The point, then, is not simply that discourse analysis requires a corpus, but that the corpus itself must be understood as an object made for the analysis. If the corpus provides access to part of it, the next question concerns the principles by which that access is constructed. This emphasis on deliberate construction is shared across corpus linguistics and historiography, although the two fields delimit the concept differently. In corpus linguistics, a corpus conventionally denotes a structured collection of texts—usually in electronic form—assembled to support systematic linguistic analysis (Kennedy 1998, 1–3; McEnery and Brezina 2022, 4–28). Corpora are, as Paul Baker puts it, "repositories of naturally occurring language" through which patterns that run counter to intuition and suggest the presence of wider discourses may become observable (2010, 123–24). Yet, as Lynne Bowker and Jennifer Pearson emphasize, they "are not merely random collections of texts but, rather, [...] collections that have been put together according to specific criteria" (2002, 41). Their internal organization may therefore become an object of analysis in its own right, since the relations visible within a corpus depend partly upon the principles governing its construction (McEnery and Brezina 2022, 74–78).

Similarly, Laurent Gauthier defines a corpus as "a set of documents assembled with a specific purpose" (2025, 3), while observing that, in historical research, "any collection of documents may be a usable corpus" and that the logic defining its composition comes "from the historian rather than from the documents themselves" (4). This means that the design of a corpus

is inseparable from the questions it is intended to address, since decisions concerning temporal range, language, genre, authorship, textual form, and document length determine which relations the model can represent. These decisions do not make the corpus arbitrary, but they do make its rationale answerable to the research question.

The question of how that answerability could be secured methodologically was posed directly within German discourse linguistics by Dietrich Busse and Wolfgang Teubert, who posed the question of whether discourse could be considered a linguistic object at all (Busse and Teubert 1994). Busse and Teubert's contribution was to make this relationship methodologically practicable by introducing corpus-linguistic procedures into discourse analysis. Textual corpora thereby became the empirical means through which analysts could move from individual linguistic occurrences towards recurrent patterns of argumentation, association, and representation. As Spitzmüller and Warnke explain, "discursivity is realized at the textual level through intertextuality": discourse acquires "cohesiveness" because some linguistic formulations are redistributed across texts (2011, 89). This feature thus describes the relational structure through which texts participate in a discourse by taking up, modifying, or contesting linguistic materials that also occur elsewhere in the corpus.

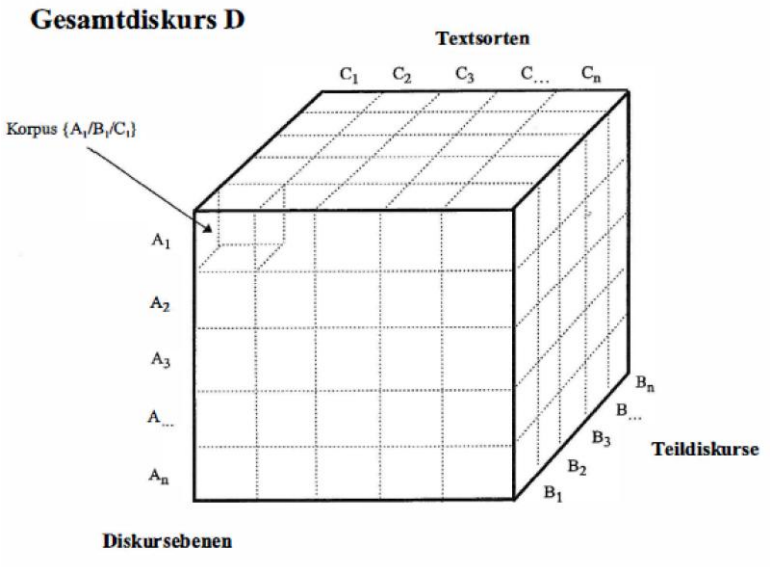


Figure 1: Diagram representing the relationship between whole-discourse and its constitutive segments in a discursive continuum according to Jung (2001, 40)

However, corpora do not stand in a one-to-one relationship with the discourse under investigation. They are deliberately constructed evidentiary objects through which only part of a broader discursive field becomes available for analysis. Matthias Jung argues that assuming the direct correspondence between corpora and discourses is fundamentally misleading since thematically more specific discourses—the "discursive segments"—remain permanently embedded in the general thematic "whole-discourse", which in turn is part of the large-scale discursive practices in a given context. In this regard, Jung explains that:

> A discourse D_1 consists of the set of all statements A_1 to A_n. They are realized within the framework of texts T_1 to T_n, which, however, also partly belong to other discourses D_1 to D_n, consisting of statements B_1 to B_n, C_1 to C_n etc. These statements are taken up again in the chronologically later texts T_1' to T_n' or T_1" to T_n" etc. as A_1' to $A_{[n]}$' etc. (Jung 2001, 40)

In essence, what Jung suggests is that discourses can be analytically modelled based on criteria of thematic individuation of discourses, as illustrated in the diagram (Fig. 1). Discourses are also

linked to the communicative genre to which they belong (Niehr 2014, 36). What is important for Jung, however, is not so much to specify which criteria are used in the delimitation of discourse, but rather to emphasize that these criteria must be systematic and internally consistent. Accordingly, texts can never be classified as isolated discourses, but rather must be understood as nodes within an intertextual network of relations among adjacent, subordinate, and superordinate discursive units (Spitzmüller and Warnke 2011, 89–90). The missing link between discourse and its textual realization and knowledge is therefore language. Within what Felder defines as "pragma-semiotic" framework (see Felder 2013; 2011; 2006), language should be viewed not merely as "a medium for recording reality" but as "a means for constructing reality" (Warnke 2009, 116). This fact, in Martin Wengeler's view, places the emphasis on the "constitutive power of language" (2003, 14), challenging any attempt to establish a strict boundary between linguistic knowledge and knowledge of the world, suggesting instead a significant overlap between the two. This assumption can be traced back to Scherner's concept of "textualized language," for which the prerequisites for understanding text are to be found in a scheme of "networked, textually evocable knowledge" (1994, 336).

Epistemologically, this means that there is a difference between the concept of discourse as a "whole-discourse" (*Gesamtdiskurs*) or as a "segment" of this greater network (*Teildiskurs*), which is created in the very process of charting them out. According to Ekkehard Felder, one needs to consider discourse as a research "construct" made of textual materials (Felder 2011, 122). A discourse's epistemological and ontological borders correspond to where each of the texts forming part of a corpus expresses one or more aspects of a single discourse (Busch 2007, 154). The basis for linguistic discourse analysis could then be the compilation of a textual corpus, which is at least approximately representative of a given discursive practice (Niehr 2014,

28). In this study, the concept is employed exclusively in reference to electronic corpora, understood as "collection[s] of texts which [are] stored on some kind of digital medium and used by linguists to retrieve linguistic items for research" and classified into various categories according to certain criteria such as its epoch, size, themes, medium, authorship, language, and annotations (Lindquist and Levin 2018).

For Sylvia Bendel Larcher, there are four nested levels of discourse according to their scope and empirical accessibility (2015, 15). At the broadest level lies the *potenzieller Gedankenraum* ("potential space of thought"), which encompasses everything that could in principle be thought or said about a given topic, including positions that are currently unarticulated or socially unthinkable. Within this, *Diskurs weit* ("broad discourse"; equivalent to the "whole-discourse" discussed above) designates all statements that are possible at a given historical moment, insofar as they conform to existing systems of knowledge and shared interpretive frameworks. A further restriction is introduced at the level of *Diskurs eng* ("narrow discourse"; equivalent to the discursive segments), which includes only those statements that are actually produced and circulated in concrete communicative forms, such as texts, media outputs, or institutionalized genres. Finally, the *Korpus* represents the most limited level: a deliberately selected and materially bounded collection of texts subjected to analysis. Bendel Larcher's theorization thus makes explicit that corpus-based research operates on only one layer of discourse, while presupposing—and indirectly inferring—broader discursive structures that extend beyond the available textual evidence. In this view, a corpus is thus always an exaggeration of an observed phenomenon—in this case, a discourse—that is limited in material terms.

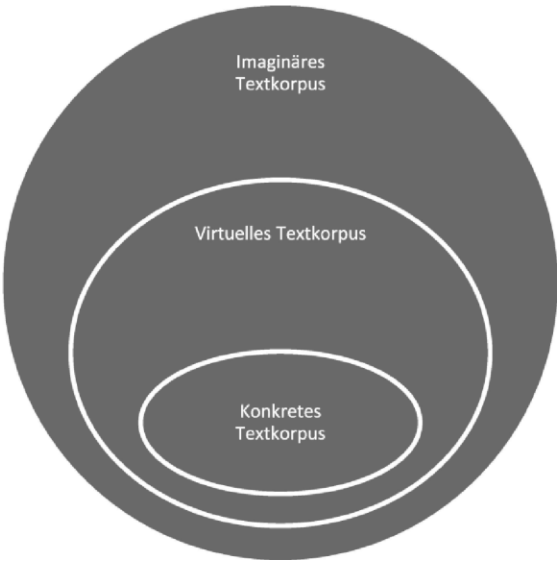


Figure 2: Schema of the three types of corpora in Hermanns (1995, 89) between imaginary, virtual and concrete corpus, starting from the top

Based on these limitations, scholars tend to distinguish among different typologies of corpora. Hermanns, in his classifications (Fig. 2), makes a distinction between *imaginäres* ("imaginary"), *virtuelles* ("virtual") and *konretes* ("concrete") *Textkorpus* (Hermanns 1995, 89). At the basis of his distinction lies the awareness that anyone who wants to undertake a linguistic analysis of speech never has access to all the existing textual sources related to that discourse and, more fundamentally, that the discursive field itself exceeds what is empirically observable. In this sense, Hermanns' taxonomy broadly corresponds to Jung's and Bendel Larcher's categorisations with the constant correspondence between the highest discursive level—i.e., the theoretically unlimited totality of all possible and actual statements related to a discourse (further extended by Bendel Larcher to include what remains merely thinkable and never articulated)—and the textual manifestation accessible through corpora. Hypothetically, both the imaginary and the virtual corpora are quantitatively boundless and materially inaccessible since the first would be constituted by all that has been thought, said, and written related to a certain discourse. The

virtual corpus would include all possible statements related to a certain discourse in a particular time span without considering other criteria such as the medium, whereas the imaginary corpus contains all the statements ever—across temporal boundaries. If such corpora of texts existed, they would ideally be so large as to become unmanageable, which would render them unusable for the purposes of analysis although preserved (Niehr 2014, 40). However, the reasons for such a claim are important, since they include the possibility that much linguistic material has never been recorded—as in the case of everyday conversations—or that it has been irretrievably lost. Therefore, any attempt to conduct a linguistic analysis of a certain discourse depends exclusively on the "concrete" kind of corpus.

As Fritz Hermanns argues, "through a targeted collection" and a process of "sifting and weighting" as well as the "elimination of the unusable and the less productive" materials, "the virtual corpus becomes the concrete corpus of all those texts related to a certain discourse which then form the basis of the linguistic-historical investigation" (1995, 90). This means that only a limited part of the available sources would form part of the corpus, considering a well-defined set of parameters as well as the possibility of retrieving them (Niehr 2014, 41). In these terms, the main criterion of validity for ensuring that a linguistic discourse analysis produces results that are truly meaningful from a scientific point of view becomes that of "representativeness" (Hermanns 1995, 90; Lindquist and Levin 2018, 43). The degree of representativeness of the texts and the linguistic forms they contain should, therefore not be underestimated, as it is precisely this that guarantees that the results of the investigation of an analysed "discursive segment" are also valid for the "whole discourse" (Felder 2011, 122) although this proves to be always a relative parameter, as it is always dependent on the corpus.

Corpora are collections of texts, meaning they provide the raw material needed to identify

the roles, structures, and functions within discourse analysis. Heiko Girnth, in this regard, outlines how texts can be categorized based on three key parameters that help determine their positioning, role, and impact within a discourse and are necessary to outline their macrostructure (2015, 89):

(a) The position in the discourse (initial, procedural, terminal, and variable) determines whether a text initiates, guides, or concludes the discourse. This indicates whether a text can appear at any point or only in a specific position within the discourse, as seen in procedural texts like those in legislative processes, where placement is generally fixed.

(b) Participation status as either primary or secondary (metatexts). Primary texts, such as political platforms or parliamentary speeches, directly shape the discourse, while secondary texts, including news articles or press reports, comment on or contextualize primary texts. Additionally, metadiscursive texts focus on the discourse itself, providing a self-referential layer.

(c) Relevance within the discourse, with some considered "discourse-dominant" because of their significant influence on the discourse's direction.

(d) Relation to other discourses: they may be discourse-immanent (addressing only one discourse), discourse-transcendent (engaging with multiple discourses), or discourse-peripheral (transcending discourses but mainly oriented toward a different discourse).

Together, these roles form a framework for analysing how texts operate and interact within and across discourses, highlighting the structural and functional complexity of discursive practices. Since corpora are made of texts, and those texts conflate all the statements related not only to one but to different discourses, corpora are an epistemological prerequisite: they represent a way to access them, to know, isolate, and identify single discourses—if not in their

entirety, at least in their partiality—starting from the elements that constitute them and that are encoded in language.

3.5 Harrisian distributional discourse

The preceding sections have argued that the post-Foucauldian archive and the linguistic corpus name different levels of the same analytical problem. The archive designates the broader heterogeneous field in which statements become possible and intelligible; the corpus provides the bounded textual field through which part of that archive can be examined. Yet, this still leaves a methodological question unresolved: once discourse has been made accessible through a corpus, how can the relations within that corpus be interpreted as evidence of meaning rather than as mere textual proximity? If the corpus is a structured field of textual relations, then its analytical value depends on the principle that relations among words, texts, and contexts are not incidental but meaning-bearing. In other words, the move from archive to corpus to discourse requires a theory of linguistic relation: a way of explaining why recurrence, co-occurrence, and patterned distribution can disclose something about discourse. This is where the linguistic tradition becomes necessary. It does not replace the post-Foucauldian account of discourse as a historically conditioned space; it operationalizes that space from below by treating meaning as something recoverable from repeated linguistic environments. The following sections therefore move through three linked stages: first, the distributional principle that makes linguistic recurrence meaningful; second, the vector-space model that translates that principle into mathematical form; and third, DIMEAN, which provides the interpretative framework through which those relations can be read as discourse analysis.

The premise of that programme is that meaning is not lodged intrinsically in words but arises from their habitual company and use. An influential philosophical formulation of this

principle appears in Ludwig Wittgenstein's *Philosophical Investigations*, which locates meaning in use rather than in intrinsic definition:

> For a *large class* of cases—though not for all—in which we employ the word 'meaning' it can be defined thus: the meaning of a word is its use in the language. And the meaning of a name is sometimes explained by pointing to its bearer. (Wittgenstein [1953] 1994, §43)

Gavin (2022, 64–66) traces a parallel literary-theoretical genealogy through Mikhail Bakhtin and Julia Kristeva, whose accounts of dialogism and intertextuality likewise understand meaning as positional, relational, and conditioned by prior language use. In the 1960s, Kristeva introduced the work of Bakhtin to French intellectual circles, and in doing so articulated what she called "intertextuality": the principle that no text exists as an autonomous unit of meaning, but is constitutively woven from prior discourse, saturated with the traces of other utterances (Kristeva 1986, 37). For Bakhtin, this was a consequence of what he termed the "internal dialogism of the word"—the observation that every act of signification occurs against a background of competing and prior uses, so that "the linguistic significance of a given utterance is understood against the background of language, while its actual meaning is understood against the background of other concrete utterances on the same theme" (Bakhtin 1981, 281)

Meaning, in this account, is not a property that words carry in isolation but an effect produced by a word's position within a historically accumulated field of use. Kristeva sharpened this spatial logic considerably. Where Bakhtin described the social dynamics of language as impinging on words from without, Kristeva located that social pressure within the structure of the word itself. Words acquire their "status," she argued, through a double orientation—horizontal, toward the writing subject and addressee, and vertical, toward the anterior literary corpus—such that "the word in the text is oriented towards an anterior or synchronic literary corpus" (Kristeva 1986, 37). The word becomes, on this account, a mediating unit whose

meaning is constituted by its differential relationships to other words across texts. Every text, in Kristeva's formulation, is "a mosaic of quotations" and "the absorption and transformation of another" (Kristeva 1986, 37), such that the proper object of literary semiotics is not the text as a closed unit but the dialogical space that the word both inhabits and produces (Gavin 2022, 66). As shown in the previous chapter, the idea of intertextuality is not merely a feature of literary language but relates to any act of linguistic production that is conditioned by the sedimented social history of the forms it employs. Bakhtin and Kristeva describe, in different but related terms, precisely the mechanism by which discourse—as a structured, socially produced, and historically variable field of meaning—is inscribed in language use. Therefore, the word carries discourse within it, while its distributional profile is the trace of that inscription.

The intuitions articulated in Bakhtin and Kristeva found its linguistic operationalization beyond literary theory, through a set of independent but convergent efforts to solve the concrete problem of polysemy across machine translation, information retrieval, and corpus linguistics (Gavin 2022, 73–79). John R. Firth crystallized this principle in the formula "you shall know a word by the company it keeps" (1957, 11), a formulation that has carried important implications for how meaning is understood since the 1950s. Central among them is the notion of *collocation,* defined as the tendency of words to co-occur—i.e., to appear together within and across texts in semantically motivated patterns. The consequences of this new theory are twofold. First, it displaces the assumption that words bear determinate meanings in isolation, relocating meaning instead within the broader linguistic contexts in which words habitually appear. The very idea of collocation depends on the premise that words appearing together with sufficient regularity form stable configurations through which semantic relationships become perceptible. Second—and perhaps more consequentially—it renders meaning in principle observable: if meaning is a

function of use expressed through co-occurrence, it can be tracked and statistically measured from actual language data.

It was in this same intellectual climate that Zellig Harris published "Distributional Structure" (1954). There, he proposed that language could be described entirely in terms of the relative occurrence of its elements—without recourse to meaning, history, or any information external to the distributional facts themselves. Harris had already coined the term "discourse analysis" two years earlier, in his 1952 programmatic article of the same name (Harris 1952). Here, Harris defines discourse analysis as an extension of descriptive linguistics "past the sentence limitation," grounded in the systematic examination of "connected speech (or writing)" through recurrent distributional patterns (Harris 1952, 1–2). Discourse, for Harris, is a formally analyzable domain whose structure can be reconstructed from observable co-occurrence relations within texts, with meaning treated not as an a priori interpretive input but as an outcome of distributional structure uncovered by systematic analysis (Harris 1952; 1954). This principle is what Magnus Sahlgren identifies as foundational for every computational linguistic approach from corpus-based applications to information retrieval, and even machine learning (Sahlgren 2024).

Harris's programme emerged in the immediate postwar period, amid intense enthusiasm for statistical approaches to language. This enthusiasm was fuelled in part by George Zipf's work on the mathematical regularities of word frequency (Zipf 1935) and by Claude Shannon and Warren Weaver's *Mathematical Theory of Communication* (1949), which demonstrated that human language could be modelled as a probabilistic system (Gavin 2022, 70). The success of Allied codebreakers had further demonstrated that language was, in principle, susceptible to machine-based analysis—a conviction that fed directly into Harris's ambition to place linguistics

on mathematically rigorous foundations. His *Methods in Structural Linguistics* (1951) proposed a systematic suspension of meaning as a theoretical category, replacing intuitive semantic judgments with "statistical measurements of linguistic distribution" (Gavin 2022, 70): if meaning was unobservable and therefore scientifically intractable, distribution was not. This suspension, however, generated a productive tension at the heart of Harris's project. His distributional method suggested two divergent lines of inquiry. One, followed most influentially by his student Noam Chomsky, pursued the hypothesis that sentences could be decomposed into semantic "kernels" governed by a finite set of transformational rules—a rule-based, generative path that ultimately displaced meaning from the corpus to an abstract grammatical competence (Gavin 2022, 70–71). The other, and the one for which Harris is now principally remembered, moved in the opposite direction: rather than formalize the structure of sentences, it proposed that the meanings of words could be approximated statistically by measuring their distribution across large bodies of actual usage (Harris 1954; Gavin 2022, 71–73). It is this second trajectory that bears directly on the present study.

In "Distributional Structure" (1954), Harris formulated what has since become the foundational axiom of computational semantics: "difference of meaning correlates with difference of distribution" (Harris 1954, 156). The distribution of an element, he wrote, is "the sum of all its environments"—the complete set of words and positions with which it co-occurs across actual utterances (Harris 1954, 146). Because "the parts of a language do not occur arbitrarily relative to each other" but rather "each element occurs in certain positions relative to certain other elements" (Harris 1954, 146), systematic semantic analysis becomes possible: words appearing in nearly identical contexts—Harris's example is *oculist* and *eye-doctor*—tend toward synonymy, while those whose contexts systematically diverge develop distinct meanings.

Meaning, on this account, can be derived and measured without appeal to psychology, intention, or prior interpretive framework, for "language is not merely a bag of words but a tool with particular properties which have been fashioned in the course of its use" (Harris 1954, 156). It followed that a sufficiently large body of text—a corpus—could serve as the primary object of linguistic investigation: "we have to investigate some actual corpus of utterances and derive therefrom such regularities as would have generated these utterances" (Harris 1954, 161). Statistical traces relevant to meaning are distributed across the data and are, in principle, recoverable.

Harris's later work on discourse and sublanguage clarifies what is at stake in this formulation for the study of discourse specifically (Harris 1957; 1959; 1982). In *Discourse and Sublanguage*, Harris explicitly distinguishes discourse from sentence grammar by defining it as the "directly observable events which constitute the occurrence of language," describable as "patterns of recurrence for the individual words in the various word-class positions in sentence structure" (Harris 1982, 231). What makes discourse analytically tractable is the fact that words recur in patterned ways across sentences, producing higher-order regularities that cannot be reduced to sentence-internal syntax alone: "the fact of word-recurrence within these positions constitutes discoursing (i.e. saying something beyond one free-standing statement)" (Harris 1982, 233). Different types of discourse—narration, argumentation, scientific reporting—are characterized by distinct patterns of word recurrence, which function as the "inherent carriers of various organizations of information" (Harris 1982, 233–234). Discursive organization is empirically recoverable from distributional regularities observable in texts. Harris further introduces the concept of sublanguage to account for the fact that discourses within specialized domains exhibit systematic constraints on word co-occurrence that exceed those of the general

language: a sublanguage possesses its own grammar, derived from corpus-specific inclusions and exclusions, such that "the grammar of a sublanguage is not merely a linguistic exercise, but a classification of the relevant terms and relations of the given subject matter" (Harris 1982, 235).

3.5.1 Meaning as space

While Harris provides the linguistic principle that meaning can be inferred from distribution, vector-space modelling offers its computational realization. The question is no longer only whether recurrent contexts are meaningful, but how such recurrence can be represented mathematically so that relations among words become measurable. Building on Harris's theoretical account of how relational patterns can be operationalized, subsequent researchers developed distributional methods for a range of NLP applications. Two conditions were especially important to this development: the increasing availability of large digital textual corpora and advances in machine learning. More precisely, this development unfolds across two closely related but distinct paradigms: traditional machine learning centred on statistical modelling and rule-based architectures (e.g., clustering, Bayesian probability, decision trees, etc.), and deep learning, centred on artificial neural networks and their capacity to learn layered hierarchical representations of data. Between Harris and the methods used in the present study, there is a long history stemming from Rumelhart, Hinton and Williams' (1986) work on distributed representations, which first established that words could be encoded as numerical vectors; this idea was subsequently applied to statistical language modelling by Bengio and colleagues (2003), and refined through a series of contributions—including Collobert and Weston (2008), Mnih and Hinton (2009), and Turian et al. (2010)—that progressively improved the quality and scalability of word representations.

It was Mikolov et al. (2013a; 2013b) who brought this approach to full maturity in their work for Google. Building on neural networks-based language modelling, they proposed two kinds of architectures: the CBOW and Skip-gram. In these algorithms, a neural network is trained to predict a word from its surrounding context (CBOW) or, conversely, to predict surrounding words from a given target word (Skip-gram), adjusting vector representations through iterative updates so as to maximize the probability of observed co-occurrences. Proposing that words could be mathematically represented as continuous dense vectors in a high-dimensional mathematical space, they demonstrated that they could "explicitly encode many linguistic regularities and patterns" (2013b, 1)—i.e., the coordinates assigned to each word are not arbitrary, but emerge from the model's attempt to compress distributional information into a form that is optimally predictive of linguistic context. Through this process, words that occur in similar contexts are drawn closer together in the space, such that "similar words tend to be close to each other" (2013a, 2), and more complex semantic relations are preserved as geometric transformations. In this way, patterns of co-occurrence are translated into spatial structure: semantic relationships become measurable as distances and directions within the learned vector space.

To understand how "meaning" can occupy a space, it helps to build the idea up from its simplest components. Starting with the example offered by Mikolov et al. of cities (Mikolov et al. 2013b, 1). In a vector space model, the word "Madrid" is not stored as a string of letters but as a sequence of numbers: its coordinates in a high-dimensional space, which cannot be visualized by a human brain. Those numbers are learned by the model from patterns of co-occurrence across billions of words of text. For this reason, a word's position in space is determined by its distribution, and those coordinates work much like latitude and longitude on a map (the first

number encoding its distance east or west, the second its distance north or south). In cartography, a vector of two numbers—each expressing a dimension—are enough to place it precisely, and to tell us something about its relationship to other places: that Paris is physically closer to, say, Orléans than Madrid, because the coordinates describing Orléans are mathematically closer to those describing Paris. From those numbers alone, using simple arithmetic, we can derive how far apart two places are, which lies further east and which further north.

Word vectors work on the same principle, but for a different kind of space—that of meanings—and in far more than two dimensions. Accordingly, each word is assigned a position in a space of hundreds of dimensions: one coordinate for each learned feature of its distributional behaviour. Words that habitually appear in similar contexts end up positioned near each other; words whose contexts diverge are placed far apart. And crucially, the directions between words encode structured relationships. To see why multiple dimensions matter, consider what a richer model of cities would need to capture. Two dimensions—latitude and longitude—can tell us that Paris is closer to London than to Madrid. But they cannot tell us that Paris and Madrid are both Romance-language capitals while London is not, or that London and New York share a financial identity that Paris does not, or that Paris is conventionally associated with love and romance in a way that neither London nor New York is. Each of these resemblances and differences would require its own dimension—linguistic family, economic function, cultural association—and a city's full identity would be the combination of its coordinates across all of them. Word vectors work in exactly this way: each of the hundreds of dimensions encodes a different learned feature of usage, and the position of a word in the space reflects all of those features simultaneously. Therefore, the vector calculation vec(*Madrid*) − vec(*Spain*) + vec(*France*) lands closest to

vec(*Paris*) (2013b, 1) because the direction from a country to its capital is geometrically consistent across the space, regardless of which country you begin from.

But the space encodes far more than this single relationship. The calculation *vec(London) − vec(finance) + vec(fashion)* lands closer to *vec(Paris)* or *vec(Milan)* than to *vec(Frankfurt)*, because the direction from financial centre to fashion capital is also geometrically consistent, even though it has nothing to do with geography. Similarly, *vec(Paris) − vec(French*) + *vec(English)* lands closer to *vec(London)* than to *vec(Madrid)*, because the linguistic dimension is encoded alongside the geographical and cultural ones. The power of the representation lies in its capacity to encode many such relationships at once, so that a single geometric space captures not one kind of similarity but the full range of distributional resemblances a word participates in. Each of the hundreds of dimensions records a different learned feature of usage—geographic, linguistic, economic, cultural—and a word's position in the space reflects all of them simultaneously. This is why distributional semantics can recover meaning from co-occurrence: the position of a word is not a single fact about it but the compressed trace of every context in which it has appeared. Meaning, on this account, moves from the metaphorical to the mathematical, becoming visualizable as an actual space.

3.6 DIMEAN as synthesis

Vector-space modelling explains how distributional relations can be represented mathematically, but it does not by itself explain how those relations should be interpreted as discourse. A model can show that words, texts, or documents occupy similar positions in semantic space; it cannot, on its own, determine whether those proximities correspond to intratextual patterning, intertextual recurrence, actor positions, genre conventions, or historically sedimented knowledge. This is why the final step in the chapter must return from computational

representation to discourse analysis proper. The preceding sections have traced two largely independent routes into discourse analysis. The post-Foucauldian route approaches discourse as a historically conditioned field structured by forms of knowledge, subject positions, institutional constraints, and rules governing what can be said. The Harrisian route begins instead from connected language and asks how discourse can be reconstructed through recurrent patterns of distribution, co-occurrence, and association. These traditions do not share a direct genealogy, nor do they define discourse in identical terms. They nevertheless converge upon a common methodological problem: how to move from materially observable language to larger structures of meaning without reducing discourse either to isolated textual features or to an abstract historical system inaccessible to empirical analysis. The difference lies in the direction from which they approach that problem. While Post-Foucauldian discourse analysis proceeds from historically constituted structures toward the statements and linguistic forms through which they become observable; Harris proceeds from recurrent linguistic forms toward the larger structures that can be reconstructed from their distribution.

Read in light of the preceding discussion of linguistic discourse analysis, moreover, Harris's programme can be understood as converging—without implying a direct scholarly lineage—on a methodological premise compatible with the DIMEAN framework developed by Spitzmüller and Warnke (2011). DIMEAN is not presented as a single method for discourse analysis or as a software procedure, but as a framework for coordinating the different levels at which discourse becomes linguistically observable, from intratextual phenomena through discourse actors to transtextual formations (197–200). Its central contribution is therefore integrative: it coordinates discourse-linguistic analysis across intratextual structures, discourse actors, and transtextual formations without prescribing one direction of analysis as primary.

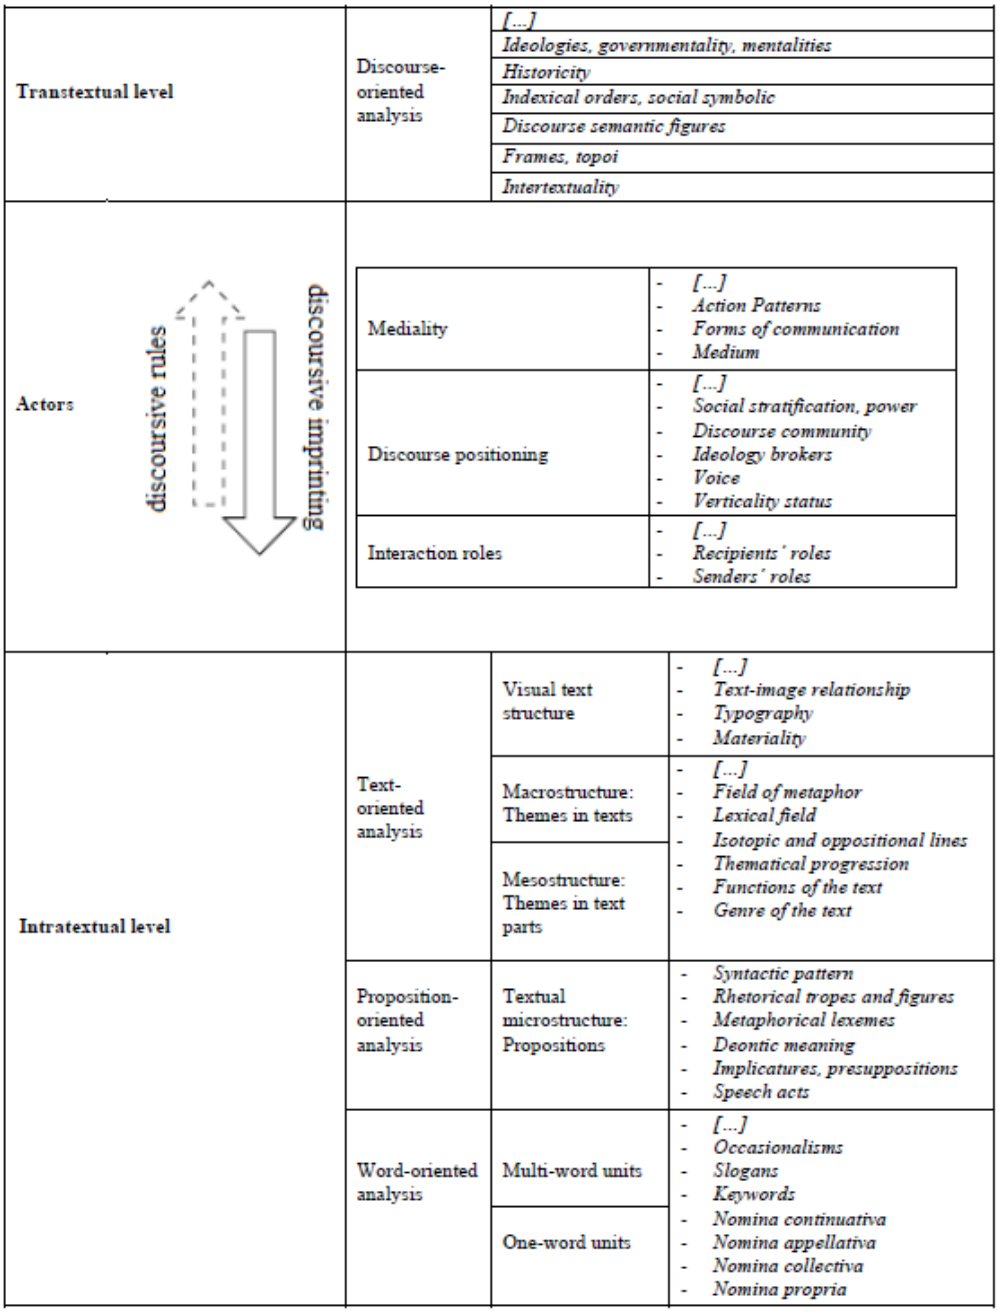


Figure 3: The articulation of the analysis levels in the DIMEAN Spitzmüller and Warnke (2011, 201)

For the present study, this is the point at which the two traditions can be brought into relation: Harris supplies a procedure for recovering patterned relations from linguistic distributions, while the post-Foucauldian tradition supplies the historical framework within which those relations

acquire discursive significance. DIMEAN provides the analytical structure within which movement between them becomes possible.

Harris's distributional hypothesis provides a procedure for reconstructing aspects of discourse from the bottom up, beginning with patterned relations among observable linguistic units. DIMEAN makes the possible movement across analytical levels explicit by locating such relations within a wider hierarchy of intratextual, actor-related, and transtextual analysis (88):

[Discourse [Text [Sentence [Word [Morpheme [Phoneme/Grapheme]]]]]]

DIMEAN treats discourse not as an immediately given object, but as something reconstructed through patterned linguistic and semiotic evidence. At the intratextual level, analysis attends to word units, multi-word expressions, propositions, syntactic patterns, metaphors, presuppositions, speech acts, textual macrostructures, themes, genre, typography, materiality, and text-image relations. At the actor level, it asks how discourse positions, discourse communities, institutional roles, and forms of social participation shape what can be said and how it can be said. At the transtextual level, it moves beyond individual texts toward intertextuality, frames, topoi, discourse-semantic basic figures, indexical orders, ideologies, mentalities, and historically sedimented knowledge formations (137–200). The framework does not prescribe a sequence of stages or a fixed direction for analysis. Instead, it identifies different points at which discourse becomes observable, while the relations among them allow analysis to proceed in either direction.

Although Spitzmüller and Warnke do not describe DIMEAN as a theory of compositionality, a broadly compositional assumption is nevertheless built into the model. It does not claim that the meaning of a whole is mechanically derived from the meanings of its

parts and their mode of syntactic combination. Rather, it presupposes a discourse-linguistic compositionality: the linguistic and discursive levels identified by DIMEAN are related without requiring the higher levels to be reducible to the lower ones. It is in this sense that DIMEAN reconnects a bottom-up analysis concerned with recurrent lexical or distributional patterns to ask how they participate in textual and transtextual formations; with a top-down analysis considering historically constituted discourses to examine the linguistic forms, actor positions, genres, and recurrent structures through which it is realized. In the case of the current thesis, the bottom-up logic of vector-space modelling with the top-down orientation of post-Foucauldian discourse analysis does not reduce one to the other. DIMEAN does not privilege either movement, which is why its usefulness here lies in allowing both to be held within the same analytical model.

Harris's distributional framework and DIMEAN converge with one another at the level of linguistic analysis, while both remain compatible with the Foucauldian conception of discourse developed above. Harris makes discourse analyzable by treating linguistic distribution as evidence of structure; DIMEAN makes that structure interpretable by specifying the levels across which it must be read. Baroni, Bernardi, and Zamparelli (2014) argue that distributional and formal-semantic approaches can be brought together through compositional distributional semantics: corpus-derived vectors represent lexical meanings, while formal operations model how those meanings combine into larger semantic structures. Their argument preserves the Fregean principle that complex meanings depend on the combination of simpler meanings, but replaces model-theoretic denotations with empirically derived vector representations. What this shows is that the distributional approach underwriting both Harris and DIMEAN need not remain atomistic: meaning can be modelled as relational, structured, and emergent from combination. At the same time, as Günther, Rinaldi, and Marelli (2019) observe, vector-space models do not

necessarily implement compositionality in a strict theoretical sense, even when they approximate compositional effects well enough to be useful semantic models.

This synthesis also addresses the concern raised by Jacobs and Tschötschel (2019) that computational approaches to discourse analysis have often remained methodologically superficial. Vector-space modelling can assist with repetition, comparison, and systematization, but only when its outputs remain connected to the epistemological categories that make them meaningful as discourse analysis. The framework identifies relations; discourse analysis determines how those relations are to be interpreted and how far the inference from them can extend. DIMEAN supplies a multi-level approach for making that connection, while Harris supplies the distributional principle from which the computational representation proceeds. Chapter 5 applies this combined framework to the discourse of love in Middle English literature.

The corpus serves as the deliberately bounded evidentiary means through which that discourse is approached; distributional models represent recurrent lexical, semantic, and document-level relations within it; and DIMEAN provides the interpretative structure through which those relations can be examined as intratextual patterns, intertextual formations, and historically situated discourse positions. The analysis therefore does not claim to recover the discourse of medieval love in its entirety. It reconstructs a partial but systematically organized representation through which inherited scholarly categories can be compared, qualified, and reconsidered.

CHAPTER 4: INTERPRETING WITH COMPUTERS

> The work covers texts ranging over three thousand years and written in several languages. In spite of all the unifying factors that we have pointed out, it is impossible for a single person to deal with such a vast material in an exhaustive and definitive way.
>
> —Erich Auerbach, *Mimesis*

4.1 Recapitulation

The preceding chapters were devoted to the identification of a scholarly problem and a method to solve it. Chapter 2 argued that the courtly love debate has reached an impasse whose source is structural: three scholarly traditions, each reasoning from a restricted archive, produced parallel accounts of medieval love that confirm themselves because the evidence from which they reason was selected to confirm them. Chapter 3 showed how this impasse can be overcome through discourse analysis—both in the Foucauldian sense, as a historically conditioned field governing what can be said and known, and in the Harrisian sense, as a distributional domain whose structure can be recovered empirically from patterns of co-occurrence across a corpus. The central claim of that chapter can be stated as a relationship between these two senses: whereas the post-Foucauldian archive names the historically conditioned field in which statements become possible and intelligible, the linguistic corpus provides the materially delimited means through which part of that field can be examined. Collectively, these chapters established that the courtly love debate is a problem of differently restricted archives analyzed through terms that appear shared but designate different things—and that discourse analysis provides the framework for addressing it.

Before that framework can be applied computationally, however, a separate theoretical problem must be confronted—one that is not raised by the discourse chapter but by the nature of the proposed intervention itself. If the thesis claims that computational methods can engage

productively with a traditional scholarly debate, it must account for the fact that computational and traditional approaches to literary texts have frequently been presented as fundamentally opposed. This opposition is most commonly framed through the distinction between close reading and distant reading—a distinction that, over the past two decades, has hardened into what often appears as a choice between incompatible methodological commitments. Close reading is associated with interpretive depth, textual particularity, and the authority of the individual critic; distant reading with computational scale, pattern extraction, and the authority of the dataset. The terms of the opposition suggest that any computational intervention in a traditional debate must either replace what came before or be dismissed as epistemologically inadequate to it.

This chapter argues that the opposition is misconceived—and that clearing it is the condition under which the computational analysis of the following chapter can be understood as a contribution rather than a displacement. The close-distant distinction, as this chapter will show, rests on symmetrical misunderstandings: close reading defines itself against a caricature of computational reductionism, while distant reading defines itself against a caricature of interpretive subjectivity. Both fetishize a privileged object—the minutiae for one, the archive for the other—and both accuse the other of mistaking a partial view for an adequate account of the whole. Answering the question requires confronting not only the rhetorical history of this opposition but the epistemological assumption that has sustained it: that computational and traditional interpretations stand in a falsificatory relationship to one another, such that one can confirm, correct, or supersede what the other established. The point is developed most explicitly through Matthew Jockers's account of computational confirmation and falsification, since his formulation makes clear the model of cumulative knowledge against which the present argument

is directed. This assumption misrepresents how humanistic knowledge works. Clearing it—and recovering, through the concept of the archive, a way of understanding computational and traditional approaches as complementary descriptive practices—is the task of this chapter and the necessary preparation for the computational intervention that follows.

4.2 The reading wars

The field of DH has generated a striking proliferation of terms to define the methodological interstice between close and distant reading (see Aledavood 2024, 6–7). The repeated effort to name the middle ground betrays an underlying unease about whether the opposition that generates it is real. As Primorac and colleagues observe, computational approaches reveal a deeper continuity with close reading than their rhetoric admits: rather than overturning the assumptions of literary theory, they extend New Critical concern with formalism into the computational domain (Primorac et al. 2023). The proliferation of middle terms, in this light, is evidence that the war was always fought over contested terrain. This chapter takes that proliferation as symptomatic of a disciplinary narrative that insists on difference where continuity discreetly persists between what I define as traditional humanists and digital humanists (or DHers). In their critique of algorithmical criticism, traditional humanists argue that the incompatibility stems from epistemological differences (e.g., see Fish 2012a; Marche 2012; Jenkins 2013; Trumpener 2009; and more generally for the Humanities, White 2013; and Hayot 2021). Their underlying assumption is that DH methods are positivistic and thus intrinsically reductionist. According to this view, DH is overly reliant on a scientific model of mathematical abstraction for its generalisations, as opposed to traditional approaches, such as close reading, which are presumed to achieve holism (especially, see Marche 2012). Rather than marking a

genuine epistemological divide, I contend that this opposition is rhetorically constructed in that each side defines itself against a caricature—or parody—of its predecessor.

As this chapter sets out to show, the question of how computational research can intervene in a traditional debate—a question raised by the diagnosis of Chapters 2 and 3, not by the discourse chapter's internal argument—has been obscured by an apparent opposition between close and distant reading that is, on closer inspection, a non sequitur. The alleged divide ultimately depends on how each approach clears conceptual space for its own methodological reorientation through what Harold Bloom calls the "anxiety of influence" ([1973] 1997). This corresponds to a productive misreading that exaggerates and distorts a predecessor in order to secure room for innovation. What differs is not the presence or absence of reduction, but the locus at which methodological authority is established as well as its size. Both approaches stabilize themselves by fetishizing their object of inquiry—i.e., by elevating a privileged object to the status of explanatory ground. Close reading anchors its authority in the form of the singularity—the line, the word, one text or few texts—treating the minute detail as the privileged site where meaning complexity and aesthetic value are secured. Distant reading, by contrast, repositions that authority at the level of the archive—i.e., the textual multitude, the corpus, the dataset, or the information object—treating large size and aggregation as the primary ground from which literary-historical claims may legitimately proceed. After Martin P. Eve's usage (2019), here, the *archive* is more of an analytical category than a simple collection of records.

The chapter argues that once this symmetry is recognized—once both approaches are seen as fetishizing different objects rather than occupying different epistemological planes—the opposition dissolves, and what actually differentiates computational from traditional work becomes visible. That difference is not scale but modelling: the explicit articulation of how

evidence in the archive stands for a larger field under stated conditions. The concept of the archive, theorized in the preceding chapter as the relationship between the Foucauldian discursive field and the linguistic corpus, is therefore not the source of this chapter's problem but its solution—the point at which the close-vs-distant opposition ceases to be an obstacle and becomes a question about how two complementary modelling practices relate to a shared evidentiary object. What actually differentiates contemporary computational work is not scale but modelling—i.e., the explicit articulation of how evidence in the archive stands for a larger field.

The tension between computational and non-computational approaches in the Humanities stems, in large part, from a misunderstanding of the nature of humanistic knowledge itself—one that presupposes a model of falsifiability that does not account for the interpretive and contingent character of literary inquiry. Something of this kind is instantiated by Nan Z. Da (2019) in her "case against computational literary studies," in which she revisits several influential studies in order to expose their methodological limitations. While her critique usefully draws attention to the practical challenges of a field in which literary scholars are not always trained as programmers, and where computational workflows may introduce new sources of error, it also reinforces the perception of close and distant reading as irreconcilable opposites—and in doing so recalls what C. P. Snow ([1959] 2001) famously described as the problem of the "two cultures": the gulf between scientific and humanistic forms of knowledge. The tension between computational and traditional literary inquiry is not, in the end, something to be overcome. As Gary Hall (2013) argues, it is something to be sustained: the encounter between computation and literary criticism generates productive friction, a space of negotiation in which new forms of knowledge become possible precisely because the two approaches do not resolve into one

another. The present study operates within that friction—and the chapters that follow are its record.

4.3 Fetishizing complexity through close reading

Paraphrasing Jonathan Culler, close reading is the *conditio sine qua non* of literary studies (2010, 20). Matthew Jockers makes a related but more specifically methodological claim, describing close reading as the "primary methodology of literary studies" (Jockers 2013, 6). In its classic formulation, it consists of an attentive, technically informed analysis of the linguistic, rhetorical, and structural features of a literary text—whether poem, novel, or play (but virtually any cultural artefact)—at the scale of the individual work or a very small corpus. Essentially, a practitioner of close reading would read a literary work closely throughout but select for detailed explication only those passages that are necessary to the argument or require particular interpretive attention. These passages function as evidence for claims about the work as a whole and may, in turn, support broader arguments about genre, theme, literary movements, or schools. The selectivity therefore belongs not to the act of reading itself but to the critic's presentation of evidence: close attention is presumed to extend across the work, even though only a limited number of passages can or need to be quoted and discussed explicitly. This sort of analysis understands and treats the text as a discrete, autonomous aesthetic object whose meanings and value emerge from the interplay of its internal elements. Whether in historical or philological readings, deconstruction, psychoanalysis, or ideology critique, the act of reading individual texts "closely" has remained the dominant practice of Anglo-American literary scholarship (Smith 2016, 57). Although referred to as a "methodology"—which implies a coherent and codified practice that close reading never quite achieved—the term "has served as a slogan more than as a name for a particular definable practice" (Culler 2010, 20). As Culler reminds us, there has never

been a single or stable form of close reading: "there are indeed different traditions of close reading," mostly stemming from both "Anglo-American New Criticism" and "the French explication de texte" before that (Culler 2010, 21). The roots of this school of literary criticism—and thus close reading as analytical practice—are commonly traced to the interwar period among a group of poets and critics—T. S. Eliot, I. A. Richards, and later the New Critics—who sought to reform literary study by turning away from philological or historical contextualism toward the disciplined interpretation of real aesthetic objects. However, their definitions of interpretation remained largely within literary studies and were shaped primarily by the practice of explicating poetry, particularly historical poetry; they were never unequivocally codified as a unified theory of interpretation. In fact, what united these disparate approaches was not necessarily a common theory but a shared tendency, which I define as the fetishization *of complexity*.

Similarly, Culler defined this ethos as a sort of attentiveness to the "stubbornness of texts," which resist easy comprehension (Culler 2010, 22). Rather than originating fully formed in the 1930s, this movement developed across a relatively long period, with British antecedents in the 1920s establishing the pedagogy and rhetoric of difficulty, and its consolidation and institutionalization in the United States from the late 1930s through the 1950s. In this U.S. phase, the movement came to be known as "New Criticism," which, as Culler notes, "focused attention on the unity or integration of literary works," explicitly opposing the dominant tendencies in universities back then (Culler 2000, 122). Specifically, these tendencies consisted of historicist and documentary scholarship, concerned with:

> [T]he production, transmission, and acquisition of facts about sets of texts. What one established as a scholar, imparted as a teacher, and learned as a student were commonly the names of historically important authors and some basic facts about their lives; the titles, publication dates, and sources [...] of their major works; relations of influence among them; and the readily observable features that distinguished forms, styles, and genres (Smith 2016, 60).

In this sense, New Criticism proposed itself as an alternative to the regime of facts about sets of texts—authors, dates, sources, influences, genres—offering a new model revolving around the analysis of individual works (Smith 2016, 60). Against that backdrop, close reading emerges intuitively as the disciplinary enactment of a new pedagogy that U.S. critics systematized, giving the practice its durable institutional form as a critical approach.

In his *Sacred Wood: Essays on Poetry and Criticism* (1920), T. S. Eliot had already elevated difficulty and impersonality into aesthetic virtues, defining critical seriousness as the capacity for rigorous attention (Smith 2016, 59). Some years later, his intuitions inspired the pedagogical experiments of I. A. Richards, whose *Principles of Literary Criticism* (1924) and *Practical Criticism* (1929) introduced the idea that the careful examination of language and form could serve as a scientific basis for aesthetic education. More precisely, in *Practical Criticism*, Richards claimed that close analysis allowed readers to gain "a much more intimate understanding both of the poem and of the opinions it provokes" (Richards [1929] 1978, 9). His goal was not primarily hermeneutical but pedagogical, seeking to train students in disciplined attention to verbal texture as a means of cultivating taste. As Joseph North notes, Richards' model was "directed towards an advanced utilitarian model of aesthetic and practical education" rather than toward interpretive theory (North 2013, 142). A decade earlier, what began as Richards' pedagogical experiment and Eliot's stylistic programme crystallized, in the hands of John Crowe Ransom, Cleanth Brooks, W. K. Wimsatt, Robert Penn Warren, Laurence Perrine, and others, into a consolidated disciplinary ideology. Their circle—articulated through *The Fugitive*, *The Southern Review*, and *Scrutiny*—gave literary criticism its modern institutional form.

It was the New Critics who turned Richards's pedagogical experiments into a consolidated school of literary criticism. The core ideology of this new inflexion of literary criticism was to demonstrate the complexity within literary works (e.g., Empson 1930; Wellek and Warren 1949; Wimsatt 1954; Brooks 1947; Brooks and Warren 1938; Perrine 1956). For the New Critics, literary hermeneutics did not correspond to a series of notions to be memorized but to the study of real aesthetic objects. Their significance resided in the intricate interplay of verbal features—ambiguity, irony, paradox, and tone—for which the critic's task was thus to show "the contribution of each element of poetic form to a unified structure" (Culler 2000, 122). The notion that literature embodies complexity—that its value resides in the reconciliation of its constitutive aporetic, paradoxical tensions—became a defining orthodoxy of twentieth-century criticism. Following this new understanding, the difficulty of a literary text was not an obstacle to meaning, but the mark of the artist's mastery inscribed in the aesthetic totality of the work of art. These intuitions provided the conditions for a profound epistemological shift, which eventually would crystallize as both close reading and the definitive disciplinary institutionalization of literary studies. In fact, Ransom's 1937 essay "Criticism, Inc." declared that the task of literary study was to develop "technical studies of poetry" that would examine "metric, inversions, tropes, and fictions by which it secures 'aesthetic distance'" (Ransom 1937, 600). His call that "professors of literature are learned but not critical men […] what we need is Criticism, Inc." (587–8) defined the ambition to make criticism systematic, collective, and technically exact, but not as in "a very exact science" (587).

The fetishization of complexity laid the foundations that made it possible for this theoretical configuration to gain credibility. For Ransom and his followers, the literary work became a "living object […] real, individual, and qualitatively infinite" (Ransom 1937, 601);

therefore, the critic's task was to preserve this object's aesthetic autonomy against reduction by mere historical or psychological notionism. This ideal achieved its most systematic formulation through their work, particularly in William Empson's *Seven Types of Ambiguity* (1930), René Wellek and Austin Warren's *Theory of Literature* (1949), and Cleanth Brooks' *The Well Wrought Urn* ([1947] 1975). Its ubiquity is such that it also shows up in people who are talking not exclusively about poetry, but arguably about reading in general. J. R. R. Tolkien's "The Monsters and the Critics" (1936), for instance, challenges prevailing habits within Beowulf scholarship, particularly the treatment of the poem as a quarry for historical and philological evidence rather than as a work of art. By defending the monsters as integral to the poem's structure and meaning, Tolkien presents Beowulf as a coherent organic whole that should not be dismantled into sources or subordinated to historical reconstruction. In this respect, Tolkien's intervention mirrors the broader critical climate: the poem becomes a unified aesthetic structure whose complexity demands preservation rather than decomposition. The defence of the work's integrity thus transcends specific schools and becomes a general orientation toward reading itself.

Together, these critics codified the conviction that great poetry resists reduction and that its meaning inheres and aesthetic significance in the very density of its formal complications. For Empson, aesthetic pleasure derives from the encounter with interpretative difficulty: "a piece of writing is felt to offer hidden riches […] one part after another catches fire, so that you walk about with the thing for several days" (Empson 1930, xi). Empson's early studies supplied the grammar for what would become the cult of textual difficulty. "I would use the term 'ambiguity' to mean anything I liked," he admits in his preface, "and repeatedly told the reader that the distinctions between the Seven Types which he was asked to study would not be worth the

attention of a profounder thinker" (Empson 1930, viii). The deliberate looseness of his key term becomes itself a methodological statement—to analyze poetry is to pursue the "puzzles" that emerge when "alternative views might be taken without sheer misreading" (Empson 1930, xiv). His insistence that "until you have done your analysis of the ambiguities, you cannot be sure whether the total effect is ambiguous or not" (Empson 1930, xi). Ambiguity manifests as an inner life within the text, a movement of meaning that perpetually exceeds itself. It promises an inexhaustible surplus, which becomes the critic's ultimate object of hermeneutic pursuit.

In line with William Empson's emphasis on ambiguity as the source of poetic richness, Cleanth Brooks redefined that complexity as a principle of aesthetic unity. In *The Well Wrought Urn* (1947), Brooks argues that a poem's "structure" is not "form" in the conventional sense of an envelope containing content, but "a structure of meanings, evaluations, and interpretations," unified through "a balancing and harmonizing" of opposites (Brooks 1947, 195)**.** This conception rejects any separation between form and meaning: the poem *is* its structure of tensions. The well-wrought urn of the title, then, serves as a metaphor for this principle: it stands for the poem as a crafted, self-contained vessel whose perfection lies in the interrelation of its parts, where beauty and truth are one. Accordingly, to translate the poem into prose is "heresy" because meaning is not separable from its textual instantiation—as its truth lies in the dynamic equilibrium of opposing forces (Brooks 1947, 201). Brooks goes beyond this formulation, elevating paradoxes to metaphysical principles, as "the language of paradox," he writes, is "appropriate and inevitable for poetry." (Brooks 1947, 3). Against historicist readings that reduce literature to document, then, Brooks insists that we must approach poetry as an autonomous act of communication, "by making the closest possible examination of what the poem says as a poem" (Brooks 1947, xi). The "miracle of communication," he warns, should not be explained away by

positivism: “the positivists have tended to explain the miracle away […] reducing the ‘poem’ to the ink itself. We had better begin with it” (Brooks 1947, xi).

René Wellek and Austin Warren translated this rhetoric of complexity into a general theory of literary unity, institutionalizing it as the guiding epistemological principle of twentieth-century criticism. In *Theory of Literature* (1942), the two scholars align with their predecessors in that they, too, regard literature as “an object of knowledge *sui generis* which has a special ontological status” (Wellek and Warren [1949] 1956, 144). Like the New Critics, they conceive the work of art as an internally coherent system whose intelligibility depends on the interrelation of its parts. However, the novelty of their contribution resides in the methodological programme that this ontology requires. “We must first make a distinction between literature and literary study,” they observe (3). “The two are distinct activities: one is creative, an art; the other, if not precisely a science, is a species of knowledge or of learning” (3). From this premise, they seek to translate the subjective experience of reading into a “coherent scheme which must be rational if it is to be knowledge” (3). Essentially, their argument stems from a double rejection: on the one hand, of the “scientific and historical” models that reduce criticism to “the mere collection of facts or the establishment of generalized historical laws,” and on the other, of the anti-scientific view that makes interpretation purely personal, a matter of “intuition” or “emotional appreciation” (6–7). Against these extremes, Wellek and Warren propose an “intrinsic” criticism concerned with “the work itself,” distinguishing it from “extrinsic” approaches that deal with its social, biographical, or psychological circumstances (61; 127). In doing so, they establish a middle ground between science and impressionism: literary study can be “rigorous without ceasing to be humanistic” (140). The significance of this intervention lies not in a new metaphysics of the poem but in a new methodology of reading. By defining literature as a

structured and intelligible object of knowledge, they provide the conceptual scaffolding through which close reading could emerge as its operative form.

Despite the ink spilt in codifying criticism as a method, the question of what critics in fact *do* remained largely unanswered. The closest account of how such a criticism would concretely resemble is offered by Paul de Man. Describing Reuben Brower's celebrated Harvard course, Humanities 6, De Man recalls that "students […] were not to make any statements that they could not support by a specific use of language that actually occurred in the text" (1986, 23). In this sense, "they were asked […] to begin by reading texts closely as texts and not to move at once into the general context of human experience or history" (De Man 1986, 23). Culler calls this "the radical nature of close reading, achieved through the analytical rigor of attention to the philological or rhetorical devices of language" (2010, 22). De Man notes that such reading begins from "the bafflement that singular turns of tone, phrase, and figure were bound to produce in readers attentive enough to notice them and honest enough not to hide their non-understanding behind the screen of received ideas" (1986, 23). This emphasis on bafflement—on dwelling within, rather than resolving, the text's resistance to comprehension—marks the decisive transformation of close reading from a humanist pedagogy of clarity to a critical practice of uncertainty. In fact, as Culler observes, for close reading "the goal is not primarily to resolve difficulties but above all to describe them, to elucidate their source and implications" (2010, 22). This mode of hyperattention has clear methodological consequences. Close reading operates at a deliberately restricted scale, privileging individual works, single authors, and tightly controlled one-to-one comparisons. Its interpretive logic can be understood as inductive: general claims about literature, form, or value emerge from sustained engagement

with a limited set of textual particulars. The authority of close reading thus rests not on breadth of evidence, but on the intensity of attention brought to a small number of carefully chosen texts.

As the numerous and often conflicting formulations of New Critical principles demonstrate, the movement never defined with precision how its theories should be realized in practice. In this sense, close reading emerged as the disciplinary enactment of a sparse epistemology—one that prized difficulty itself as the surest sign of value. To read closely in this sense is to approach the work as a structure of latent coherence that solely the critic's sensitivity can bring to light. Based on an exegesis of F. R. Leavis' *The Great Tradition* (1948), Mulhern notes that the nature of New Critical criticism was "maieutic," which he defines as "to criticize was to bear witness to meanings that were already adequately constituted in the words on the page, needing only to be 'realized' in the consciousness of the reader" (Mulhern [1979] 1981 , 172). The very "closeness" that defined this kind of reading thus took shape, first, as a minute attentiveness to verbal detail, and second, as an assumption that meaning is fully contained within the text's formal structure, awaiting recovery through critical labour. Whereas I. A. Richards had conceived close reading as a pedagogical exercise to refine judgment and aesthetic sensitivity, the New Critics transformed it into the scholarly practice we know. Unity and complexity were thereby joined: the more intricate the relations within the poem, the more perfectly it could stand as evidence of its own organic wholeness. This is the moment when complexity becomes fetishized—when the text's resistance to paraphrase, its ambiguity and tension, are reimagined as the very marks of its autonomy. The critic's task, accordingly, was not to resolve difficulty but to preserve it, endlessly re-enacting the discovery of coherence within contradiction.

The methodological discipline of close reading thus entails a particular scale of attention—a hyperfocused scrutiny of the individual text or a narrowly delimited set of works. Its labour of description and inference depends on an inductive logic: meaning is built from the minute observation of linguistic particulars, which the critic then elevates to general claims about literary value, authorship, or period style. This small-scale, one-to-one sort of scholarship—monographs on single authors, studies of single poems, and finely tuned comparative readings—became the model of rigour for the Humanities. From this perspective, Franco Moretti polemically observed that "the most problematic aspect of close reading is that it necessarily depends on an extremely small canon" (Moretti 2013, 57). For Moretti and other advocates of distant reading, this limitation appeared to arise from the scale at which conventional hand-and-eye methods could operate. Their critique consequently casts close reading as a form of telescopy: it may disclose considerable detail within its field of vision, but it cannot reveal what falls beyond that field (e.g., see McCarty 2012). Within this polemical account, the intimacy of attention that had traditionally guaranteed interpretive insight becomes evidence of methodological narrowness, while distant reading and macroanalysis promise to extend literary attention across a much larger body of texts.

This chapter, however, argues that such claims misidentify the principal methodological difference. Computational interpretation does not simply reproduce close reading on a larger scale or translate its attentiveness into statistical form. It constructs a model of the texts under examination, selecting and formally representing particular features while necessarily excluding others. The decisive distinction therefore lies not only in the quantity of material surveyed but in the procedures through which textual phenomena are modelled, made visible, and interpreted.

4.4 The polemic against distant reading

Few concepts in DH have been debated as persistently as distant reading. Although the term entered wide disciplinary circulation in the early 2000s, most visibly through the work of Franco Moretti at Stanford, who had been developing his core critique of canon-based literary history since the early 1980s. In Moretti's own words, non computational studies rely on an "extremely small canon" (Moretti 2013, 57). Consequently, distant reading cannot but indicate "the reading of individual texts with the analysis of large-scale patterns" (Moretti 2013, 48). This scholarly endeavour stems from a single, persistent, and relatively simple, critical claim that literary history cannot be explained through close reading of canonical texts because the canon is structurally unrepresentative of the literary field (Bode 2012, 7–9). This view stems from Moretti's early argument that traditional literary history resembles an incomplete map, where the familiar coastlines of canonical works surround a vast, largely unknown interior composed of ordinary, forgotten texts (1983, 15). What follows from this account, however, is not just a project of expanding the map, but the idea that literary history's incompleteness is due to the epistemic insufficiency of the practices which sustain it. What might have been understood as the addition of a new descriptive layer instead becomes narrativized as methodological successor—in a way that distant reading presents itself not as a supplement but as a corrective to the field itself.

Conceptually, it is at this point that the debate hardens into outright opposition with explicit epistemological contours. Synthesizing positions linked to that resistance to computational methods, Martin Paul Eve argues that close readers often attach on a quasi-theological character because close reading is treated not merely as a technique but as an ethical posture toward literature—one that assumes the text demands sustained, attentive presence and rewards such attention with access to its meaning. In this way, close reading becomes a form of

reverence, grounded in the belief that a limited canon of texts merits protection from instrumental or reductive handling (Eve 2019, 5; see also Goldstone 2017, 639 on "the doxa of reading"). This seems to be a response to Moretti's position and furthered by Stephen Ramsay with his ideas of *algorithmic criticism* (Ramsay 2011). This approach defines a distinctively new kind of criticism for computational literary studies, which "[is] derived from algorithmic manipulation of text" (2011, 2). Through distant reading, DH thus promotes a method in which computers function as reading machines rather than as aids to human interpretation. This widespread assumption has, over time, generated the perception that distant reading was a viaticum for negligent—rather than distant—reading.

As Culler puts it, what distant reading risks enabling is not a different scale of interpretation but "sloppy" or "casual" reading (Culler 2010, 20). Close reading, he insists, requires attention to the strategies by which texts generate meaning and effect—attention that cannot be reproduced by computational procedures (Culler 2010, 22). Culler's concern is not merely methodological but disciplinary: computational methods, he suggests, make it possible to conduct literary research without reading at all, substituting pattern extraction for interpretive engagement (Culler 2010, 24). Several critics have argued that distant reading risks adopting a god's-eye perspective, transforming interpretive decisions into authoritative descriptions of the literary field (Spivak 2003; Frow 2008; Trumpener 2009; Rothberg 2010; discussed in Bode 2012, 10–17). Moretti's graphs, described as "snapshots" that place the literary system "literally in front of our eyes," seemed to substitute seeing for explaining (Bode 2012, 17). The danger was not simplification per se, but the naturalization of mediation: categories, corpora, and visualizations appeared as discoveries rather than constructions.

After Moretti, the distance between literary computing and established interpretive practice remained evident in the early scholarly discussion. Susan Hockey observes that literary computing remained largely peripheral to the central activities of literary criticism (Hockey 2004, 10). Even as late as 2007, Stephen Ramsay could still claim that the digital revolution had "not penetrated the core activity of literary studies," which he described as grounded in subjective interpretation rather than algorithmic transformation (Ramsay 2007, 478). These assessments are relevant not because disciplinary acceptance determines methodological validity, but because they register a persistent division between computation and what literary scholars understood interpretation to entail. The epistemological objections raised against computational criticism help to explain that division.

Similar concerns are raised by Trumpener, who characterizes statistically driven literary history as governed by an impersonal logic that identifies patterns without explaining them (Trumpener 2009, 164). Although computation may reveal large-scale regularities, she argues, it displaces the forms of causal inquiry that depend on archival labour, manuscript study, and historically grounded judgment. For Trumpener, interpretation requires what she calls "real footwork," that is, reconstructing decisions made by authors, publishers, and institutions on a case-by-case basis. At large scales, she warns, individual texts risk becoming statistically negligible, and the aleatory dimensions of literary history are flattened (Trumpener 2009, 164–171).

Marche develops a related objection to the treatment of cultural artefacts as data (Marche 2012). In line with scholars such as Jenkins, who contends that the value of art is not measurable and that no quantity of data can explain what a work of art is, what it means, or why it matters (Jenkins 2013), Marche offers insights into the traditional prescription of the epistemological

nature of reading in the Humanities. He, in fact, accuses computational methods of sacrificing complexity, specificity, and contextual depth in exchange for scale, automation, and descriptive patterning. Interpretation, on this view, cannot be abstracted from deep disciplinary knowledge or delegated to machines without loss. Literature, Marche essentially argues, is irreducibly contextual and expressive, possessing qualitative dimensions that resist formalization. By reducing texts to sequences of tokens, algorithmic analysis erases distinction, taste, and reception history, treating all literature as if it were fundamentally the same.

On this account, turning literature into data is an act of epistemic violence, which strips texts of context, value, reception, and affect, reducing them to mere interchangeable strings of tokens to substantiate a hypothesis based on a series of visible relations. Drawing from Marche's critique, it is possible to infer that what traditional Humanities feel is lost in this process is not solely nuance but the very point of criticism itself—consisting of the evaluative, historical, and affective, and thus subjective, dimensions that define literary judgment. In this sense, the literary work is sacred and reserved for the less invasive and more respectful close reading (see Eve 2019, 5). Marche's critique is more philosophically explicit, and for that reason more revealing. When Marche states that "[t]he algorithmic analysis [...] is necessarily at the limit of reductivism," the force of this objection lies in its appeal to a long-standing humanistic suspicion of positivism (Marche 2012).

Implicitly, Marche reveals a traditional humanistic stance by opposing the reductionism he associates with DH to an epistemic commitment to holism. Reductionism, strictly speaking, is not the mere analysis of parts, but the explanatory claim that the behaviour or meaning of a whole can be accounted for entirely in terms of its constituent elements (see Wimsatt [2007] for his critique of naïve level-reduction in scientific practice); holism, by contrast, assigns

explanatory priority to the system, structure, or set of relations within which those elements are embedded, such that meaning or function cannot be exhausted by reference to parts alone (see Nagel [1961] for his account of reduction and emergent wholes). Marche's critique rests on the further assumption—shared by much anti-computational polemic—that close reading uniquely secures this holistic grasp, while distant reading necessarily forfeits it by fragmenting texts into data. Close readers, therefore, accuse distant readers of producing undertheorized, positivistic results—of substituting measurement for meaning and confusing correlation with explanation (Marche 2012). This accusation has a broader corollary, in fact, if computation is reductionist by nature, then distant reading is condemned in advance for neglecting the very qualities criticism is meant to preserve. Tiffany Jenkins—writing from a perspective sympathetic to Marche's predicaments on epistemological reductionism—makes this extension explicit. For Jenkins, the positivistic, reductionist posture is not a contingent feature of particular computational projects but appears inseparable from the computational enterprise itself (Jenkins 2013). Both positions, however, further a kind of *belles-lettrist* exceptionalism, without realizing that they are valid only insofar as one accepts the further premise that close reading enjoys privileged access to totality.

4.5 The anxiety of distance

Distant reading has its own way of critiquing close reading, which begins from the claim that close attention to individual works cannot, by itself, account for literary history at scale. From Moretti's polemical perspective, the problem was not that close reading lacks interpretive subtlety, but that its evidentiary horizon is too narrow for the kinds of historical claims it often supports. Distant reading therefore treats close reading's selective presentation of evidence as though it were equivalent to the whole basis of its knowledge.

Bode summarizes the distant-reading programme as seeking "fewer elements, hence a sharper sense of their overall interconnection," privileging forms, structures, and models over detailed interpretation (2012, 9). Moretti's *Conjectures on World Literature* takes this further by criticizing close reading for its reliance on the exceptional to explain the typical—treating the one per cent of the archive as representative of the remainder. Exceptional works are, by definition, poor evidence for systemic processes of literary production, circulation, and change. Moretti's response is what he calls "a little pact with the devil," for which one must "lose something—details, individuality, concreteness—to gain a broader view" (2013, 48). The compromise consists of upscaling the analysis: larger corpora, quantitative methods, texts operationalized as sets of measurable features, and hypotheses about distribution, similarity, or change over time tested through statistical modelling, pattern detection, and comparison across samples. The procedures resemble those of the social sciences, where inquiry proceeds by defining variables in advance, formulating testable hypotheses, and evaluating them against statistically aggregated data rather than deriving conclusions solely from individual readings.

Underwood extends this argument by asking what experimental methods mean when applied to historical material. To experiment on the past stretches the everyday association of experimentation with intervention and control—but the constraint is not unique to literary history. The natural sciences routinely conduct experiments on fixed datasets (Underwood 2017). Distant reading therefore belongs to what Cleland (2001) describes as a historical science: one that tests hypotheses by seeking to protect them from misleading confirmations rather than by manipulating variables in real time. Once the inquiry is framed on these terms, several familiar objections lose their force.

Moretti himself recognized that close reading at the scale of national or transnational literary archives is practically impossible. His alternative reconceptualizes literature as a collective system rather than an aggregation of individual works (Bode 2012, 8–9). That system, he argues, must be grasped through abstraction—which is "not an end in itself" but a means of widening the domain of literary history and enriching its internal problems (Moretti 2005, 2). In this sense, distant and close reading appear less as general theories of interpretation than as complementary tools, each suited to questions that the other cannot adequately address alone. In the various chapters of his book *Graphs, Maps, Trees*, Moretti instantiates what he means by that. The first chapters draw on quantitative book history to analyze publication trends, genres, and authorship across national contexts, while later ones move inside texts, constructing datasets from features such as plot, character, setting, and device, which are then translated into visual models (Bode 2012, 9–10). Moretti's focus on abstraction gave the field the theoretical ground to begin fetishizing feature-patterning across much larger collections of data. Once literary texts were reconceptualized as systems composed of extractable features—rather than as singular objects demanding holistic interpretation—the possibility of large-scale comparative analysis became methodologically plausible. It is in this sense proper that Moretti's theorizations proved especially generative: they legitimized abstraction as a scholarly operation.

But why must distant reading appear not simply as an extension of literary criticism, but as a corrective to it? As noted by Underwood, literary studies are saturated with succession narratives—New Criticism, New Historicism, Poststructuralism—each framed as displacing what came before (Underwood 2019, 1). The persistence of succession stories in literary studies, Underwood suggests, is historically grounded. Critical debates have often been struggles over a scarce resource—readerly attention—and the discipline's organisational stability has reinforced

the sense that approaches compete for the same intellectual territory (Underwood 2019, 1–2). From this perspective, Moretti's *Slaughterhouse of Literature* appears as a new flag planted on recognized ground, claiming to displace close reading rather than to open a genuinely new object of knowledge (Underwood 2019, 2–3).

One possible explanation lies in adapting Harold Bloom's concept of the *anxiety of influence* to the level of method rather than poetry. In *The Anxiety of Influence*, Bloom describes the psychological struggle of aspiring poets to overcome the burden of their precursors. What he terms *strong* successors, he argues, engage in "poetic misreading or misprision proper" (Bloom [1973] 1997, 14) by reinterpreting, distorting, or strategically revising the work of their antecedents in order to clear imaginative space for themselves. For Bloom, every disciple necessarily takes something from the master, and this inheritance generates the anxiety that drives creative revision ([1973] 1997, 6). Essentially, the anxiety compels the disciple to appropriate and distort that inheritance in order to avoid the fate of the epigone and assert independent authority. Whereas Harold Bloom develops his theory in relation to poetry, I argue that its underlying logic of agonistic succession clarifies the dynamics at work in the close-vs-distant reading debate. Distant reading does not emerge *ex nihilo* within literary studies; it arises within a discipline already structured by the authority of close reading. To secure methodological legitimacy, it must therefore differentiate itself from the paradigm it inherits.

In Moretti's polemical formulations, the differentiation between distant reading and its predecessor must take the form of a parody. Parodies operate as engines of renewal, since *strong* reading—like *strong* poetry—must risk excess to escape repetition (Bloom [1973] 1997, 43–44; 95). If close reading appears as canon-bound and scale-blind, distant reading presents itself as the necessary corrective capable of restoring historical breadth. In light of Moretti's work, distant

reading has been forced into a would-be successor to close reading. Yet, Underwood argues, this rhetoric reflects a peculiarity of literary studies rather than a universal feature of scholarly inquiry. Other disciplines tend, in fact, to narrate their histories cumulatively: bioinformatics did not replace traditional biology; it instead expanded biologists' capacities to pose *broad* questions through statistics (Underwood 2019, 1; also, see Stevens 2013, 242). If the antagonism between close and distant reading can be understood as an anxiety-driven exaggeration of differences—being able to formulate broader or narrower research questions—then the value intrinsic in both scales of reading is not anymore at stake.

4.6 Distant reading's fetish for the *archive*

What emerged from the discussion carried out above is that the opposition between close and distant reading reveals a deeper reciprocity: close reading—no less than distant reading—positions the scale of the other as epistemologically disqualifying. The advocates of distant reading accuse close reading of treating conclusions drawn from a small number of texts as though they applied to a much larger literary field, whereas critics of distant reading reverse this accusation: although computational methods examine far more texts, they argue, those texts are represented only through the limited features that can be formalized and counted. In other words, for distant readers, the sample of works is too narrow, whereas for close readers, the representation of each work is too reductive (see Jockers 2013). In this way, both reproduce the very gesture they criticize in the other—that is, both accuse the other of mistaking a partial view of literature for an adequate account of the whole. It is here that Eve's notion of fetishization becomes particularly illuminating.

Eve introduces the idea that distant reading is a form of "fetishi[zing] the curation of textual archives" (Eve 2019, 1). This argument closely echoes Dominick LaCapra's warning that

when the archive is treated as a substitute for truth, it becomes a "fetish" rather than an object of interpretation (LaCapra 1985, 92). Rosanne Potter claims that "until everything has been encoded, or until encoding is a trivial part of the work, the everyday critic will probably not consider computer treatments of texts" demonstrates this fetish early on in the history of the field (Potter 1988, 93). Similarly, Underwood sees the expansion of digital libraries as having enabled broader historical questions and provided quantitative inquiry with a stronger social foundation (Underwood 2019, xi–xii).

If close reading historically fetishizes textual complexity at the level of the minutiae, distant reading does the same at a larger scale, locating the promise of interpretive completeness in the archive. This fetishization is explicitly instantiated by Gavin's *Literary Mathematics*, where he places corpora at the centre of literary inquiry as the principal objects from which knowledge is produced (2022, 4). Read in these terms, Eve's invocation of the archive appears to condense several distinct stages of knowledge production: the field of available materials, the corpus selected from it, and the computational operations through which that corpus is made to yield patterns. As employed by Eve, the term *archive* consequently refers not only to a collection of texts but also, more expansively, to the range of operations through which knowledge may be extracted from them—mathematical distributions, statistical patterns, correlations, and other outputs of computational modelling. This expansive usage, however, risks obscuring the distinction between the materials available for analysis, the corpus constituted from them, and the model through which their relations are represented—which is the main point of the fetish proper.

The preceding chapter argued, in fact, that the two traditions of discourse analysis—historical and philosophical as well as linguistic—could be reconciled by recognizing that both

operate on the same object: patterned regularity recoverable from a corpus of texts, where knowledge is not simply extracted from an archive but produced through the selection, arrangement, and representation of its materials. Gauthier's account of the corpus in historiographical research makes this distinction explicit. A corpus, Gauthier reminds us, is always "assembled with a specific purpose" (2025, 3), and in historical research the logic governing its constitution comes "from the historian rather than from the documents themselves" (2025, 4). The same principle applies to computational literary research. Decisions concerning temporal range, language, genre, authorship, textual form, and document length determine which relations a computational model can represent. These decisions do not render the corpus arbitrary, but they make its construction answerable to the claims made from it. A corpus is valid not because it reproduces the totality of a discourse, but because its criteria are explicit, internally consistent, and appropriate to the inquiry. Corpus construction is consequently constitutive of the analytical framework rather than incidental to the analytical process: it transforms a broader and materially inaccessible discursive archive into a bounded relational field within which recurrence, association, and intertextuality can be operationalized. The corpus therefore does not overcome the selectivity associated with close reading; instead, the selection occurs in the decisions that determine which texts are included, how the corpus is delimited, and how those texts are represented for analysis.

Once this relocation is recognized, the polemic between close and distant reading appears in a different light. The dispute is not, at its core, about computation versus interpretation, but about competing fantasies of holism. Distant reading is faulted for failing to deliver a total view it never claims to provide, while close reading is exempted from the same charge despite relying on equally selective and reductive procedures operating at a different scale. Once this is

recognized, the accusation of reductionism loses much of its force. Both approaches remain situated within reductionist epistemologies: either hypothetico-deductive in the case of computational methods, or inductive in the case of traditional ones. The distinction between them therefore does not lie in their relation to meaning, but in the horizons of evidence they are designed to traverse.

Recent work in DH increasingly acknowledges that the field has increasingly shifted away from framing close and distant reading as mutually exclusive paradigms and toward theorizing their productive combination, particularly in the study of literary history (Aledavood 2024). Resulting reflections within the field have complicated this stance, with scholars who have emphasized that large-scale analysis neither eliminates interpretation nor renders close reading obsolete, but instead reorganizes when, where, and how interpretive labour takes place (e.g., see Bode 2012; Piper 2015; English and Underwood 2016; Booth 2017; Eve 2019). The proliferation of terms describing mixed or multi-scalar methodologies—such as "scalable reading" (Mueller 2014), "parallax reading" (Sample 2017), "mid-range reading" (Booth 2017; Booth 2022), "telescopic reading" (Unamuno 2017), "meso-analysis" (Saint-Amour 2019), "middle-distanced reading" (Craig 2013) and "middle reading" (Chun and Elkins 2022)—signals a point of convergence for computational literary studies. What these approaches share is a rejection of the reading wars themselves. Rather than asking whether texts should be read closely or distantly, they reframe the problem as one of research process, scope, and iteration: how movement between scales enables discovery, testing, and revision (Rockwell and Sinclair 2016).

The possibility of coexistence between close and distant reading is not a new argument. In his genealogy of the term, Ted Underwood affirms that, historically, distant reading was not originally conceived as a rejection of qualitative interpretation (2017). It arose instead from a

longer tradition of organizing literary inquiry as an experiment, drawing less from literary criticism proper than from book history, structural anthropology, and sociology, where inquiry proceeds by defining features in advance, reading systematically for those features, and allowing patterns to emerge across a controlled sample (e.g., see Rockwell 2003; Sinclair 2003; Ramsay 2003 on the notion of "disciplined play"; Ramsay 2014 on that of "hermeneutics of screwing around"; and Van Zundert 2015 on "screwmeneutics"). He traces this experimental logic most clearly through the work of Janice Radway regarding genre fiction in the context of the 1980s academic discourse on race, class, and gender, where the study of popular romance challenged canonical hierarchies and unsettled prevailing assumptions about literary value and readership. In *Reading the Romance*, Radway does not seek exhaustive interpretation of individual texts, but rather a structured account of how particular narrative features function across a corpus, using questionnaires, coding schemes, and binary distinctions (see Radway 1984). Moretti's early work—especially *The Slaughterhouse of Literature*—shares this methodological orientation to track the presence or absence of narrative clues in detective fiction, and his construction of classificatory trees from a limited sample, closely parallels Radway's procedure (Underwood 2017; see Moretti 2000). Over time, this sense of shared methodological identity was gradually eroded by Moretti's misreading of close reading. Moretti's insistence on differentiation contributed to the perception that distant reading functioned as an alternative rather than a supplement to close reading. One further dimension of this posture lies in the logic of abstraction and experimentation discussed above, which informs much of the rhetoric of exceptionality—and therefore legitimization—surrounding distant reading. The persistence of Moretti's parody becomes visible precisely in those autonomous accounts seeking to move beyond antagonism.

Even when scholars explicitly advocate coexistence, they frequently do so within a framework already shaped by Moretti's rhetoric.

Matthew Jockers' *Macroanalysis* is exemplary in this regard (2013). In an era of big data, he suggests, traditional sampling strategies no longer apply: the availability of massive corpora transforms the conditions under which literary knowledge may be produced, reducing dependence on small, controlled samples and enabling new forms of pattern detection (Jockers 2013, 7). However, Jockers is careful to acknowledge the limitations of both scales. From close distance, something significant will inevitably be missed; from a great distance, something equally significant will also be obscured. For this reason, he maintains that micro- and macro-scale analysis "should and need to coexist" (Jockers 2013, 9). The concept of macroanalysis is introduced through an analogy with economics. Prior to the 1930s, Jockers notes, there existed no distinct field of macroeconomics, only microeconomics, which examined the behaviour of individual agents. Macroeconomics emerged to address phenomena irreducible to the aggregation of individual cases. Jockers proposes a parallel distinction for literary studies: close reading, like microeconomics, focuses on individual texts, whereas macroanalysis seeks to understand literature at the level of the system (Jockers 2013, 24). Jockers explicitly prefers the term *analysis* to *reading*, since macroanalysis foregrounds systematic, quantifiable examination and deemphasizes the hermeneutic act traditionally associated with reading. "This is no longer reading that we are talking about," he writes, even if computational terminology continues to employ *read* to describe file operations (Jockers 2013, 25). Attempts to generalize about periods or genres by synthesizing multiple close readings remain, in his view, a form of microanalysis: qualitative, selective, and case-based. Macroanalysis, by contrast, is "numbers-driven," grounded in aggregation and quantification rather than interpretive synthesis (Jockers 2013, 25).

At the same time, Jockers resists framing macroanalysis as a replacement for close reading. Just as macroeconomics presupposes microeconomic knowledge, macroanalysis depends upon contextual understanding at the level of individual texts. He advocates a "blended approach," insisting that the macroscale perspective should inform close readings by providing "a fuller sense of the literary-historical milieu" (Jockers 2013, 28). It is only through the combined application of both approaches, he suggests, that scholars arrive at "a new and better-informed understanding of the primary materials" (Jockers 2013, 28). The macroscale, in this formulation, functions heuristically: patterns become visible only at a certain scale (Jockers 2013, 32). Its value lies in drawing attention to general trends and missed configurations that demand further explanation. As Jockers himself puts it, the method "calls our attention to general trends and missed patterns that we must explore in detail and account for with new theories" (Jockers 2013, 29). The computer identifies and organizes evidence; interpretation remains a human responsibility. However, Jockers seems to gesture toward the fact that close reading appears epistemologically fruitful so long as it remains local and qualitative. Once it begins to generalize inductively—once it aspires to make claims about systems, periods, or large-scale tendencies experimentally—its evidentiary basis is implicitly rendered suspect. At that juncture, macroanalysis, by virtue of its scale and sample size, is positioned as methodologically preferable. The implication is that inductive generalization properly belongs to the macroscale, whereas close reading is confined to interpretive explication. Scale thus becomes conflated with epistemic legitimacy in the same way Moretti had originally proposed it.

4.7 Melville, Milton, and Watts as examples

Three cases show the shortcomings of this rhetoric of antagonism between close and distant reading and its epistemological entailments. The first is an anecdote told by Jockers in the

foreword to *Text Analysis with R for Students of Literature* (2014); the second is the exchange between Stanley Fish and Mark Liberman that followed Fish's 2012 *New York Times* column on DH; and the third one is Ian Watt's *Rise of the Novel* [1957] 1987. Collectively, they reveal that the falsification framework does not merely surface in polemical arguments for or against computation but rather shapes the terms on which even sympathetic accounts of coexistence are built.

4.7.1 Jockers and the UMD English Department

In the foreword to his handbook, Jockers recounts an exchange with colleagues from the English department at the University of Maryland (UMD), which cuts to the heart of the epistemic nature of the problem:

> During a lecture about macro-patterns of literary style in the nineteenth-century novel, I used an example from *Moby Dick*. I showed how *Moby Dick* is a statistical mutant among a corpus of 1,000 other nineteenth century American novels. A colleague raised his hand and pointed out that literary scholars already know that *Moby Dick* is an aberration, so why, he asked, bother computing an answer to a question we already know? My colleague's question was good; it was also revealing. The question said much about our scholarly traditions in the humanities. It is, at the same time, an ironic question: as a discipline, we have tended to favor a notion that literary arguments are never closed. Do we really know that *Moby Dick* is an aberration? Maybe *Moby Dick* is only an outlier in comparison with the other 20 or 30 American novels that we have traditionally studied alongside *Moby Dick*. My point in using *Moby Dick* was not to pretend that I had discovered something new about the position of the novel in the American literary tradition, but rather to bring a new type of evidence and a new perspective to the matter and in so doing fortify (in this case) the existing hypothesis. If a new type of evidence happens to confirm what we have come to believe using far more speculative methods, shouldn't that new evidence be viewed as a good thing? (Jockers 2014, vii–viii).

Jockers furthers the point with an analogy: just as a new Mars rover returning further evidence that the planet could once have supported life would be scientifically significant even without the drama of a first discovery, computational confirmation adds a piece to a larger puzzle. And beyond confirmation, he suggests, computation carries the potential to present contradictory evidence—evidence that challenges "traditional, impressionistic, or anecdotal theories"—and in

this sense offers "some opportunity for the kind of falsification that Karl Popper and post-positivism in general offer as a compromise between strict positivism and strict relativism" (Jockers 2014, viii).

The comparison is instructive, but not in the way Jockers intends since it belongs to a specific model of knowledge production—one in which evidence accumulates toward a more accurate picture of a domain, in which hypotheses are either confirmed or disconfirmed, and in which the long-run tendency is convergence. Each new finding adds a course to a structure that is, in principle, getting closer to the truth so that falsification is the regulative ideal that keeps the structure honest. However, literary scholarship does not work this way, and the difference is not a deficiency to be corrected. When I. A. Richards close-read poems for ambiguity, he was not building toward a more accurate measurement of English poetry that later scholars would refine. When deconstruction arrived, it did not falsify Richards—it opened a different set of questions, made different features of texts generative, produced a different kind of interpretive productivity. And then New Historicism did not falsify deconstruction and so on. The previous paradigms in each case did not become less true; they became less intensely pursued. A new paradigm in the Humanities thus does not replace its predecessor; it supplements it, running alongside it, available to be taken up again when the questions it generates become interesting once more. The structure looks like something closer to a series of lenses, each of which makes certain features of texts visible and others invisible, none of which cancels the others.

This distinction matters for how we understand what computational confirmation actually does. Jockers' colleague was not making a naive objection when he pointed out that we already know *Moby Dick* is an aberration. He was registering, however imprecisely, that the existing consensus is not an underpowered empirical claim awaiting better instrumentation—it is a

different kind of claim altogether. The claim that *Moby Dick* is exceptional is not a hypothesis in the scientific sense; it is an interpretive judgment, grounded in decades of reading, accumulated through a tradition of close attention to what the novel does that other novels do not, and productive precisely because it opens further questions about form, ambition, and failure. To restate it as a statistical finding about distribution across a corpus of a thousand novels is not to confirm it in a stronger sense; it is to translate it into a different register—one that is valid and potentially illuminating, but that does not stand in a confirmatory or falsificatory relationship to the original claim. The fact that Jockers showed statistically, according to some set of criteria, that *Moby Dick* is, as his colleagues had argued, an aberration is largely irrelevant to future literary research on the topic: it is neither generative nor dispositive to literary arguments that the book is unusual (or, almost certainly, in some cases not unusual) in various different thematic, historical, narratological and other senses.

What is significant, however, is that the consequence of this treatment of humanistic knowledge as hypothesis-based in the scientific sense is that computational analysis gets positioned as a corrective—one that either strengthens existing claims by giving them a firmer evidentiary foundation, or exposes them as impressionistic by showing they cannot survive empirical scrutiny. Both moves misrepresent what humanistic knowledge is and what computation can contribute to it because they point to an idea that our hermeneutics accumulate toward accuracy rather than proliferate toward richness. The value of computational analysis lies not in its capacity to confirm or falsify literary-critical claims but in its capacity to make new descriptive layers available—layers that interpretive reading can then work with, extend, complicate, or use to generate new questions. This distinction also establishes the terms on which the analysis in Chapter 5 must be read. When the document-level model is placed beside Lewis's

categories of love, agreement between them does not amount to confirmation, nor does divergence amount to falsification. The comparison asks instead what becomes visible when an inherited interpretive framework is represented through a different evidentiary model.

4.7.2 Fish and Liberman

Drawing from Daniel P. O'Donnell's work, the Fish-Liberman exchange, represents a second case in which the confusion surrounding the mandate of computation applied to humanistic inquiry (O'Donnell 2018). Fish, a prominent American literary theorist associated with reader-response criticism and long-standing debates about interpretive authority, has frequently engaged questions concerning the limits of formalism and the conditions under which meaning is produced. Liberman, by contrast, is a linguist and co-founder of the *Language Log* blog, whose work often bridges linguistic analysis and quantitative methods. In January 2012, Fish published a column in the *New York Times*' "Opinionator" series titled "Mind Your P's and B's: The Digital Humanities and Interpretation," in which he sought to demonstrate what he took to be the limitations of digital humanities approaches to literary meaning (Fish 2012a). Shortly thereafter, Liberman responded on *Language Log* with a post entitled "The 'Dance of the P's and B's': Truth or Noise?", subjecting Fish's claims to statistical scrutiny (Liberman 2012).

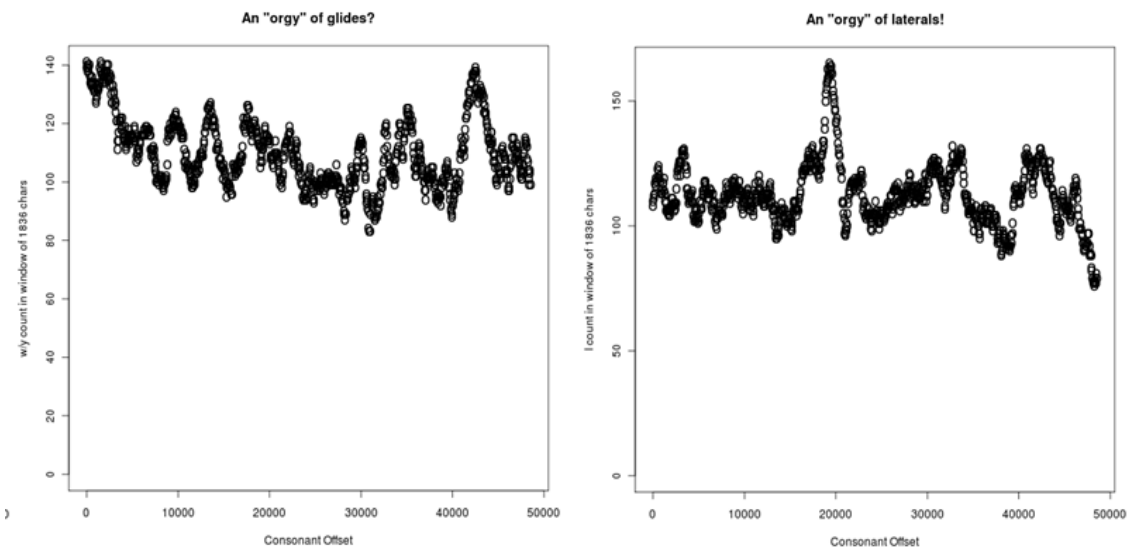


Figure 4: Liberman's diagrams representing distribution of selected consonant classes across Milton's *Areopagitica* (Liberman 2012)

As visible in the diagram (Fig. 4), Fish's close reading of *Areopagitica* turns on Milton's claim that "Bishops and Presbyters are the same to us both name and thing," which Fish argues is formally reinforced by a chiastic pattern in the distribution of the consonants p and b (Fish 2012a). The near-identity of these bilabial plosives, mirrored in *bishops* and *presbyters*, functions for Fish as a phonological enactment of Milton's political argument, a play on sameness that culminates in a local intensification of plosive sounds at precisely the moment where Milton's polemic turns.

Mark Liberman's response is not to deny the plausibility of Fish's interpretation, but to test its empirical distinctiveness. Using simple distributional analysis, Liberman shows that the passage Fish isolates is only one of several peaks in the density of those consonants across the text, and that other letters and letter combinations exhibit comparable fluctuations elsewhere, as shown in Fig. 5. From this perspective, the pattern Fish elevates as meaningful does not stand out statistically from background variation. Liberman therefore suggests that Fish's claim, while

rhetorically appealing, is indistinguishable from noise when measured against the full distribution of the text (Liberman 2012).

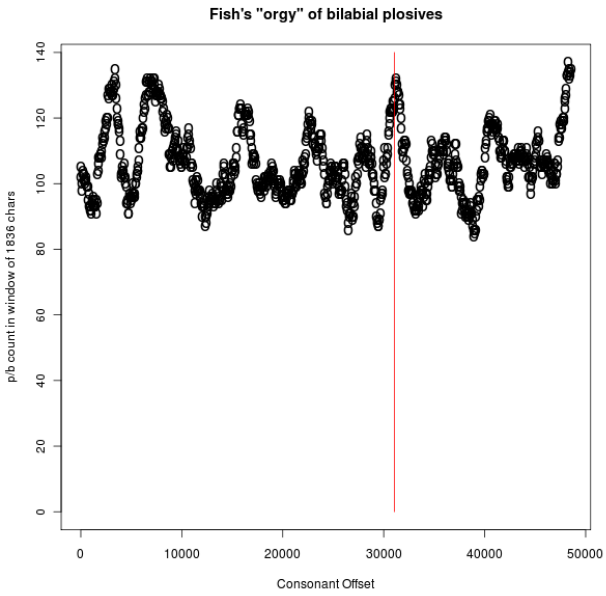


Figure 5: Liberman's chart 1 showing the bilabial consonant distribution (Liberman 2012)

The irony is that Liberman's findings do not contradict Fish's claim so much as refine it. Conversely, it confirms that the cluster Fish identifies is indeed one of the more pronounced peaks, while also drawing attention to other regions—approximately at 1300, 2100, 3500, and 4100 words—where comparable concentrations occur (O'Donnell 2018). Instead of invalidating the reading, the chart effectively proposes further sites for investigation, suggesting that one might look more closely at the interplay between argument and sound in those locations as well.

What emerges is the recognition of layered description. Statistical analysis identifies distributions; interpretive reading assigns significance. The former does not nullify the latter, nor does the latter render the former irrelevant. Rather, each operates at a distinct descriptive register. What differentiates them is not that one proceeds from hypotheses and the other does not—as both do—but the register in which those hypotheses operate and the medium through

which they are articulated. Computational approaches formalize selected features and relations in machine-operable terms; interpretive reading builds its case through contextual and rhetorical argument. The difference is one of descriptive mode, not of analytical rigour. Computational approaches tend to formalize hypotheses through clear feature selection and modelling, thereby defining in advance what counts as evidence; interpretive reading, by contrast, advances hypotheses more tacitly, without formal specification of evidentiary conditions (see McCarthy 2005; Bode 2012; Piper 2018). In this respect, the episode reveals a characteristic feature of Humanities research. Unlike certain scientific paradigms in which competing hypotheses are mutually exclusive, literary scholarship frequently sustains multiple, non-exclusive accounts operating at different scales and levels of abstraction.

4.7.3 Watt and Jockers

Another instructive case is Ian Watt's *The Rise of the Novel* ([1957] 1987), which builds its account of the genre's early development primarily on close readings of five eighteenth-century novels by three authors: Defoe, Richardson, and Fielding (O'Donnell 2018). Watt's project is literary history, not distant reading—but it is literary history that reasons inductively outward from a small number of closely read texts toward claims about the genre as a whole, and that inductive reach is precisely what makes it a revealing test case. Jockers' response to this fundamentally selective endeavour is revealing, as he wonders:

> What are we to do with the other three to five thousand works of fiction published in the eighteenth century? [...] Watt had no yardstick against which to make such a measurement. He had only a few hundred texts that he had read. Today things are different. The larger literary record can no longer be ignored: it is here, and much of it is now accessible. (Jockers 2013, 7–8)

Watt himself recognized deeper methodological problems that extend beyond mere sample size. As he notes, Defoe, Richardson, and Fielding were not simply responding to shifts in the reading

public; their works were "more profoundly conditioned by the new climate of social and moral experience which they and their eighteenth-century readers shared" (Watt [1957] 1987, 7). Moreover, one cannot account for the emergence of the novel as a form without determining what its distinctive literary features are. In other words, Watt's project was never reducible to numerical coverage, but it concerned the articulation of form, experience, and social transformation, as O'Donnell points out (2018).

In a way, Jockers is correct that access to thousands of additional titles transforms the evidentiary landscape. But his characterization of Watt's work risks measuring it against an ambition it did not claim. Watt's argument is fundamentally about his chosen authors and their relation to their context—he thinks that they responded to the pressures of their time in revealing ways. Greater access to five thousand additional titles would not invalidate that focus; it would enrich the contextual background against which it is staged (O'Donnell 2018), but it might also reduce the analytic power of his close reading of this context's impact. The broader issue, then, is not what the traditional Humanities does *not* do—to insist that distant reading is primarily a correction is to make a weak argument for it (O'Donnell 2018). But if distant reading is not fundamentally opposed to traditional humanistic inquiry, then its value lies in support and augmentation rather than substitution. O'Donnell argues that the existence of larger corpora does not render exemplary reading obsolete; it provides a broader field within which exemplarity can be assessed (2018). Watt would not necessarily have written worse history with access to five thousand novels. He might instead have written differently—more historically grounded, perhaps more attentive to distribution and less attentive to detail—but still interpretive in orientation. Distant reading can illuminate tendencies across a field; it cannot eliminate the need for scholars to "cut through the mess" and articulate why certain works crystallize broader transformations.

What emerges from this case is not the obsolescence of Watt's approach, but the persistence of that deeper anxiety of influence Moretti introduced into the debate. This arises only when expansion is mistaken for replacement; when the opposition between micro and macro scale dissolves into a question of methodological supremacy.

According to Underwood, the discovery of large-scale patterns does not create another approach; it operates at a different level of historical description altogether (Underwood 2019, 3). The problem is not that close reading is wrong, but that—as an alternative, equally valid way of doing humanistic research—it cannot, by itself, reveal how the pieces of literary history fit together over centuries (Underwood 2019, 6, 31). Conjoining strong local knowledge with profound global ignorance produces a cognitive dissonance that literary scholars—both close and distant readers—have often resolved by retreating to familiar interpretive scales (Underwood 2019, 31–32). A similar recalibration is undertaken in *Reading by Numbers*, where Katherine Bode resituates quantification within, rather than against, qualitative inquiry. Rather than presenting quantitative work as deductively conclusive, she frames it as speculative and abductive—in continuity with the interpretive practices of the Humanities rather than epistemologically alien to them (Bode 2012, 24).

Drawing on book history and the French *Annales* school of social history, Bode highlights that quantitative methods have long been used with an explicit awareness of their partiality. Statistics, like maps, are "lies to some extent," but they are "lies that tell a truth that would not otherwise be evident" (Joshi 2002, 264). Quantitative approaches necessarily flatten individuality, but for Bode this flattening is not a methodological failure so much as a conditional cost, one that can be epistemically fruitful when it enables access to patterns and trends otherwise inaccessible to traditional literary history (Bode 2012, 14–15). It is at this point

that the opposition of scales—close versus distant—can be recognized as a non sequitur rather than a necessity.

Ted Underwood explicitly identifies this mistaken inference, arguing that the acknowledgement of a blind spot in literary history does not entail the displacement of existing interpretive practices (Underwood 2017). From the perspective of many literary critics, it may seem obvious to respond to distant reading by defending close reading or demanding reconciliation between methods. But from within the project of distant reading, this response misfires because distant reading is not an argument against close reading at all, it is a way to point to a blind spot in literary history—a region of the archive and of historical time that has remained poorly understood because it is not easily accessible at the scale of individual reading (Underwood 2017). A confession of ignorance about that blind spot is not a normative claim about how literature ought to be read; conversely, it calls for deeper scrutiny. Therefore, the appropriate question, Underwood suggests, is not whether distant reading should replace existing methods, but whether the blind spot it identifies turns out to contain anything interesting. Underwood's identification of this non sequitur is persuasive insofar as it exposes a categorical mistake—i.e., the assumption that distant reading is claiming interpretive norms, rather than about the limits of existing historical knowledge (Underwood 2017). However, this reframing also risks underplaying a second, equally consequential slippage—one that operates not at the level of disciplinary norms, but at the level of the terminology utilized. As Johanna Drucker has argued, the term *distant reading* is not merely imprecise but structurally oxymoronic. Reading designates a hermeneutic act grounded in situated human cognition, whereas computation performs operations of sorting, counting, and matching that do not, in any meaningful sense, interpret (Drucker 2017). To describe computational processing as "reading" is therefore to

conflate mechanical procedure with hermeneutic engagement—an objection that holds most strongly when computation is treated as a substitute for interpretation rather than as an adjunct to it. In cases where computation and interpretation are understood as equivalent, the epistemic conditions under which meaning is produced as well as those under which it displaces human agency in the interpretative process are obscured.

In a different intervention, Underwood sharpens this tension further, observing that debates about complexity and uncertainty are often rhetorical rather than epistemological, and that critics too frequently demand that distant reading reproduce the virtues of close reading instead of judging it on its own terms (Underwood 2016). Yet the reverse dynamic is equally visible: as for Moretti, close reading is implicitly measured against the ambitions of distant reading and found wanting—criticized for its limited scale, anecdotalism, or canon-bound selectivity. If distant reading operates in a conceptual space between processing and interpretation, its naming has encouraged a false symmetry between the two. In fact, Drucker suggests that the terminology of distant reading nevertheless invites this misinterpretation deliberately. The problem is not merely rhetorical confusion among critics, but the way the metaphor of reading masks the extent to which computational analysis depends on prior acts of human modelling, parameterization, and interpretive decision-making (Drucker 2017). In this sense, the non sequitur is not external to the project of distant reading; it is, at least in part, generated by the language through which the project has been articulated. What emerges, then, is a more precise account of the non sequitur. The mistake is not simply to oppose distant reading to close reading, but to assume that both belong to the same epistemic category. Distant reading does not compete with close reading because it does not read at all in Drucker's sense; nor does it replace interpretation, as Underwood insists. Instead, it reorganizes the conditions under which

interpretation becomes possible by relocating human judgment upstream—into the design of models, the selection of features, and the framing of hypotheses.

The task of the Humanities is therefore not to reject numbers, but to "recognise them as a form of representation" and to interrogate how they function epistemologically (Bode 2012, 12–13). This entails rejecting the view of computers as neutral "knowledge jukeboxes" that simply deliver results, and instead treating computation as a set of representational practices with profound consequences for humanistic inquiry (Bode 2012, 22–23). At this point, Bode is able to replace Moretti with Willard McCarthy's concept of modelling (McCarthy 2005). Models, she argues, are explicitly provisional, reductive, and manipulable. They are built not to be final representations but to be modified, tested, and challenged. Modelling foregrounds the contingency of data and emphasizes knowledge as a process rather than the result of a process (Bode 2012, 23–24). As McCarthy puts it, models are "temporary states in a process of coming to know rather than fixed structures of knowledge" (McCarthy 2005, 26–27). What Bode ultimately offers is not a reconciliation between close and distant reading, but a rearticulation of their relationship within a shared hermeneutic framework, one in which different scales of readings are complementary amongst many other forms of interpretation.

4.8 Models Beyond Scales

If the debate over distance can be reframed as a question of descriptive layering rather than methodological rivalry, then the next step is not to move beyond scale altogether, but to situate scale within a broader account of modelling. Lev Manovich frames computational cultural analytics as a means of augmenting close reading through new interfaces, visualizations, and iterative techniques for observing cultural datasets and flows, explicitly drawing upon Douglas Engelbart's notion of "augmenting human intellect" (Manovich 2020, 11–15).

Algorithms, in this account, can measure many characteristics at speed, but they cannot determine what matters: humans remain better at recognizing significance, anomaly, and contextual relevance, even as their ability to survey cultural fields unaided collapses under scale. This assumption leads Manovich to observe that cultural analysis often requires attention to diversity, rarity, and infrequency rather than the clustering of common features (Manovich 2020, 15). Scale, therefore, is not inherently epistemically virtuous; its significance depends on how it is embedded within a model that determines what counts as a meaningful pattern. Some cultural datasets are too large to be surveyed through unaided human cognition (Manovich 2020, 17), while the interpretative value of smaller datasets depends upon whether they can support the kinds of generalization being attempted, and the critical task lies in navigating this threshold rather than assuming that "bigger" necessarily produces insight (Manovich 2020, 17). Underwood similarly argues that the preference for bigness over insight misconstrues what distant reading—and, more broadly, reading at any scale—sets out to accomplish (2017). Manovich describes the relationship between computation and close reading principally as one of augmentation. The argument developed here differs in one important respect: computational modelling need not improve, extend, or complete an existing close reading in order to possess interpretative value. It can instead provide a complementary description of evidence at a different representational register.

This displacement of scale as the governing methodological distinction finds a more systematic articulation in Andrew Piper's *Enumerations* (2018). Piper does not ask whether we should read more texts or read them from further away. Instead, he asks how computation participates in meaning-making. Data and modelling "unquestionably" matter for literary study, he argues, but the dominant framings of DH—distance, bigness, and objectivity—rest on binary

models of reading that remain insufficiently grounded in earlier humanistic practice. His aim is corrective: to "align new kinds of models with old kinds of questions" (Piper 2018, ix). Piper's argument permits a further inference central to this chapter: claims about what a change of scale allows a critic to see depend upon a prior model specifying what is being counted, compared, or represented. Scale names the quantity, scope, or granularity of the material under examination; a model specifies how that material is constituted as evidence and which relations within it become analytically significant. This reframing addresses what Piper identifies as the Humanities' persistent part–whole problem: literary studies routinely make expansive claims—about "the novel," "realism," or "world literature"—on the basis of evidence whose representativeness is assumed rather than demonstrated (Piper 2018, xi). Computation renders quantity, recurrence, and distribution newly tractable, but these constitute different forms of cultural evidence rather than corrections of prior interpretations (Piper 2018, 2). What is at stake, therefore, is not scale in isolation, but the way scale has functioned as a proxy for more fundamental disagreements about evidence, representation, and interpretation.

Underwood extends this shift from scale to modelling in *Distant Horizons*, where the stakes of representation and evidence are made explicit. His well-known analogy of the traveller who sees mountains but not the curvature of the earth underscores not a deficiency of the traveller's vision but a constraint on what any single vantage point can make visible (Underwood 2019, x). What changes is not simply the quantity of evidence but the vantage from which relations among texts are represented. For Underwood, numbers become meaningful only when framed through models of literary concepts—genre, audience, character—whose meanings arise within interpretive communities (Underwood 2019, xii–xiii). Quantitative work degenerates into "data dredging" only when measurement precedes hypothesis (Underwood 2019, 17). A model is

not a transparent, exhaustive, or final representation of its object, but a heuristic whose limitations help clarify what the object resists being represented as. Models do not eliminate ambiguity; they render particular forms of uncertainty—variation, distribution, and indeterminacy—explicit. Every computational output is, in this sense, what Rockwell and Sinclair call an "interpreted artifact": not a neutral representation of the text but an object that itself requires reading, shaped by what the tool made visible and what it set aside (Rockwell and Sinclair 2016). Disagreement between models is therefore interpretively productive rather than methodologically fatal (Underwood 2019, 27).

The position advanced here is therefore not that modelling supplies the methodological rigour that close reading lacks. Both traditional interpretation and computational analysis construct models of their evidence, select relations to foreground, and reason from partial representations of larger literary fields. The relevant distinction lies not between modelling and interpretation, since both practices model, but between different modes of modelling: how their representations are produced, which relations they preserve, and what kinds of claims they are equipped to support. Non-computational inquiry can develop historically and contextually situated accounts of passages, texts, and intertextual relations through discursive argument. Computational modelling formalizes selected textual features in machine-operable terms and represents their distribution across a deliberately constructed corpus. Neither practice exhausts its object, and neither makes all the assumptions governing its claims fully transparent.

Scale becomes significant only within these different modes of modelling. A larger corpus does not by itself produce a more adequate interpretation, just as sustained attention to an individual text does not provide access to every relation within the wider field. What changes across scales is the descriptive register within which evidence becomes visible. Traditional

approaches can describe semantic complexity, rhetorical ambiguity, and historical particularity at the level of a passage or text; computational modelling can describe recurrence, variation, proximity, and distribution across a corpus. These capacities may overlap, but they are not interchangeable. Nor are they successive stages in which computational modelling verifies, enlarges, or improves the findings of close reading. They produce partial and complementary representations whose correspondences and divergences must themselves be interpreted.

This distinction also clarifies what is meant by the inspectability of computational models. Formalization requires particular decisions—about corpus boundaries, textual features, parameters, categories, and measures—to be stated in a form that a machine can execute. It does not follow that computational interpretation is free from tacit assumptions, nor that traditional criticism depends only upon personal authority. Traditional humanistic research makes its reasoning answerable through historical evidence and argumentative justification; computational modelling makes selected procedures answerable through their formal specification and reproducibility. Each possesses forms of explicitness and forms of opacity specific to its medium. Jerome McGann argues in *Radiant Textuality* that formalizing an analytical procedure can produce a heightened awareness of particular methodological decisions; this is a consequence of translating interpretation into a machine-operable form, not evidence that close reading lacks its own mechanisms of critical accountability (McGann 2001). Formalization makes certain assumptions inspectable while leaving others embedded in the construction of the corpus, the selection of features, and the interpretation of the resulting model.

The argument therefore moves beyond the opposition between close and distant reading without claiming that scale has become irrelevant. The analytical significance of scale depends upon the construction of an evidentiary model: what is included, how the evidence is delimited,

which textual features are represented, and which relations are treated as meaningful. The following chapter extends this argument through discourse analysis by examining how a necessarily bounded corpus can stand in for a wider discursive field and how its construction determines which relations become available for analysis. Questions of representativeness, exclusion, and evidentiary scope are therefore not preliminary technical matters; they establish the conditions under which the computational description can make any historically defensible claim.

The later analysis puts this complementary relation into practice. When vector-space models are compared with C. S. Lewis's categories, their divergence does not establish that the computational representation is more accurate than Lewisian interpretation. The two accounts proceed from differently delimited evidentiary bases and represent different kinds of relation. Lewis supplies historical and interpretative distinctions that the vectors cannot derive independently; the vectors supply a distributional representation that Lewis's reading neither sought nor was designed to produce. The vectors therefore do not lead to Lewis's interpretation on their own, just as Lewis's categories do not anticipate the distributional organization produced by the model. The contribution emerges from placing these partial descriptions in relation. Their correspondences and divergences allow Lewis's categories to be criticized, qualified, and sharpened—without supposing that the computational model either confirms or supersedes them. Complementarity lies not in their convergence upon a single, improved reading, but in the critical perspective produced by comparing representations that remain methodologically distinct and mutually irreducible.

CHAPTER 5: APPLICATIONS AND RESULTS

> Meaning is not a property of individual words or groups of words but an ongoing performance of the world in its differential dance of intelligibility and unintelligibility.
>
> —Karen Barad, *Meeting the Universe Halfaway*

5.1 The heart of the matter

The preceding chapters have answered, in sequence, the questions that this particular computational intervention required. Chapter 2 argued that the courtly love debate has reached a cul-de-sac whose source is structural rather than interpretive: three scholarly traditions—philosophico-theological, sociological, and psychological—each built their accounts from overlapping but differently delimited bodies of evidence, producing parallel inquiries whose internal coherence derives partly from the evidentiary fields through which courtly love has been constituted as an object of inquiry. Schnell's reframing of courtly love as a discourse—in other words, an ongoing, contested negotiation about the meaning of love distributed across texts, authors, genres, and periods—opened the possibility of addressing that impasse through methods capable of recovering discursive structure from a corpus rather than from within any single strand's preferred texts.

Chapter 3 examined two major traditions of discourse analysis that are often treated as largely independent of each other—one philosophical, centred on Foucault and the historical analysis of discursive formations, the other linguistic, centred on Harris and the distributional recovery of meaning from co-occurrence patterns in a corpus—and showed how they can be brought into relation without collapsing their distinct genealogies or conceptions of discourse. The post-Foucauldian tradition provides the concept of discourse as a historically conditioned field of statements and relations; Harris's distributional hypothesis provides the empirical

principle by which recurrent linguistic patterns can be treated as evidence of semantic relation; vector-space modelling provides the computational means of representing those relations; and DIMEAN specifies the levels across which they may be interpreted as evidence of discursive structure. Together, these traditions help establish that the courtly love debate's impasse—produced by scholars reasoning from differently restricted archives—could in principle be addressed by analyzing a broader corpus through distributional methods that do not presuppose the categories they are testing. In the analyses that follow, DIMEAN is therefore adapted as an interpretative framework rather than treated as a computational procedure in its own right: the word-level analysis begins from lexical and intratextual relations, while the document-level analysis moves toward transtextual relations among works, with genre, authorship, and historical position reintroduced at the point of interpretation.

Chapter 4 then asked on what epistemological terms such an intervention could proceed. It argued that recent debate has often treated computational and traditional approaches as incompatible, but that this incompatibility rests to a great extent on symmetrical misunderstandings—close reading defining itself against a caricature of computational scale, distant reading defining itself against a caricature of interpretive subjectivity—and that the apparent opposition can be resolved by recognizing that both are modelling practices distinguished not by scale but by the medium through which their analytical decisions are formalized. Computational modelling does not replace the interpretive tradition, it enters; conversely, it operates at a different descriptive register, producing a complementary representation whose contribution emerges from being placed in relation to existing scholarship rather than from superseding it.

The question this chapter addresses is thus whether computation, working through the distributional discourse-analysis framework introduced above, and interpreted through DIMEAN, can contribute something meaningful to a debate that has been in motion for over a century—and if so, what kind of finding it produces and under what conditions. The test undertaken here is deliberately bounded: its purpose is to establish whether such an intervention is methodologically productive, not to provide an exhaustive reconstruction of medieval love discourse. Medieval literary data is not a stable, standardized object that can simply be submitted to a computational pipeline. It is fragmentary, unevenly preserved, radically variable in orthography and dialect, and mediated through layers of scribal transmission and editorial decision-making that introduce noise at every level. Any computational analysis that does not account for these conditions risks mistaking the noise for the signal—or, worse, building an analytical framework on a corpus whose construction has not been examined. The chapter therefore opens by addressing corpus design, textual instability, and orthographic normalization—each of which required decisions whose consequences for the analysis are stated and justified. These are not preliminary technical matters but the point at which the thesis' theoretical commitments become operational: the preceding chapters established archival restriction as the source of the impasse, corpus-based analysis as the method for addressing it, and the construction of the corpus as the modelling decision that determines what the analysis can legitimately claim.

Only once that foundation is established does the chapter move to the analysis itself, which proceeds in two stages. The first works at the word level, using fastText to map the distributional field of love across the corpus without imposing any prior categorical framework. This stage addresses the archival problem of the preceding chapters directly: if each scholarly

tradition was built from a differently delimited evidentiary field, does a deliberately broad and heterogeneous Middle English corpus show those registers coexisting, crossing, and interpenetrating in ways that become difficult to perceive when each is examined primarily through its own evidentiary base? The second stage moves to the document level, using sentence embeddings to engage with a specific and influential inherited framework: C. S. Lewis's taxonomy of love. As Chapter 2 showed, Lewis's synthesis was the single most consequential intervention in the courtly love debate—the framework from which subsequent scholarship either developed or defined itself against, and the point at which the evidentiary lineage established by Paris was consolidated into the categories through which medieval love has since been interpreted. Testing those categories against the distributional organization of a broader corpus is therefore not an arbitrary choice of benchmark but a direct engagement with the framework around which much of the subsequent debate was organized. The purpose is not to falsify Lewis, but to determine which distinctions remain coherent when represented at the level of document distributions and where the inherited framework becomes less adequate to the organization of the corpus. The two stages ultimately answer the thesis' central question not in the abstract but in practice: whether computation, working through discourse analysis, can sharpen a longstanding and still unresolved scholarly debate by intervening in it on terms both traditions recognize.

5.2 The corpus construction

The conventional requirements for constructing a corpus of Middle English texts cannot be applied straightforwardly. Corpus studies ordinarily assume that the texts entering the collection are stable and comparable units, that their constituent elements—principally word-tokens—can be identified consistently and set in relation to those of other texts, and that the

resulting selection can be assessed against some larger referential totality. None of these assumptions holds for Middle English without substantial mediation. The surviving textual record does not arrive in a form ready for computational analysis: it consists of materially and editorially mediated witnesses whose linguistic forms, chronological identities, and categorical boundaries remain unstable. Several problems are especially consequential: (a) representativeness, given that the surviving textual record offers no stable totality against which a corpus could be measured; (b) periodization and dating, given that the boundaries of the relevant literary moment remain contested and that composition, transmission, and editorial production belong to different historical layers that cannot be collapsed into a single date; (c) the treatment of long and composite works, whose internal heterogeneity require individual decisions about which editions to use and how to segment them for analysis; and (d) the radical orthographic and dialectal variability of Middle English itself, which demands a normalization strategy capable of recovering the semantic signal without erasing the linguistic record. From the outset, criteria of corpus construction must therefore cope with the absence of any complete referential totality against which the resulting corpus could be measured.

5.2.1 Representativeness of the evidence base

As established in Chapter 2, any concrete corpus is necessarily a bounded approximation of a larger discursive field. The analyst never has access to the totality of statements constituting a discourse and must work instead from a deliberately selected and materially limited collection (Hermanns 1995, 89–90). The practical question here is therefore not whether the corpus is exhaustive, but what claims its construction can reasonably support. A corpus is conventionally considered “representative of a language when the distribution of linguistic phenomena in the corpus is almost identical to that of the language as a whole” (Stefanowitsch 2020, 30; see also

McEnery and Wilson 2001, 76–79). This definition, however, presupposes the existence of a relatively stable and accessible linguistic totality against which such representativeness can be measured, a condition that cannot be satisfied straightforwardly for Middle English literary evidence. Representativeness must therefore be understood here as bounded and purposive rather than as statistical correspondence with Middle English literature as a whole. The present study operates at the level of what Hermanns calls the "concrete corpus": the materially available subset upon which analysis is actually performed (1995, 89). It does not claim to represent Middle English literature in its entirety, nor Middle English language use as a whole. Rather, it comprises a deliberately selected body of surviving, edited literary texts through which the discourse of love can be examined comparatively. Its coherence derives not from completeness but from the analytical question governing its construction. The validity of the corpus for the present study consequently rests on the explicitness of those selection criteria and on their suitability for the methodological test being undertaken, rather than on a claim that the resulting collection exhausts the literary field.

The surviving material also complicates Sinclair's influential description of corpus data as the "genuine communications of people going about their normal business" (1991, 18). In literary-historical research, the relevant "normal business" is not ordinary conversation but the production, copying, adaptation, and preservation of culturally meaningful texts. In this sense, the corpus does not provide direct access to ordinary linguistic practice, but to historically situated forms of representation. While this criterion of authenticity is relatively unproblematic to satisfy in modern corpora, it becomes considerably more problematic in historical contexts. What is available to the present study is therefore a stratified and mediated record of how love

was represented within surviving Middle English cultural artefacts, rather than an unmediated record of medieval linguistic or social practice.

Bárbara Bordalejo's distinction between *document*, *text*, and *work* provides a useful account of this mediation (2013). As she argues, what we encounter in the material record are not stable texts, but what she terms "the text of the document," defined as "the complete sequence of marks present in the document" (Bordalejo 2013, 67). These marks include not only words, but also all meaningful signs—insertions, deletions, marginalia—from which readers construct what Bordalejo calls "variant states of the text," i.e., different interpretative configurations derived from the same document (Bordalejo 2013, 68). In this sense, textuality at the level of the document is inherently unstable: a single manuscript witness may give rise to multiple textual realizations, and no individual document can be said to capture the text it transmits fully. Against this instability, Bordalejo posits the concept of the "work" as an abstract but necessary category. The work is not identical with any of its material instantiations; rather, it is "a conception in the mind of an author […] that serves as a minimal denominator to identify its remaining physical manifestations" (Bordalejo 2013, 71). It is therefore not directly recoverable from the documentary record, but inferred from it: a construct that allows different textual witnesses to be recognized as belonging to the same discursive object. As she further notes, "traces of the work […] can be found in the documents and the texts they hold, but the work itself is none of these instances while, at the same time, is somehow present in all of them" (Bordalejo 2013, 72). No single manuscript or edition is therefore identical with the work it transmits.

This is important because the present thesis does not attempt to model the full multiplicity of documentary witnesses. Such variation is indispensable to textual criticism and manuscript study, but treating every witness as an independent analytical object would introduce

a level of granularity that exceeds both the aims and the analytical scale of this research. This is not because variation ceases to matter, but because its significance changes with the level of observation. While differences between manuscript witnesses are crucial to close reading and textual scholarship, at scale this degree of granularity becomes less informative, since such variations may be indistinguishable from noise. For this reason, the analysis proceeds on the basis of established editions, which mediate between individual witnesses and the abstract work and provide an analytically workable textual unit for the level at which discourse is modelled here.

While large corpora are often associated with greater reliability, corpus linguistics has repeatedly shown that "bigger is not always better" when dealing with specialized domains (Bowker and Pearson 2002, 46). For targeted investigations, a smaller corpus that is carefully designed and thematically coherent may yield more meaningful results than a larger but less focused dataset. This principle, however, sits in tension with the data requirements of distributional modelling, since recurrent semantic relations can only be estimated reliably when the relevant forms occur with sufficient frequency. Vector space models, in particular, benefit from a sufficiently large evidentiary base to stabilize distributions and reduce sparsity, placing practical pressure on corpus size without making scale itself a guarantee of analytical validity. For the purposes of this study, the main selection criterion has been that of a text's relatedness to the theme of love. This choice is motivated by its ubiquity across medieval literary production, which indexes the unfolding discourse, while also supplying enough textual material for recurrent distributional relations to become observable. As a discourse distributed across a wide range of medieval literary forms, love allows for the inclusion of texts, authors, and genres that

would otherwise rarely be brought into a single analytical field while retaining a coherent object of inquiry.

The concrete corpus for this project consists of a total of 121 works divided into 155 text files comprising approximately 4 million tokens, selected through a combination of metadata-driven filtering and domain-informed curation from the METS and MED *Corpus of Middle English Prose and Verse*—which respectively comprise over 800 edited texts in the first case and almost 400 in the second. Because of this large amount of data and the fact that the MED does not have the same level of metadata encoding, the initial selection was based on relevance to the research question and checked against the genre and thematic metadata provided by the METS, which allows for the identification of texts potentially relevant to the discourse of love across a wide range of literary forms. The categories retained for inclusion span a broad spectrum of genres and textual types, including, but not limited to, romance (e.g. Breton lai, crusade romance, Arthuriana), lyric poetry (including Marian lyric and chanson-derived forms), dream visions (including Boethian and love visions), debate poetry, complaints, *fabliaux*, fables, sermons, hagiography, and mystery drama. Additional categories, such as chronicles, courtesy books, and wisdom literature, are also included where they contain sustained engagements with affective or relational themes. Once selected, the texts are organized within a structured dataset in which each entry is encoded through a set of metadata fields reflecting both editorial and analytical considerations. Each text is associated with its title, author (where available), and corpus provenance (MED or METS), alongside a triple dating system that records the traditional composition date, the date of the source (i.e. manuscript witness), and the date of the modern edition where relevant.

The limits of this selection are also visible in works that a more comprehensive study of later Middle English love discourse might reasonably be expected to include. Lydgate's *Fall of Princes*, for example, is absent from the present corpus. In the terminology of computational literary studies, this is an *opportunistic corpus* (Schöch 2017): it is assembled from readily available machine-readable and legally reusable texts rather than sampled from the full population of relevant works. Which texts have been transcribed and released under open licences, and which remain locked in print or proprietary editions, necessarily affected which works could be incorporated. As Bode (2020) has argued, the availability biases that result are constitutive of the evidence a computational corpus produces, and they require characterization rather than correction. From a close-reading perspective, the most consequential absences are texts that, like the *Fall of Princes*, occupy established positions in the courtly love debate and whose inclusion might have altered the distributional relationships the vectorization surfaces. Whether those gaps can be addressed in future iterations of this corpus or must be conceded as limits of the present proof of concept is itself a question the methodology is designed to make visible.

5.2.2 Periodization and dating of texts

The corpus is centred on the period conventionally associated with Ricardian literature, but neither the temporal limits nor the meaning of that category are self-evident. Overall, the texts are all meant to reflect the discourse community of literature in Middle English between the thirteenth and the fourteenth centuries, regardless of the actual geographical origin of the texts. The temporal structure of the corpus is derived from J. A. Burrow's formulation of Ricardian poetry. This was a longstanding problem for specialists in the field, who struggled long to find a single label for the period spanning roughly 1370 to 1400, i.e., the later years of Edward III's

reign and the reign of Richard II (Burrow 1971, 1–3). Burrow's key claim is not simply that this period is historically significant, but that it may be understood as a literary "time section, dominated by a whole system of norms, which no work of art will ever realize in its entirety" (9). The term Ricardian, borrowed from historiography, is thus proposed as a way of naming these emergent norms, allowing the period to be conceptualized as a structured literary field (3–4).

At the same time, I make two important modifications to this framework. First, rather than limiting the Ricardian period to poetry, I broaden it to include prose. Second, although I take the final three decades of the fourteenth century to be the centre of the Ricardian, in agreement with Burrow's periodization, I situate it within a broader tripartite temporal structure: pre-Ricardian, Ricardian, and post-Ricardian. This expansion allows for the identification of both continuities and transformations in the discourse of love, situating the Ricardian moment within a longer trajectory rather than isolating it as a peak of literary achievement. This broader use of the category also differs from Burrow's suggestion that, for some poets who fall chronologically within the Ricardian period, its literary norms "do not apply" (10). I do not treat such departures as grounds for exclusion. If individual texts diverge from the tendencies Burrow identifies—including the stronger focus on human life, the preference for rhyme over alliterative versification, reflexive irony, and heightened textual self-consciousness—those divergences remain part of the literary field the corpus is intended to represent. The periodization is therefore used as an analytical coordinate rather than as a criterion of stylistic homogeneity. Moreover, although the corpus includes works whose narrative or textual sources may date back to the twelfth century or earlier, the periodization adopted here is based on the Middle English instantiation of these texts rather than on their source traditions. In other words, what is taken into account is not the origin of the material, but its articulation within the Middle English

literary system. This choice reflects the study's focus on discourse as it is realized in a specific linguistic and cultural context, rather than on the diachronic history of motifs or narratives. As a result, texts such as translations or adaptations of earlier French or Latin works are dated according to their appearance in the Middle English literary record, even where their sources belong to an earlier period.

The categorization system is adopted from MED's treatment of textual chronology, which its final editor, Robert E. Lewis, terms "double dating," i.e., giving the manuscript date of each text first with its estimated composition date (Lewis and Williams 2007, 35). More precisely, the manuscript date is treated as primary, while the composition date is supplied—where it can be established with reasonable certainty—only if it is significantly earlier (Lewis 2002a, 82). This practice reflects the MED's fundamental orientation towards the material witness: its aim is not to reconstruct an abstract original text, but to describe "what the medieval scribes put down on parchment and paper," thereby privileging the linguistic evidence preserved in manuscripts (82). At the same time, this system acknowledges the limits of our knowledge. As Lewis notes elsewhere, relying solely on composition dates is untenable given the uncertainty surrounding such datings, even if the use of double dating can sometimes obscure the historical development of a word (Lewis 2002b). The resulting compromise is therefore methodological: it preserves both the instability of textual transmission and the partial recoverability of earlier linguistic states. From a linguistic perspective, this position had already been articulated by Sherman M. Kuhn, who observes that a text composed in one period but transmitted in a later manuscript "has obviously lived through some changes in the language," since scribes frequently modernized the forms they copied, albeit inconsistently (Kuhn 1982, 24). The manuscript date thus captures not only the act of copying but also a layer of linguistic updating.

As detailed in Appendix A, chronological details on the composition, manuscript or printed source, and publication of the edition used are recorded alongside title, authorship, corpus provenance, genre, and a broad classification of the treatment of love, as well as a periodization field based on the abovementioned Burrowian categorization. Bibliographic references are provided for edited texts in order to document their provenance and the editorial basis of the corpus, while accompanying notes record text-specific issues such as the choice of base witness, attribution, translation, alternative titles, and cases in which a corpus text was replaced or supplemented by another edition.

5.2.3 Special treatment of sources

In addition to determining the temporal bounds to my corpus, a number of practical choices have been made. Several records were added or replaced because relevant works were absent from the MED corpus, because the available copyright-free editions were inadequate, or because newer documentary evidence justified the use of different editions. These interventions form part of the construction of the concrete corpus rather than an attempt to reconstruct exhaustively every text that might bear upon the discourse of love. The first group of missing works included some of Chaucer's fundamental texts, such as *Book of the Duchess*, *Parliament of Fowls,* and *Legend of Good Women.* For the case of the *Book of the Duchess*, I chose McGillivray's critical edition (1997) based on multiple manuscript witnesses—Bodleian Library MS Bodley 638, MS Fairfax 16, and Glasgow University Library MS Hunter 197—as well as Thynne's 1532 printed edition. *The Parliament of Fowls* is included through *The Riverside Chaucer* (Benson 1987) as well as *The Legend*. Moreover, I added *Patience*, an alliterative religious poem attributed to the Gawain-poet, transmitted in the late fourteenth-century manuscript London, British Library MS Cotton Nero A.x., as it was the only text from that codex

missing in the MED (Shawver 1996). Another important intervention concerns Malory's *Le Morte Darthur*. Although the work is represented in the MED corpus, the version adopted here is Peter J. C. Field's (2017) edition, based on both the Winchester Manuscript and Caxton's 1485 print. The decision not to rely on the MED version reflects the poem's textual tradition tied to the later discovery of London, British Library, Additional MS 59678, which shifted scholarly attention away from the Caxton edition as the sole textual witness, making early editions—such as the one used by the MED—obsolete.

The periodization and linguistic boundaries adopted for this corpus are likewise intended to organize heterogeneity rather than demonstrate homogeneity. In line with Burrow's suggestion, the corpus excludes records in languages other than Middle English, including Latin and French (Burrow 1971, 10). The sole exception to this rule is represented by Old Scots. The consensus regarding what constitutes Middle English for the MED editors was never fully established. As the principal editor of the Middle English Dictionary, Sherman M. Kuhn affirms that "[Middle] English was spoken from the Channel to the foothills of the Scottish Highlands, as well as in parts of Ireland" (Kuhn 1982, 17). The MED editorial board nevertheless ultimately defined its corpus as comprising "[t]exts in English written or copied between 1100 and 1500 in England, Ireland, and Wales," while excluding those "written in Scotland, that is, Middle Scots texts" (Lewis and Williams 2007, 7). According to Kuhn, this decision was framed as "a concession to Scottish patriotism," highlighting the respect toward another lexicographical initiative at the time, i.e., the *Dictionary of the Older Scottish Tongue* edited by William Craigie and his successor A. J. Aitken (Kuhn 1982, 18). In contrast to the MED's approach, however, the METS editorial board states the inclusion of "all forms of literary production significant to

medieval Britain and Ireland's literary and cultural development between the twelfth and sixteenth centuries" (Hahn et al. 2024).

The reservations I have with this choice—to put it with Norman F. Blake—regard "[t]he decision on Scots […] does have implications" at different levels (Blake 2002, 64). First, from an evolutionary linguistic standpoint, Scots emerged from the northern varieties of Middle English, ultimately deriving from the Northumbrian dialect of Old English. To exclude it, therefore, is not simply to draw a boundary between distinct languages, but to interrupt a continuum of linguistic development that remains historically and structurally connected. Second, both the MED and METS include Robert Henryson as part of their corpus of Middle English texts—though they do so in ways that reveal a degree of methodological tension. The MED includes Henryson despite its stated exclusion of Scottish material, thereby introducing an inconsistency between principle and practice, whereas METS incorporates his works while explicitly labelling them as Old Scots in the metadata. In light of the need to reconcile these discrepancies, the present corpus includes Henryson's main works: *The Testament of Cresseid*, *Orpheus and Eurydice*, and the (strongly attributed) *Minor Poems* (Parkinson 2010). This inclusion is motivated by the fact that Henryson's texts participate directly in the same discursive field as their Middle English counterparts, particularly in relation to the literary treatment of love. For the purpose of this analysis, excluding them on a lexicographical boundary would remove material that is linguistically contiguous with Middle English and directly engaged with the literary traditions being modelled.

Finally, longer works—particularly those of Chaucer and Gower—have been subdivided into smaller analytical units based on the boundaries of the single *Tales*. This decision is motivated by the recognition, supported by existing scholarship, that such works do not present a

single, unified intervention in the discourse of love, but rather contain multifaceted and sometimes divergent positions (e.g., see Benson 1986). In Chaucer, each tale articulates distinct and often competing configurations of love—ranging from courtly idealisations to cruder fabliaux—for this reason, treating the work as a single unit would risk flattening these internal differences. Similarly, Gower's *Confessio Amantis* organizes its exploration of love within a confessional and didactic architecture that ostensibly gives rise to a series of exempla that engage with love from multiple moral, philosophical, and narrative perspectives. Considering the textual tradition of the Chaucerian canon, the links to the tales follow F. N. Robinson's edition (1957)—based on Ellesmere and Hengwrt as *codici optimi*—where these elements have been classified as epilogues, therefore, guiding the positioning before or after individual tales. For Gower, the situation is simpler as the *Confessio* presents the dialogue between Amans and Genius as framing strategies before and after each tale.

5.2.4 Linguistic variation and etymological normalization

As is well known, Middle English is "*par excellence*, the dialectal phase of English" (Strang 1970, 224). For this reason, as Michael Piotrowski argues, historical corpora present a fundamental challenge for computational analysis precisely because they lack standardized orthography, meaning that "variation is the norm rather than the exception" (Piotrowski 2012). This diversity reflects the absence of a stabilized orthographic norm alongside the coexistence of multiple, frequently used regional and social varieties, with Middle English in particular permitting authors to write in their own local dialects to a far greater extent than earlier or later periods. As a result, texts such as *Piers Plowman* and S*ir Gawain and the Green Knight* can coexist alongside Chaucer's London-based usage and Henryson's Scots within the same broader literary field without converging upon a limited set of standard written forms. The result is an

exceptional degree of (often dialect-driven) orthographic variability, with common lexical items attested in hundreds of different spellings, shaped by regional pronunciation, scribal training, and inherited conventions. At a more granular level, this includes the coexistence of obsolete graphemes (<þ>, <ð>, <ȝ>, <ꝑ>, <æ>), shifting conventions (e.g. <u>/<v>, <ou> for long vowels), and inconsistent morphological marking, all of which contribute to a highly heterogeneous textual surface.

For the machine, the variation intrinsic to historical texts represents an insurmountable problem: formally distinct spellings may represent the same lexical item, creating data sparsity and functioning as noise by obscuring the recurrent lexical and semantic relations the analysis is designed to recover. Standard NLP pipelines, developed principally for modern languages with stabilized orthographic norms, are poorly equipped to handle this degree of variation. Forms that belong to the same historical lexeme may therefore enter the model as distinct observations simply because their surviving spellings differ. Here, the signal consists of recurrent relations among words and meanings across the corpus, while noise consists of surface variation that causes related forms to be represented as computationally distinct. This distinction is artificial rather than intrinsic to the texts themselves. Dialectal spelling, scribal practice, and morphological variation would themselves constitute signal in a study of linguistic geography or textual transmission; they function as noise here only because the present inquiry is concerned primarily with semantic distribution. Following Dennett's account of patterns as regularities that permit compression, the aim is therefore not to eliminate variation indiscriminately, but to reduce those forms of irregularity that fragment semantically related items and obscure recurrent lexical structure (Dennett 1991; see also Camps, Clérice, and Pinche 2021 regarding the discussion of the same problem for medieval data).

The principal methodological decision made in this study was therefore to represent Middle English lexical forms through the spellings of their Present-Day English etymological descendants wherever such descendants remain available and semantically appropriate. To understand why, it helps to consider the alternative. One conventional hand-and-eye way of reducing orthographic variation is to construct a lookup table that maps attested spelling variants to a shared lexical form, so that the computer recognizes *grene*, *greene*, *grenne*, and *grone* as instances of the same lexeme. While, for a modern standardized language this is straightforward, since the number of variant forms is small, for Middle English, where a single word may appear in dozens of dialectal and scribal spellings across a corpus of this size, the task is prohibitively labour-intensive—it would require manually sorting thousands of forms into their respective lemmas before any computational analysis could begin. The procedure adopted here produces a comparable consolidation without requiring that lookup table to be constructed manually. Rather than building a lookup table by hand, it uses an AI-assisted process to replace each Middle English form with the spelling of its Present-Day English etymological descendant: *grene* becomes *green*, *chirche* becomes *church*, *knyght* becomes *knight*. The result is that forms separated only by dialectal or scribal variation enter the same computational lexical space—which is all the distributional model requires. This consolidation is appropriate to the present analysis because its object is semantic distribution across contexts rather than the grammatical or dialectal variation encoded in the original spellings.

The purpose, in other words, was not to translate the corpus into contemporary English or to modernize its style, syntax, or conceptual vocabulary. It was to provide recurrent Middle English lexemes with a shared orthographic representation so that forms separated only by dialectal or scribal variation could enter the same computational lexical space. Thus, the Present-

Day English form functions primarily as a normalized spelling base rather than as an instruction to read the word according to its modern semantic range. This distinction is important because the procedure differs from modernization in the traditional editorial sense—what makes sense for a computer in order to allow processing would likely make the text semantically unreadable to a human. The resulting texts preserve the word order, syntactic relations, and local co-text of their Middle English sources. A token assigned the spelling of its Present-Day English descendant continues to occur within the grammatical, rhetorical, and semantic environment supplied by the Middle English text. The distributional model consequently learns from those historically specific contexts rather than from the word's contemporary usage. In effect, the corpus remains Middle English at the level of syntax and contextual meaning, while its lexical surface is rendered more regular through Present-Day English spelling.

As Manfred Markus observes, editorial practices involving Middle English have traditionally used *normalization* in ways that overlap with both regularisation—the substitution of rare variants with dominant forms—and modernization—the replacement of unfamiliar forms with contemporary equivalents (Markus 1997, 211). In computational contexts, however, normalization assumes a more strictly methodological role: reducing orthographic variation becomes a practical precondition for comparison and automatic processing (Hickey 1993, 90). Marcel Bollmann's typology of historical spelling normalization, for example, includes rule-based methods that use manually curated substitution lists to collapse variant forms into a common representation (2019, 3885). Such methods can be effective in controlled settings, but their coverage depends upon forms and rules encoded in advance, and they generalize poorly to variants not represented in the system. The procedure developed here may therefore be described more precisely as AI-assisted etymological normalization. It is a constrained token-level

operation in which each Middle English source token is mapped, where possible, to the spelling of a semantically continuous Present-Day English reflex. Where no direct reflex survives, the closest one-word semantic equivalent is supplied. The first operation constitutes the bulk of the process and performs orthographic consolidation by etymological descent; the second is a limited fallback used to prevent obsolete or non-corresponding forms from remaining isolated in the computational vocabulary.

The novelty of the method lies principally in this normalization principle and in its application to a large and highly variable Middle English corpus, rather than in the use of a particular language model. The same operation could, in principle, be carried out manually on a smaller body of texts or implemented through a sufficiently extensive rule-based system. An LLM makes it practicable at corpus scale because it can adjudicate among historically variable forms using contextual information rather than relying exclusively upon a predetermined list of substitutions. The substitutions were generated with Gemini 3.1 Flash Lite model—optimized for translation tasks—which was prompted to perform lexical replacement only and explicitly prohibited from paraphrasing, expanding, omitting, or stylistically rewriting the source (Appendix B). The model was instructed to preserve word order and syntactic structure and to select a direct Present-Day English reflex wherever one remained semantically equivalent. Where no such reflex existed, it selected the closest one-word semantic equivalent. Latin strings exceeding five tokens were treated separately, but otherwise the procedure aimed to preserve a near one-to-one correspondence between source and output tokens. AI therefore functions here as a means of carrying out a constrained substitution procedure at corpus scale rather than as an editor of the texts. Its outputs remain part of an auditable workflow in which the normalized

corpus can be compared against the raw source files and the cumulative effects of processing measured at corpus and document level.

This approach reduces the orthographic fragmentation that would otherwise cause historically related forms to be treated as unrelated tokens. Retaining every Middle English spelling unchanged would leave many lexical items thinly distributed or isolated within the dataset, increasing OOV problems—that is, instances in which rare or previously unseen forms fall outside the model's usable vocabulary—and weakening the distributional representation of forms that are functionally or etymologically related (see Guo, Zheng, and Jiang 2026). Mapping such variants to a shared descendant spelling allows their occurrences to contribute to a common lexical representation. The natural objection to this procedure is that replacing a Middle English form with a Present-Day English descendant might import modern meanings into a historically distinct lexical system. The method mitigates this risk by treating the descendant primarily as an orthographic label. Meaning is not assigned from a modern dictionary in isolation; it is inferred distributionally from the Middle English contexts in which the normalized token occurs.

The procedure nevertheless does not preserve every feature of the source. It deliberately removes evidence carried by dialectal spelling, scribal practice, phonological variation, and some morphological distinctions. It may also collapse forms whose historical senses or derivational histories are not perfectly equivalent, while the use of a one-word semantic substitute where no descendant survives introduces a more explicitly interpretative decision. Polysemy presents a related risk: a single normalized spelling may gather occurrences whose senses differ across texts or periods. These are losses incurred in relation to the analytical purpose of the corpus. The present study does not investigate dialect geography, scribal language, phonology, or orthographic history; it investigates recurrent lexical and discursive relations across texts. For

that purpose, reducing surface variation makes comparison possible while the preserved contexts continue to supply the historical semantic information on which the models operate. The resulting object should therefore not be described as a corpus translated into Modern English. It is a corpus of Middle English texts represented through a partially modern spelling system: orthographically consolidated where etymological continuity permits, minimally translated where it does not, and otherwise left structurally intact. “Modernized” is retained below as a convenient shorthand, but *etymologically normalized* more accurately describes the dominant operation.

Because each stage of textual preparation can introduce alterations that accumulate across a corpus, normalization was audited against the raw files rather than treated as a transparent preprocessing step. To assess the effect of cleaning and normalization on textual noise, I compared each raw file with its corresponding normalized version using token-count deltas. Across the 155 texts, the raw corpus contains 3,799,408 tokens, while the normalized corpus contains 3,754,957 tokens, a net reduction of 44,451 tokens, or 1.17 per cent of the raw total. This reduction is expected because the preparation process removes paratextual material, editorial residue, Latin and French passages, glosses, repeated fragments, and other non-target textual matter. Token delta is therefore used here as an indicator of the extent of cleaning and structural alteration, not as a direct measurement of semantic loss. At the document level, normalization altered text length only slightly. The median difference between a raw text and its normalized version was a reduction of seven tokens. Expressed proportionally, the median change was −0.11 per cent. When the direction of change is disregarded—that is, when increases and decreases are treated simply as changes in length—the median difference was 0.22 per cent. Most normalized files therefore remain very close in length to their raw sources. Larger negative

deltas do not by themselves demonstrate textual breakage; rather, they identify texts in which a substantial quantity of non-target material was removed and which consequently required closer inspection.

The normalization audit also tracked cases in which a Middle English form could not simply be mapped to the spelling of a semantically continuous Present-Day English descendant and instead required a one-word semantic equivalent. Across the 155 texts, the median document contains 2,179 such substitutions, representing 23.68 per cent of its tokens. The frequency of these substitutions varies considerably across the corpus, particularly in longer texts and in works whose vocabulary is more distant from Present-Day English. This figure therefore describes the extent of the intervention rather than the amount of semantic change introduced: some substitutions involve relatively transparent correspondences, while others require a closer approximation because no direct descendant survives. The resulting corpus consists of 155 etymologically normalized text files corresponding to 121 distinct works (see Appendix A for the full list of files and their corresponding works). The corpus-level type–token ratio (TTR = 0.0141) is low, as expected for a corpus of this size, while the median document-level TTR of 0.1625 indicates considerably greater lexical diversity within individual texts. Document length also varies substantially: the median text contains 8,783 tokens, with lengths ranging from 783 to 348,686 tokens and a standard deviation of approximately 42,020. This distribution reflects the coexistence of short works and extracts with much longer texts.

A median sentence length of 25.67 words across a median of 346 sentences per text indicates that sufficient syntactic structure has been retained to support distributional analysis at both word and document levels. The median hapax-legomenon count is 755, with a median hapax rate of 50.24 per cent. In a typical text, therefore, approximately half of the distinct word

forms occur only once as an exact character sequence. These hapax forms include proper names, genuinely infrequent vocabulary, and morphological or orthographic variants that the normalization procedure did not consolidate fully. Their persistence is a reminder that the procedure reduces rather than eliminates historical variation and that the resulting corpus remains an analytically constructed approximation rather than a perfectly standardized linguistic object.

5.3 The space of meanings

Chapter 2 argued that the principal interpretations of courtly love can be usefully reorganized around three explanatory stakes, showing that these accounts reason from overlapping but differently delimited evidentiary bases, with much of the debate repeatedly returning to Provençal troubadour lyric, Old French romance, the Tristan and Lancelot traditions, Andreas Capellanus, Ovid, and a limited body of religious and philosophical writing. The present chapter operationalizes its consequences: the concentration of scholarship upon a recurrent evidentiary base does not invalidate the patterns identified by the three traditions, but it limits the textual field within which those patterns have been observed. The religious strand does not derive from a single archive, but from several overlapping routes that nonetheless converge on the same small cluster of canonical texts. De Rougemont read Provençal troubadour lyric and the *Tristan* legend through the lens of Cathar and Manichaean dualism, producing courtly love as a spiritualized passion defined by distance, suffering, and the attraction of death (de Rougemont 1939); Denomy moved from Avicenna and Neoplatonic philosophy of the soul through the problem of heterodox thought, arriving at Andreas Capellanus' *De amore* as evidence of an Averroist double truth that allowed Christian restraint and worldly passion to coexist (Denomy 1965); Russell read Cathar documents and anti-heresy chronicles against the same *De amore*, treating courtly love as a possible expression of religious dissent and anti-clerical tension

(Russell 1965). The sociological strand built its case on an equally restricted archive of French romance, Provençal lyric, chivalric culture, and feudal society: Paris derived his definition of *amour courtois* from Lancelot's relation to Guinevere in Chrétien de Troyes' *Chevalier de la Charrette* (Paris 1883, 519); Ker moved from the *Chanson de Roland* through Chrétien, Benoît de Sainte-Maure, and Ovid's *Heroides* and *Ars amatoria* in order to frame romance as the literary sign of a broader chivalric transformation (Ker [1897] 1957, 345–50); Wilcox narrowed this line through the court of Marie de Champagne, the later romances of Chrétien, and Andreas Capellanus' *De amore*, defining courtly love through the worship of woman, free love, and sublimation through chivalric deeds (Wilcox 1930, 322–23); and Köhler remained within the troubadour lyric, reading *fin'amor* against feudal and class relations (Köhler 1976). In the psychological strand, Moller combined Provençal and *Minnesang* lyric with the social-demographic record of southern France in order to explain courtly love as frustration, sublimation, and social training (Moller 1959, 48–49); the Lacanian school then read troubadour lyric as a scene of fantasy structure, treating the unattainable lady not as a social person but as an object produced by desire, and treating historical variation as secondary to the transhistorical logic of that structure (Lacan [1959–60] 1986; Gaunt 1995; Frelick 2002; Žižek 2005).

All three draw primarily from Provençal troubadour lyric, Old French romance, and a handful of Latin theoretical texts—material composed between the late eleventh and early thirteenth centuries, in languages and literary traditions at considerable remove from the Middle English cultural milieu that is the object of the present study. The generalizations described above are a structural consequence of working from a restricted archive: each theory confirms itself because it was built from the texts most likely to confirm it. This convergence on an early Occitan and French archive reflects the foundational role that Paris, C. S. Lewis, and their

contemporaries played in establishing the terms of the debate, and the extent to which subsequent scholarship, even when it challenged their conclusions, continued to work within the lineage they had established.

Recognizing the limitations of the evidentiary bases within the scholarship is therefore the point at which computational analysis can intervene productively. By representing the selected Middle English texts within a common lexical and document-level space, the analysis first relocates the explanations developed largely from Provençal, Old French, and Latin materials into a different linguistic, literary, and historical field. It can then examine whether the relations identified by the three traditions remain tenable within Middle English, and whether they retain the same boundaries when observed across a broader and differently constituted evidentiary base. The common computational representation also makes it possible to compare relations that earlier traditions had generally examined separately. What follows is therefore not another explanation of why the three strands differ, but an investigation of how far their distinctions correspond to, overlap within, or diverge from the organization of the corpus for this study.

What computation offers is not a correction of that scholarship but a different model of the evidence: one that represents lexical and semantic relations across a deliberately constructed Middle English corpus within a common distributional space. The corpus's breadth matters, but scale alone does not confer greater authority since its analytical significance derives from how the model constitutes the evidence, formalizes relations among texts, and makes those relations available for comparison.

5.3.1 Operationalizing the impasse

The argument advanced in Chapter 1 produced a specific, testable expectation. If the theological, sociological, and psychological strands each built their accounts from restricted archives, and if those archives made certain conclusions more available than others, then a corpus broad enough to span the full range of Middle English literary production should show something none of the individual strands could see on its own: whether the registers they identified coexist within a single discursive field—and if so, how they relate to one another. Do the theological, sociological, and psychological vocabularies of love occupy separate regions of the corpus, each confirming the strand that foregrounded it? Or do they overlap, cross, and interpenetrate in ways that explain why the debate has remained both persuasive and inconclusive? The distributional framework established in the discourse chapter provides the means of answering these questions empirically. Because distributional semantics represents words as points in a geometric space where proximity encodes similarity of use—words that habitually appear in comparable textual contexts are placed close together, words whose contexts diverge are placed far apart—the coexistence or separation of the three registers can be tested directly: if the theological, sociological, and psychological vocabularies of love occupy distinct and non-overlapping regions of that space, the three strands describe genuinely separate phenomena; if they share a continuous neighbourhood, the strands describe different dimensions of a single discourse, and the impasse is a consequence of perspective rather than of the evidence itself.

The first figure of this section (Fig. 6) provides a preliminary answer. It visualizes the words that appear most frequently in the same textual contexts as *love* across the corpus—i.e., the lexical company that *love* keeps. Each word on the chart is positioned according to how similar its pattern of usage is to *love*'s: words that appear in similar contexts are placed close

together, words whose contexts diverge are placed far apart. What the chart represents, in non-technical terms, is a map of what *love* means in this corpus as measured by the words it habitually accompanies—not a definition imposed by the analyst but a portrait of usage recovered from the texts themselves.

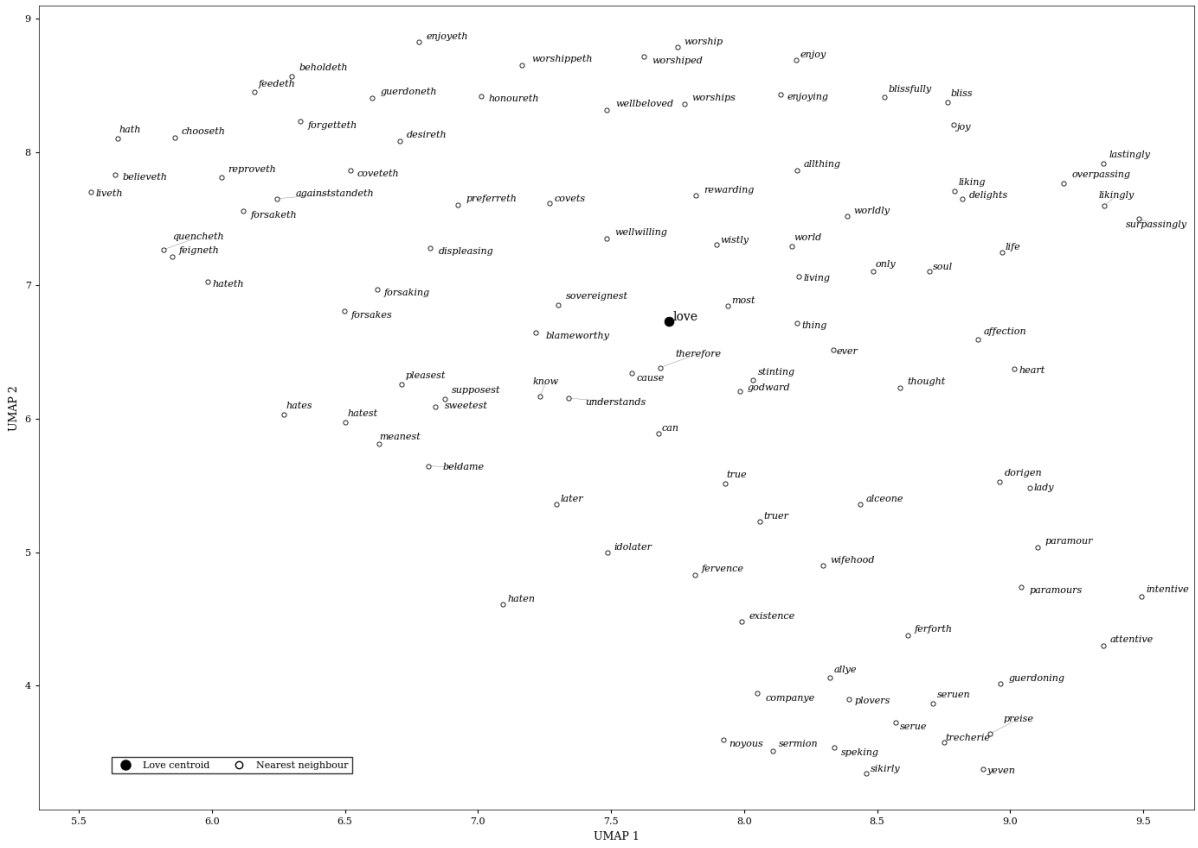


Figure 6: UMAP projection of the 100 nearest neighbours of the love centroid, ranked by cosine similarity. The filled marker is the centroid; open markers are individual neighbour terms

What this map shows is a single, continuous field in which all three scholarly stakes are present simultaneously, organized around a shared discursive centre. The theological-devotional vocabulary is visible in one region: words such as *worship*, *worshippeth*, *believeth*, *forsaketh*, *soul*, *grace*, *bliss*, and *godward* mark the devotional and ethical field through which love is

drawn into relation with salvation, obedience, sin, and divine order. This is the region of the discourse that Denomy, Robertson, and the theological tradition foregrounded—and the model confirms that it is distributionally real. The sociological-courtly vocabulary occupies a contiguous but distinct region: *lady*, *paramour*, *wellbeloved*, *sovereignest*, *guerdoning*, *serve*, and *wifehood* situate love as a social relation structured by rank, gender, reward, and obligation—the field from which Paris, Wilcox, and Köhler drew their explanatory force. The psychological-affective vocabulary is equally present: *heart*, *affection*, *liking*, *desire*, *fervence*, *enamoured*, *sweeting*, and *joy* point toward interiority, longing, pleasure, and emotional fixation—the aspect of love that Moller and the Lacanian tradition treated as central. All three registers appear in the same figure because they coexist in the same corpus. The scholarly strands were each responding to something real—but each was responding to one dimension of a multilayered concept.

The technique that produces this map works as follows. A word embedding model is trained on the full corpus using fastText, an algorithm based on a skip-gram architecture that learns to predict the surrounding context of each word within a fixed window of eight words (Bojanowski et al. 2017). FastText extends earlier embedding models by incorporating subword information, representing each word as a composition of character n-grams: the model decomposes each word into overlapping sequences of three to six characters, learns vector representations for each of those sequences, and constructs the final word vector as their sum together with the whole-word vector. This design choice is particularly consequential for a corpus whose spelling has been regularized but not lemmatized. Morphological and orthographic variants of the same root remain distinct as full word forms, and a conventional embedding model would treat them as unrelated tokens, generating unreliable representations for any form

encountered too infrequently to accumulate stable distributional statistics—the OOV problem discussed in the preceding section. FastText addresses this directly: because the character n-grams of *loveth* and *loved* overlap substantially, the model can construct a meaningful vector for either form even when neither appears frequently enough in isolation to be learned from whole-word co-occurrences alone. The resulting vectors have 256 dimensions and encode each word as a point in semantic space where proximity reflects similarity of usage across the corpus (Mikolov et al. 2013a; Yin and Shen 2018).

The map itself is produced by constructing a centroid—the geometric centre of all the morphological forms of *love* attested in the corpus, calculated by averaging their vectors dimension by dimension. To use a geographical analogy: if each inflected form of *love* (*loveth*, *loved*, *lovyng*, *loves*) were a city on a map, the centroid would be the single location equidistant from all of them—a point that represents the average position of the word rather than any one of its particular appearances. The centroid is not a word that occurs in any text but a constructed analytical object. Its nearest neighbours in the semantic space—the words whose vectors lie closest to it—can then be read as the lexical company *love* keeps across the corpus as a whole. Fig. 6 visualizes the hundred nearest neighbours of this centroid, projected into two dimensions using UMAP (see McInnes, Healy, and Melville 2018). If the centroid neighbourhood shows that the three registers coexist around *love*, the next question is what happens when love is pulled toward a particular scholarly concern. To test this, the analysis uses composed vectors. The idea is simple: if *love* on its own sits at the centre of a multi-register field, what happens to its neighbourhood when it is drawn toward *god*? Toward *lady*? Toward *desire*? Toward *sin*? Each composition adds the vector of an anchor term to the love centroid and asks which words become most proximate. The operation functions as a discursive probe, showing which region of the field

becomes most visible when the shared discourse of love is oriented toward a particular anchor. Fig. 7 shows the results for several such probes. For instance, when the anchor is *god*, the nearest words shift toward devotional and theological vocabulary: *grace*, *Jesus*, *Christ*, *heaven*, *almighty*, *soul*, *worshippeth*, *forgetteth*, *preferreth*, and *maker*. This is the field that the theological strand of the courtly love debate foregrounded—and the model shows that it is distributionally present. When the anchor is *lady*, the field shifts toward courtly-romantic vocabulary: *queen*, *paramour*, *daughter*, *wife*, *enamoured*, *fair*. When the anchor is *desire*, the affective-psychological field emerges: *desirest*, *fervence*, *fervent*, *covet*, *liking*, *heart*, *joy*. Each of these register-confirming compositions behaves largely as the scholarship would predict. The three scholarly strands were not arbitrary since their central claims correspond to real distributional tendencies in the corpus.

The more revealing results, however, come from the register-crossing compositions—probes whose anchors do not belong neatly to any one strand. *Vec(love) + vec(lord)* offers the clearest example. The resulting neighbourhood brings *Jesus* and *lady* into the same semantic region, with *Christ* and *grace* also close by. This result reflects the semantic range of lordship in Middle English, where the word may name Christ or the beloved as a feudal lord. The composed vector makes visible a fusion that the tripartite scholarly division tends to separate artificially: the discourse of love does not keep theological lordship, social lordship, and amatory lordship apart—it allows them to reinforce one another. *Vec(love) + vec(service)* makes this connection even more explicit. Its neighbourhood draws together *serve*, *servant*, *servitude*, *worshipped*, *worshippeth*, *godward*, *guerdoneth*, *rewarding*, and *pleasance*. Service is therefore not simply feudal, religious, or erotic—it emerges as one of

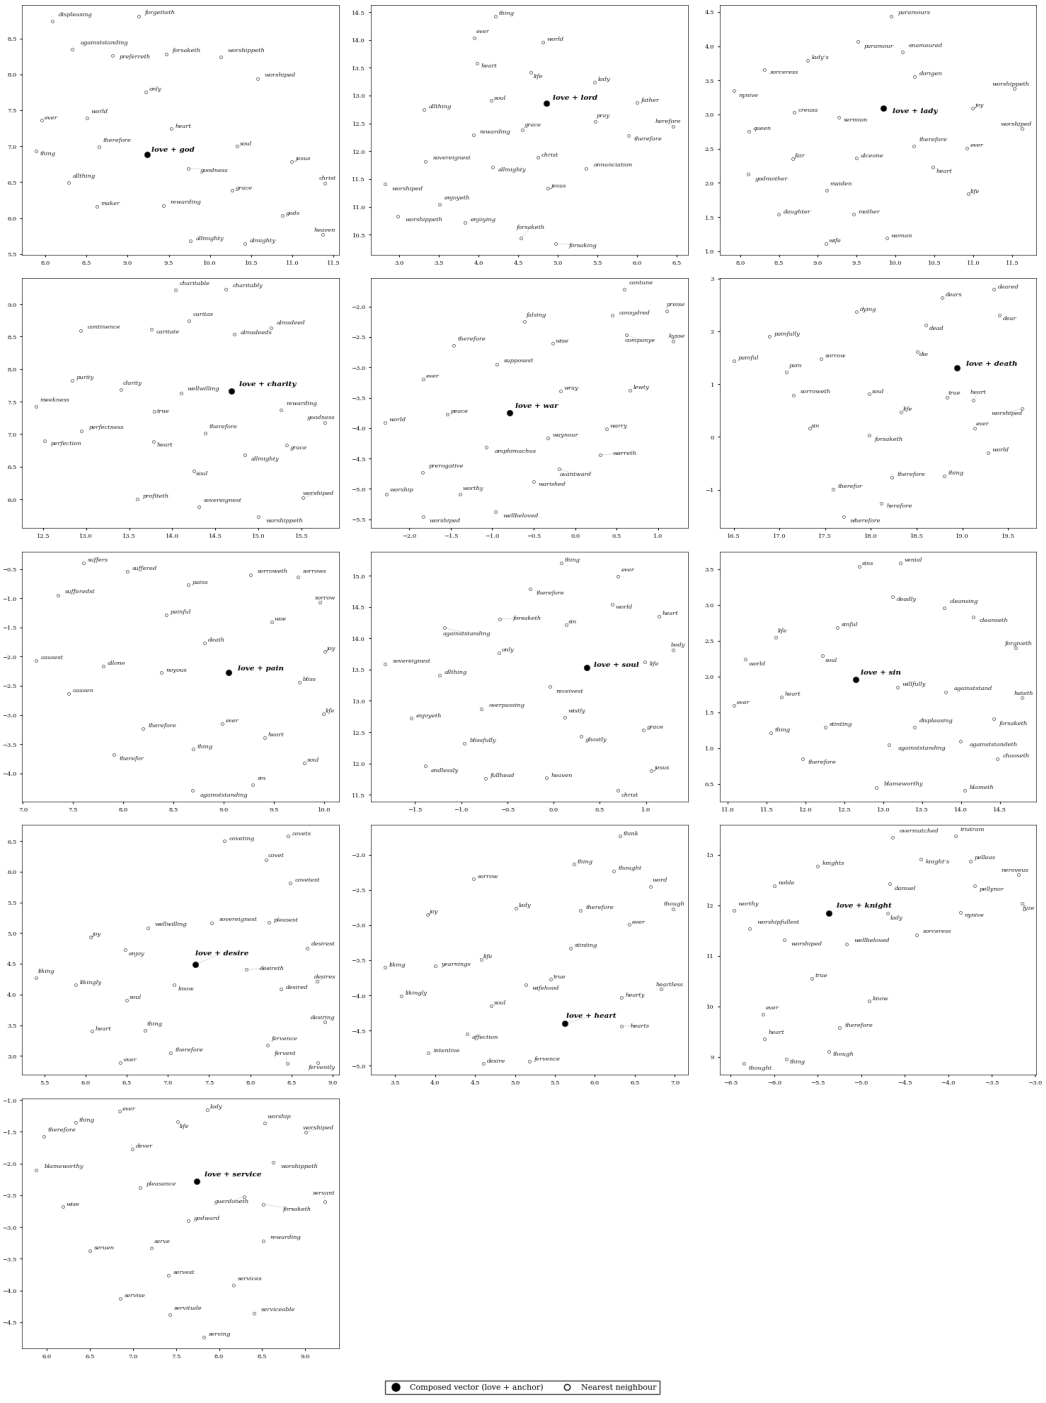


Figure 7: Per-anchor UMAP projections of the 25 nearest neighbours of each composed vector vec(*love*) + vec(*anchor*). Each panel shows an independent two-dimensional reduction; the filled marker is the composed vector, open markers are its nearest neighbours ranked by cosine similarity

the discourse's central mediating terms: a word through which devotion to God, obligation to a superior, and submission to a beloved become mutually intelligible. This is precisely the kind of lexical continuity that selective scholarship tends to divide into separate explanatory traditions.

Vec(love) + vec(knight) similarly pulls love into the chivalric field: its neighbourhood, in fact, combines courtly terms such as *lady*, *damsel*, *paramour*, and *sovereignest* with Arthurian proper names such as *Pellynor*, *Pelleas*, *Tristram*, and *Nynive*, alongside evaluative and religiously inflected terms such as *worshiped*, *worthy*, and *true*. The result corresponds to Ker's sense that romance transforms the heroic world through love, but it also shows that chivalric love is sustained by a mixed vocabulary of worship, worthiness, loyalty, gendered address, and narrative exemplarity. *Vec(love) + vec(pain)* produces an affective-somatic field (*sorrow*, *woe*, *painful*, *death*, *joy*, *bliss*) that recalls de Rougemont's emphasis on suffering and passion but complicates any simple account: the presence of *bliss* beside *pain*, *sorrow*, and *death* suggests that suffering is not outside the discourse of love but one of the forms through which love becomes intensified, moralized, and narratively meaningful. Vec(love) + vec(soul) produces a religious-psychological interiority (*ghostly*, *body*, *grace*, *heart*, *sin*, *Christ*, *heaven*) showing that medieval inwardness is often structured through the soul's relation to sin, grace, and salvation, and that the psychological reading of love is partly right to attend to interiority but that interiority in this corpus is not separable from theology. *Vec(love) + vec(sin)* makes the moral dimension explicit (*deadly*, *sinful*, *venial*, *cleanseth*, *forsaketh*, *againststanding*, *displeasing*, *hateth*, *blameworthy*), confirming that the corpus sustains a moral-theological understanding of love as something that may need to be disciplined, resisted, purified, or redirected. Its vocabulary is

active and agonistic: love is something one may forsake, withstand, cleanse, or blame. The moral register does not condemn love from outside; it participates in shaping love as a field of struggle.

Taken together, the two stages of analysis clarify what the courtly love debate has been arguing about—and why it has remained both persuasive and inconclusive. The first stage shows that the vocabularies of devotion, social hierarchy, and affective interiority—represented by vector-words—all cluster around love within a single distributional field rather than occupying separate and self-contained regions—i.e., the neighbouring centroids. Each of the three scholarly traditions was responding to something real in the texts: each identified a genuine dimension of the medieval discourse, but the same field also shows why no single tradition could settle the debate: a scholar reasoning from any one register will inevitably encounter evidence that confirms the chosen framework while remaining unable to account for the evidence that the other two foreground. The limitation was explanatory priority, given that each tradition granted one region of the discourse authority over the whole.

The second stage takes this further by showing how the field reorganizes itself when love is drawn toward particular concerns, and that the discourse extends into regions the tripartition does not fully capture. *Lord*, *knight*, *service*, *pain*, *soul*, and *sin* function as hinge-terms connecting registers that the inherited scholarship treats as separate: lordship bridges theological devotion and feudal hierarchy; service mediates between divine obligation and amorous submission; soul connects interiority to salvation. These crossings explain why the same passages in medieval texts can sustain theological, sociological, and psychological readings simultaneously—and why a debate conducted from restricted archives, each privileging one register at a time, could not see that coexistence. The computational analysis therefore does not refute the scholarship on courtly love. It explains the impasse. Love in the corpus is theological

without being only theological, sociological without being only sociological, and psychological without being only psychological. Its meaning emerges from the coexistence and movement of these registers across the textual field. That coexistence is what the space of love is designed to model.

5.4 Why Lewis, and why *The Four Loves*

The previous section demonstrated that, at the word level, theological, sociological, and psychological registers coexist around love, and that the discourse extends through hinge-terms whose cross-register operation no single-archive reading could have predicted. This finding reconstructs the semantic field of love without requiring the model to begin from any one of the explanatory positions examined in Chapter 2: the three scholarly registers emerge as coexisting dimensions of a single discursive field rather than as competing definitions of it. For precisely that reason, however, it already provides a partial answer to the thesis's central research question by confirming the assumptions that underlie it, while leaving the larger implications of that answer to be tested at the next level of analysis—whether computational discourse analysis can productively intervene in a longstanding scholarly debate. A more direct intervention requires moving from the semantic neighbourhood of a word to the organization of whole texts and placing the resulting document-level patterns in explicit relation to an established interpretative framework.

To test this claim, the section uses C. S. Lewis's fourfold taxonomy of love, and the principal finding can be stated in advance: three of Lewis's categories produce recognizable textual formations in the corpus, while the fourth—Affection—does not. Rather than gathering texts organized around familial tenderness or domestic attachment, storge absorbs works whose profiles are insufficiently aligned with the other three categories, functioning as a residual rather

than an independent discursive formation. This asymmetry does not invalidate Lewis's recognition of affection as a form of human experience, but it demonstrates that affection does not organize whole literary texts in the way that romantic devotion, fellowship, or divine charity do. The section that follows presents the evidence for this finding, explains how the document-level analysis was constructed, and prepares the ground for the revised three-category scheme proposed afterward.

As Chapter 2 showed, Lewis's synthesis became the principal reference point for the subsequent courtly love debate, with later scholarship repeatedly developing from, revising, or defining itself against the historical model established in The Allegory of Love. His usefulness here follows from that position within the field rather than from any claim that his account is neutral or comprehensive. Its assumptions about gender, sexuality, marriage, and heterosexual attachment are historically situated, but Lewis remains the framework against which a substantial part of the debate has taken shape. Testing Lewis against the distributional organization of the corpus therefore provides a direct way of examining an inherited interpretive framework at another analytical scale. Yet the categories Lewis developed in *The Allegory of Love*—humility, courtesy, adultery, and the religion of love—cannot be converted straightforwardly into distributional categories. They describe the conventions through which courtly love is represented—the conduct, social relations, and literary postures associated with it—rather than distinct semantic domains against which whole documents can be measured. To operationalize them directly would require translating interpretive claims into categories that the original framework does not itself define, thereby risking the imposition of the result in advance.

The *Four Loves* provides a more tractable point of entry. Chapter 2 established the continuity between its taxonomy and concerns already present in the *Allegory*; what matters here

is that Affection, Friendship, Eros, and Charity are stated as explicit categories of love. They can therefore be represented computationally without first converting the literary conventions of the *Allegory* into categories Lewis did not himself formulate. This does not mean that *The Four Loves* supplied the historical taxonomy through which medievalists debated courtly love; as Chapter 2 showed, that role belonged overwhelmingly to The Allegory. Its value here is methodological. Tissari demonstrates that this fourfold distinction can be brought into contact with linguistic evidence rather than remaining solely a philosophical taxonomy (2002; 2003). The four categories consequently provide a workable proxy through which Lewis's broader organization of love can be tested at document level. The purpose is not to falsify Lewis's model but to interrogate and complicate it: to ask whether the four categories correspond to coherent regions of document space, whether texts associated with each category resemble one another consistently, and whether the boundaries Lewis draws are reproduced by the distributional evidence.

Fig. 8 maps each of the 155 texts in the corpus according to how strongly its language aligns with each of Lewis's four categories. Texts placed near one another treat love in linguistically comparable ways; texts placed far apart treat love through different vocabularies, arguments, and thematic orientations. The figure should be read as a map of the corpus organized by the relative pull of Affection, Friendship, Eros, and Charity across whole works.

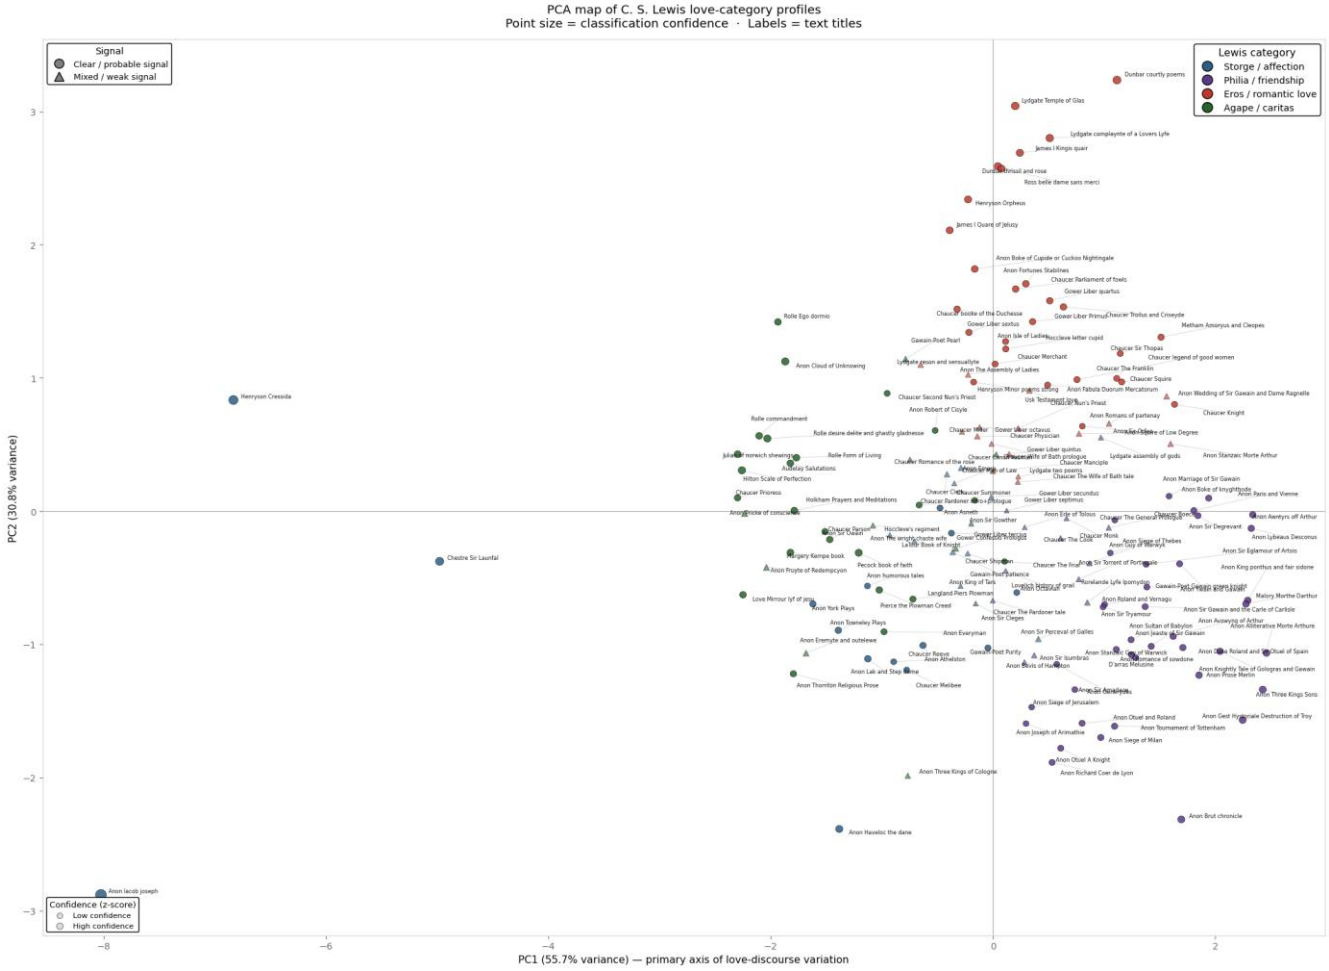


Figure 8: PCA projection of document-level love-discourse profiles under C. S. Lewis's four-category scheme. Each point represents a document projected onto the first two principal components. Colour indicates the assigned Lewisian category; marker size represents z-score-based confidence; and marker shape indicates the strength of the signal

Three of the four categories produce broadly recognizable textual groupings, although with different degrees of internal compactness. The most compact cluster belongs to Charity: Julian of Norwich's *Showings*, Rolle's *Ego Dormio* and *Form of Living*, Hilton's *Scale of Perfection*, and the anonymous *Cloud of Unknowing* occupy a coherent devotional-contemplative region of the document space. Their proximity indicates that divine love, charity,

grace, spiritual transformation, and contemplative devotion constitute a sufficiently consistent semantic orientation to organize whole texts. The texts associated with eros form a second recognizable grouping, concentrated in the upper portion of the projection where courtly complaints, love lyrics, dream visions, and narratives organized around longing for a beloved converge—including Lydgate's *Temple of Glas*, Dunbar's courtly poems, and James I's *Kingis Quair*. Philia occupies a largely distinct region as well, gathering works structured around companionship, fellowship, counsel, and shared moral commitment. The projection therefore supports the principal distinctions in Lewis's framework: divine charity, romantic longing, and companionate fellowship each identify a recognizably different kind of textual discourse, visible not only in individual words but in the aggregate semantic organization of complete works.

Where Lewis's framework succeeds is in these three separations. Where it fails is in the fourth. The distribution of Affection is markedly different from the other three categories. Rather than gathering texts concerned primarily with familial affection, domestic attachment, nurture, or habitual intimacy, Affection is dispersed across almost the entire projection. Its assignments include texts positioned near each of the other categories as well as conspicuous outliers such as Henryson's *Cresseid*, *Sir Launfal*, *Havelok the Dane*, and *Jacob and Joseph*. Storge therefore possesses no equivalent textual centre. It does not identify a coherent region of document space so much as absorb works whose profiles are insufficiently aligned with Eros, Friendship, or Charity.

This scatteredness suggests a limitation not necessarily in Lewis's recognition of affection as a form of human love, but in its usefulness as a category for organizing the literary landscape. Familial tenderness, familiarity, nurture, and domestic belonging may be present within a work without supplying the dominant discourse through which that work represents

love. Unlike romantic longing or divine charity, affection operates as a relational element distributed across several genres and registers rather than as an independently organizing textual formation. The corpus therefore presents an asymmetrical structure: Charity constitutes a comparatively coherent discursive formation; eros and philia each identify recurrent orientations; while Affection disperses across a more fluid field of secular, affective, and social relations without coalescing into a region of its own. The document-level evidence therefore invites a refinement of Lewis's taxonomy: not the rejection of Affection as an experience, but the reconsideration of its capacity to function as a distinct category of literary discourse. The revised three-category scheme proposed in the following section addresses this problem directly: it introduces the register Lewis's framework structurally excluded, absorbs storge's relational elements into the registers where they actually operate, and preserves the distinctions the corpus sustains cleanly.

5.4.1 How the document-level analysis was constructed

The analysis moves from word-level vectors to document-level representations through text embeddings, using a different model from the fastText embeddings employed in the previous section. Whereas fastText learns 256-dimensional word vectors from corpus-internal patterns of co-occurrence, this stage uses e5-mistral, a 7-billion-parameter text-embedding model based on the Mistral architecture and trained for semantic retrieval (Wang et al. 2024). Each encoded passage is represented by a vector of 4,096 dimensions. Sentence-BERT provides an earlier precedent for representing sentences and passages in a shared semantic space in which cosine similarity can be used to compare their semantic content (Reimers and Gurevych 2019), although e5-mistral itself does not belong to the Sentence-BERT architecture.

The model does not interpret texts in the sense that a literary critic does. It measures the contextual similarity between passages—in other words, two stretches of prose whose vocabulary, argument structure, and thematic orientation resemble one another will produce similar vectors, regardless of whether they share specific keywords. The reliability of this measurement can be assessed by checking whether the resulting document representations place texts near one another that scholars would independently recognize as semantically related. If Julian of Norwich's *Showings* and Rolle's *Ego Dormio* occupy similar positions while Chaucer's *Miller's Tale* remains distant from them, that relation derives from their encoded semantic profiles rather than from a prior assignment of the texts to devotional or comic genres.

Each of the 155 texts is divided into sentence-aware chunks of approximately 3,500 tokens. The chunks are encoded separately as 4,096-dimensional vectors, normalized, and averaged to produce a single document-level representation. The resulting vector represents the semantic profile of the document as a whole. The analysis is therefore not determined by the frequency of *love* or any other individual lexical item: it incorporates the broader semantic environment in which the discourse of love occurs, including devotion, service, desire, kinship, sin, grace, marriage, suffering, and social obligation.

Lewis's categories are represented in the same embedding space. For each of the four categories—Affection, Friendship, Eros, and Charity—I composed five short seed sentences characterizing the category from related perspectives. One formulation remains close to Lewis's own account, while the others describe the category through general experience, literary-historical vocabulary, moral or theological terms, and social or relational terms. For example, one seed sentence for Eros reads: *Romantic love characterized by longing, desire, passionate attachment, courtly service, suffering, and the idealization of the beloved*. One for Charity reads:

Divine love expressed through charity, grace, mercy, forgiveness, sacrifice, and spiritual transformation. Each seed sentence is encoded with the same model as the documents, and the five vectors belonging to each category are averaged to produce a single normalized category centroid.

Using several formulations reduces the extent to which the representation of a category depends upon the wording of any one sentence. The centroid represents the shared position of the five descriptions rather than a single formulation taken in isolation. The seed sentences should therefore not be understood as prompts asking the model to classify the corpus. They are embedded texts whose vectors provide four reference points against which the document representations can be compared. E5-mistral uses a query-document retrieval objective and requires inputs to be framed according to its retrieval instruction format. The same encoding conditions are consequently used for the document chunks and the category descriptions so that their vectors remain directly comparable.

For every document, cosine similarity is then calculated between its 4,096-dimensional representation and each of the four category centroids. This produces a four-value profile consisting of one similarity score for Affection, one for Friendship, one for Eros, and one for Charity. Two distinct operations are then performed on this profile, and Fig. 8 represents both simultaneously. The first concerns the relative position of the documents. PCA takes the complete four-value profiles and reduces them to two dimensions, determining where each document appears on the map. The second concerns categorical assignment. Each document is assigned to the Lewisian category with which it has the strongest relative similarity, and this assignment is represented by colour. Position and colour therefore arise from different operations on the same underlying similarity profile: PCA uses the relations among all four scores to

position the document, whereas the colour records which individual category provides its strongest assignment.

The strength of that categorical assignment is represented separately from both position and colour. Marker size indicates the z-score-based confidence of the assignment, while marker shape distinguishes stronger from weaker or more uncertain signals. A text may therefore be assigned the colour of one category while lying close to texts assigned another, or may receive a categorical assignment whose weak signal is indicated by its marker. Such cases are analytically important because they reveal where the four Lewisian categories do not separate the corpus consistently. The figure thus combines spatial position, categorical assignment, and uncertainty without treating them as the same measurement.

PCA is a linear dimensionality-reduction technique that identifies the combinations of variables accounting for the greatest variation among observations (Jolliffe 2002). In the present analysis, it is applied to the four-dimensional similarity profiles rather than directly to the original 4,096-dimensional document embeddings. Under the Lewisian scheme, each document therefore enters PCA with four values, one for each category, which are reduced to two principal components for visualization in Fig. 8. Unlike the UMAP projection used in the preceding word-level analysis, the resulting figure represents variation specifically in relation to the four Lewisian concepts.

The horizontal and vertical zero lines divide this projection into four quadrants, but the quadrants have no classificatory meaning of their own. They arise from the positive and negative values of the first two principal components and provide a convenient way of describing relative locations within the figure. The upper-right quadrant, for example, is not intrinsically an Eros region; it becomes relevant only if documents classified as Eros happen to concentrate there.

Likewise, the intersection of the two axes at zero does not indicate equal similarity to the four categories. It marks the centre of the PCA coordinate system. The orientation of the axes is also mathematically arbitrary in the sense that the sign of a principal component may be reversed without changing the relationships among the documents. The interpretation therefore depends on relative proximity, separation, and overlap among points, and on how consistently those spatial relations correspond to the colours and confidence markers assigned by the separate classification procedure.

5.5 Modelling the love discourse

The refinement proposed here does not simply remove Affection and preserve the remainder of Lewis's scheme; it changes the principle of classification. Lewis distinguishes kinds of love primarily according to their objects, relations, and phenomenological qualities. The revised model instead distinguishes the literary-discursive registers through which texts organize, represent, and evaluate love. It reorganizes the corpus around three coordinates: the secular-erotic, the courtly-romantic, and the theologico-philosophical. The corpus and its document embeddings remain unchanged from the preceding analysis. The same 155 texts are encoded under the same conditions into the same 4,096-dimensional representations. What changes are the conceptual reference points against which those representations are compared. Each revised category is again characterized through five seed sentences covering definition, social relation, phenomenology, literary manifestation, and moral orientation. Their embeddings are averaged to produce a single normalized centroid, following the same multi-seed procedure used for Lewis's four categories. The comparison therefore tests whether a different organization of the discourse produces a more intelligible account of the same document-level evidence, rather than generating a new semantic representation of the corpus.

The revised categories operate at a different analytical level from Affection, Friendship, Eros, and Charity. They do not deny the presence of those forms of relation within particular works. They ask instead how such relations are incorporated into broader textual formations. Affection, for example, remains legible wherever kinship, familiarity, nurture, or domestic attachment occur, but those features need not constitute an independent document class. They may participate in courtly, secular, or theological treatments of love while being organized by a larger discursive register. The same principle applies to the categories that produced clearer profiles under the Lewisian scheme. Features associated with Eros and Friendship contribute to a courtly-romantic register structured by longing, fidelity, service, fellowship, social hierarchy, and chivalric obligation. The material associated with Charity is retained within a broader theologico-philosophical register encompassing charity, contemplation, moral order, spiritual transformation, and philosophical reflection on the proper direction of desire. The change therefore does not simply rename Lewis's categories. It asks which distinctions are useful for describing the organization of whole literary texts rather than the phenomenology of particular human attachments.

The most substantial addition is the secular-erotic register, which gives independent status to a textual formation that Lewis's fourfold taxonomy does not isolate. It encompasses the ribald, appetitive, transactional, and desacralizing dimensions of love: fabliaux, comic tales, and bawdy narratives in which bodily desire, sexual exchange, marital conflict, deception, and appetite become organizing principles of the text. Lewis defines Venus as "the carnal or animally sexual element within Eros" (1960, 106) and calls her "a partly comic spirit" (114). He therefore recognizes bodily and comic sexuality, but retains it as a constituent of Eros, whose defining condition remains that of "being in love." The difficulty appears when appetite rather than

romantic attachment organizes a text's language, action, and evaluation of desire. Venus receives no independent status within the fourfold taxonomy.

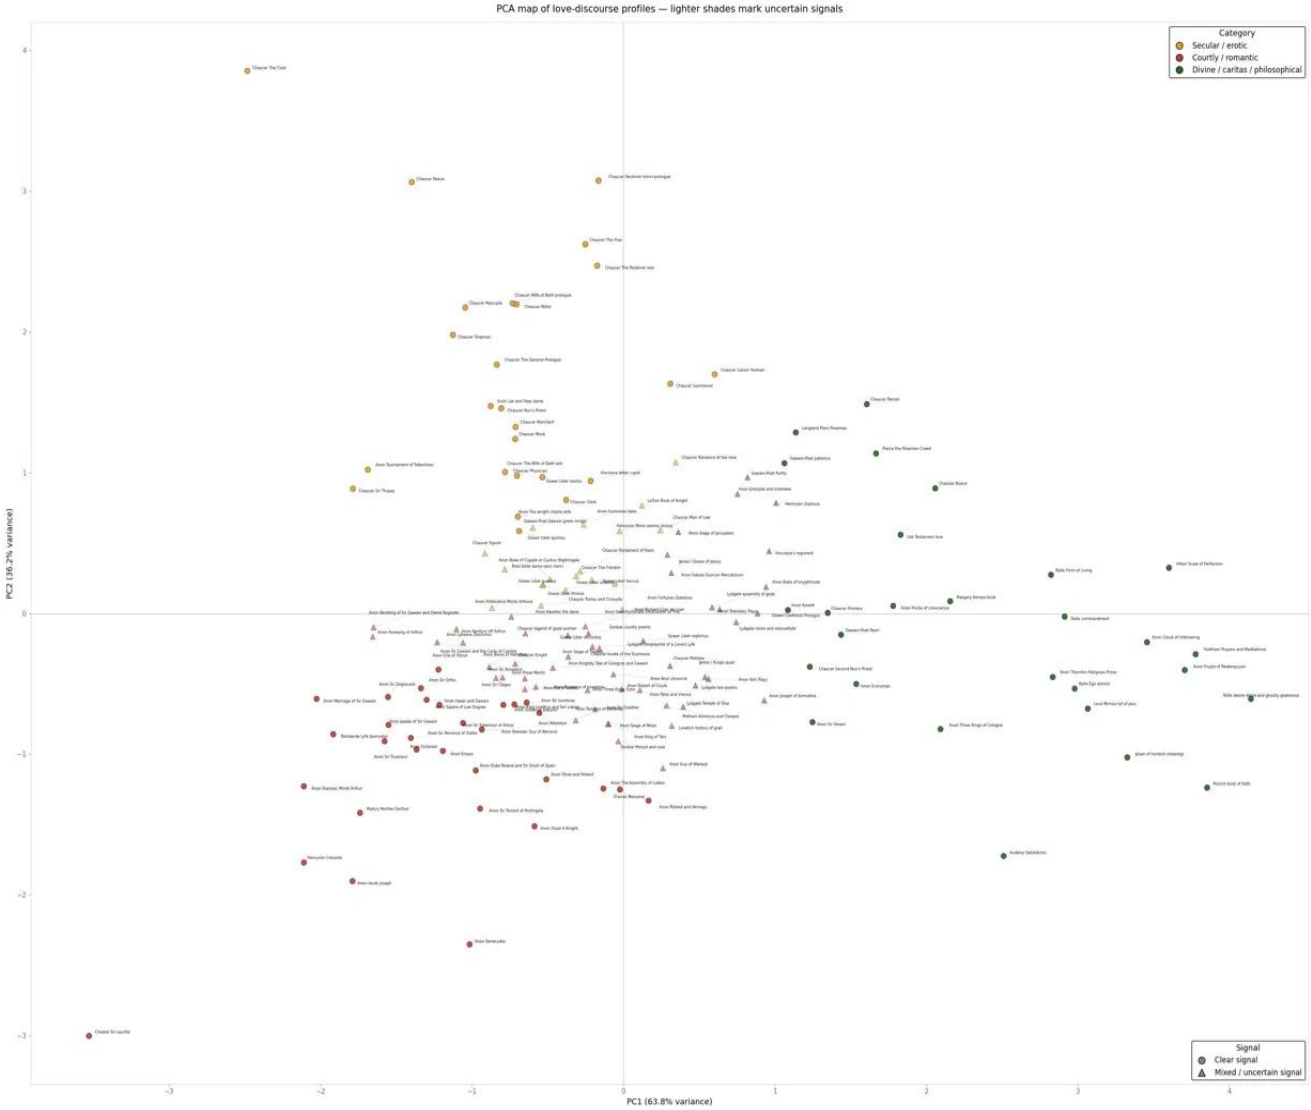


Figure 9: PCA projection of document-level love-discourse profiles under the revised three-category scheme. Colour indicates the category of strongest similarity; position represents the PCA projection of the complete three-score profile; lighter triangular markers identify mixed or uncertain assignments, while circular markers indicate clearer signals

The secular-erotic category is intended to make that distinction available at document level. It describes texts organized around appetite, comedy, transaction, or sexual exchange without requiring those features to be subordinated to the idealization, longing, and self-transformation that Lewis treats as characteristic of Eros. The question posed by Fig. 9 is therefore whether these revised coordinates produce a more interpretable organization of the corpus than the four-category Lewisian scheme. The diagram also combines several visual encodings that should not be conflated. Each document is compared with the three revised category centroids, producing a three-dimensional profile of cosine-similarity scores. PCA operates on the complete three-score profile and determines the relative position of each document in the two-dimensional projection. Colour is assigned separately from that projection according to the category with which the document has the strongest similarity. Marker form records the strength or uncertainty of that assignment. Position, colour, and uncertainty therefore represent different aspects of the same underlying profile.

In this implementation, the three standardized profile values are linearly dependent, so only two principal components carry variance. PC1 accounts for 63.8 per cent and PC2 for 36.2 per cent, together exhausting the observed variation. Unlike Fig. 8, the two-dimensional projection therefore preserves the complete variation of the revised three-category profile. Each relation visible among the documents in the full profile space can consequently be represented in the figure without loss through dimensional reduction. The apparent orientation of Fig. 9 should not be compared literally with that of Fig. 8. The sign and orientation of PCA axes are arbitrary: a principal component may be reversed without altering any of the underlying relations among documents. Separate PCA solutions may therefore appear reflected or inverted even when their relational structures remain comparable. The fact that a group appears on the left in one figure

and on the right in another consequently carries no interpretative significance. What matters is the proximity and separation of documents within each projection, the degree of overlap among the category profiles, and the relation between those spatial patterns and the categorical assignments represented by colour. Fig. 9 should therefore be read as a map of recurrent relations among texts within the discourse of love, rather than as a taxonomy of discrete types. Its horizontal axis captures the broadest opposition represented in the revised profile: toward one pole lie texts characterized primarily by appetite, bodily comedy, and social or sexual transaction; toward the other lie texts characterized by divine charity, contemplation, moral order, and spiritual transformation. The second component differentiates the secular side of that opposition more clearly, separating courtly-romantic treatments organized around longing, service, idealization, and chivalric obligation from more explicitly erotic, comic, or transactional treatments.

The revised model is more interpretable not because it produces perfectly discrete clusters, nor because the colours simply appear neater on the page. Its advantage lies in the clearer correspondence between the proposed discursive distinctions and recognizable textual formations in the corpus, while mixed or uncertain documents remain visible as intermediate cases. Theologico-philosophical texts form the most concentrated region; courtly-romantic texts occupy a recognizable but more heterogeneous formation; and secular-erotic texts acquire a position that was not independently available within Lewis's scheme. The remaining overlaps are therefore part of the interpretation rather than failures to be removed. The theologico-philosophical texts form the clearest concentration. Julian of Norwich's *Showings*, Rolle's *Ego Dormio*, *Form of Living*, and *Commandment*, Hilton's *Scale of Perfection*, the anonymous *Cloud of Unknowing*, and the *Holkham Prayers and Meditations* all occupy a tightly grouped region of

the projection. Their proximity reflects a recurrent semantic organization around grace, mercy, contemplation, purification, sacrifice, divine union, and the ordering of desire. This concentration is consistent with the earlier Lewisian analysis, in which Charity already produced the clearest document-level distinction. The revised model therefore preserves a relation that the previous classification represented successfully rather than replacing it merely for the sake of a new taxonomy.

At the opposite pole lie texts whose treatment of love is neither courtly-idealized nor devotional-charitable: Chaucer's *Cook's Tale*, *Reeve's Tale*, *Nun's Priest's Tale*, *Miller's Tale*, *Shipman's Tale*, *Manciple's Tale*, the *Friar's Tale*, and the *Pardoner's Tale*. These works are not simply texts that Lewis's framework fails to recognize. They foreground a recurrent sexual and social discourse that his fourfold classification does not isolate as an independent category. Their concentration under the secular-erotic coordinate clarifies one source of the instability visible in Fig. 8: material previously distributed among categories such as Affection or Eros can here be represented according to the comic and uncourtly, that seem to organize the texts.

The courtly-romantic texts occupy a third recognizable region. *Troilus and Criseyde*, the *Knight's Tale*, Lydgate's complaint poems, and the Arthurian romances cluster away from both the theologico-philosophical and secular-erotic poles. The Arthurian group—the *Alliterative Morte Arthure*, *Sir Gawain and the Green Knight*, *Ywain and Gawain*, *The Avowyng of Arthur*, *The Wedding of Sir Gawain and Dame Ragnelle*, and the other Gawain-associated texts—appears alongside non-Arthurian romances such as *Sir Orfeo*, *Sir Degrevant*, *Sir Launfal*, and *Emaré*. Their proximity suggests that the relevant relation is not Arthurian subject matter alone but a broader romance register in which desire is articulated through service, social hierarchy, gendered conduct and fidelity.

5.5.1 Individual authors across the discursive space

The preceding paragraphs described the map's three broad formations and the textual formations gathered at each. The projection also reveals how individual authors and works are distributed across the space, and in several cases works by the same author occupy markedly different positions. These distributions are interpretively significant because they indicate that the model is responding less to authorial identity in itself than to the particular discursive organization of individual works.

Chaucer is the most widely distributed author in the corpus. His texts span almost the entire projection, from the *Cook's Tale* at the secular-erotic extreme to the *Second Nun's Tale* and *Parson's Tale* on the theologico-philosophical side, with the *Knight's Tale*, *Troilus and Criseyde*, the *Franklin's Tale*, the *Wife of Bath's Prologue*, and the *Merchant's Tale* distributed across the middle. The breadth of this distribution is consistent with the generic and discursive range of Chaucer's corpus. The *Canterbury Tales* cannot be assimilated to any single tradition of love: the collection places competing representations of desire, marriage, fellowship, charity, appetite, and moral obligation into sustained proximity. The *General Prologue*, however, does not occupy a neutral centre from which these positions radiate; it falls on the secular side of the projection and therefore participates in one region of the field rather than standing outside it as a purely framing document. What matters at the level of authorship is consequently Chaucer's dispersion across the space rather than the location of any single representative text.

Gower presents a different pattern. The books of the *Confessio Amantis* occupy a comparatively compact but internally mixed area around the central region of the projection. The figure does not show a simple progression in which the earlier books begin in the courtly-romantic region and the later books move steadily toward theologico-philosophical love. Their positions instead suggest a repeated interaction among amorous, ethical, political, and didactic

registers across the poem. This distribution is compatible with the *Confessio*'s larger structure, in which Amans' predicament is repeatedly subjected to moral and exemplary interpretation, but the projection is better understood as showing that coexistence than as tracing a chronological movement through the books.

Henryson's *Testament of Cresseid* sits in the courtly-romantic region but at an unusual distance from the main romance cluster—isolated in the lower left, far from the Arthurian romances and the Chaucerian courtly texts. The Testament engages with courtly love precisely by narrating its destruction: Cresseid's punishment by the planetary gods after her abandonment of Troilus is not a celebration of the courtly code but a commentary on its collapse. Its position suggests that participation in a discourse does not require affirmation of its values. The distance between the poem and the main romance cluster is consistent with a critical relation to the courtly register rather than an absence from it. Henryson's *Orpheus*, by contrast, occupies a substantially different part of the projection, toward the theologico-philosophical side and with a more mixed signal. The separation between the two works provides a particularly clear example of the work-specific character of the model: shared authorship does not impose a stable semantic position when the texts organize love through different discursive relations.

Lydgate is distributed across the projection, but not in the clean generic division that might initially be expected. His complaint poems do not form a compact courtly-romantic cluster: *Temple of Glas* and the *Complaint of a Lover's Life* lie much closer to the contested central region, while *Reason and Sensuality* moves more clearly toward the theologico-philosophical side. This distribution is itself revealing: works identifiable through close reading as courtly complaints may nevertheless draw sufficiently upon moral, allegorical, philosophical, or devotional vocabularies to produce mixed document-level profiles. Lydgate therefore

illustrates the distinction between generic identification and discursive organization: a text may participate recognizably in a courtly literary tradition without being semantically confined to the courtly-romantic coordinate.

A comparable differentiation appears in the works associated with the Gawain-poet. Sir *Gawain and the Green Knight* lies with the courtly-romantic material, whereas *Pearl* and the religious poems occupy the theologico-philosophical side of the projection. Pearl lies closer to the central region than the most strongly contemplative texts, which is consistent with its combination of affective grief—the grief of the dreamer—and theological argument—the heavenly consolation of the Pearl-maiden. Its profile is nevertheless more clearly theologico-philosophical than evenly suspended between the two registers. The contrast among these works again shows that shared authorship, and even shared manuscript context, does not require a common position within the discursive space. The same qualification applies to texts previously described as occupying the contested middle.

Piers Plowman lies on the inner edge of the theologico-philosophical formation rather than at a midpoint between courtly-romantic and theological love. Its position is nevertheless closer to the centre than that of the strongly devotional and contemplative texts farther to the right, which is consistent with a poem that repeatedly brings charitable, social, moral, and worldly forms of love into negotiation. Chaucer's *Parson's Tale* likewise occupies the theologico-philosophical side without falling within the more tightly concentrated mystical and contemplative group. Its position therefore suggests a distinction within that larger register between penitential and doctrinal discourse and more explicitly contemplative treatments of divine love. The relative stability of writers such as Rolle, Hilton, and Julian of Norwich provides a useful contrast. Their texts remain concentrated within the theologico-philosophical

region, although they occupy different positions within it. The model is therefore capable of producing authorial consistency where the represented works share a comparatively stable discursive orientation. The dispersal of Chaucer, Henryson, Lydgate, and the Gawain-poet cannot consequently be explained simply as instability in the model; it corresponds to differences among the works being represented.

Collectively, these distributions suggest that the most useful unit of interpretation is not the author understood as a single semantic identity, but the individual work as a particular organization of discursive relations. Authorship remains interpretively relevant, but it does not determine where a text must appear in the space. The revised categories consequently operate as coordinates rather than containers: individual works may occupy concentrated regions, move between them, or combine several registers. The intermediate and unexpected positions are especially productive because they identify concrete cases in which the relations made visible by the model can be returned to close reading. In this sense, the projection does not resolve the interpretation of individual texts; it identifies where further interpretation becomes necessary.

5.5.2 What the revised map adds to the debate

The three categories are therefore not proposed as a replacement taxonomy of kinds of love, but as coordinates for describing the organization of love at the level of literary discourse. This distinction clarifies what the revised model changes in Lewis. His categories remain intelligible as descriptions of relational and affective experience, but they do not discriminate whole texts with equal consistency. Eros and Charity produce the clearest opposition in the four-category model, while Friendship substantially overlaps with the charitable region and Affection remains more diffusely distributed. The problem exposed by the map is consequently one of

analytical resolution: distinctions that remain meaningful at the level of human relations need not correspond to equally distinct formations at the level of literary documents.

The computational analysis therefore tightens Lewis's categories rather than falsifying them. It shows which distinctions survive the change of analytical scale and where finer distinctions become difficult to sustain when the same categories are applied across a heterogeneous corpus. Friendship and Affection do not disappear from the texts; their relational features remain visible within the broader formations through which those texts organize love. Conversely, the secular-erotic register gives independent visibility to a literary formation that Lewis recognizes through Venus but subordinates to Eros. The revised scheme thus compresses distinctions that the corpus does not consistently maintain at document level while introducing a distinction that the literary evidence makes difficult to ignore.

This refinement also bears directly on the longstanding courtly-love debate. The question is no longer which single definition of love best describes the Middle English evidence, but how texts position romantic attachment, bodily desire, social relation, and divine charity in relation to one another. The revised map makes those relations more legible while preserving the intermediate cases in which they overlap. Its contribution is therefore not the discovery of a new taxonomy of medieval love, but a more precise account of where an inherited taxonomy succeeds, where its distinctions become redundant, and where the corpus requires another distinction to be made.

In this sense, the result answers the methodological question with which the thesis began. Computation contributes to an established humanistic debate not by replacing its interpretations with a supposedly more objective classification, but by representing their distinctions in another form and showing where those distinctions can be sharpened. The intervention in Lewis is

therefore also the proof of concept: the computational model produces a result that remains answerable to the existing debate while making a distinction within it more intelligible.

CHAPTER 6: CONCLUSION

> To find signals in data, we must learn to reduce the noise—not just the noise that resides in the data, but also the noise that resides in us. It is nearly impossible for noisy minds to perceive anything but noise in data.
>
> —Stephen Few, *Signal: Understanding What Matters in a World of Noise*

6.1 Implications for this study

This thesis began with a methodological question: under what conditions can computational methods intervene productively in an established humanistic debate while remaining answerable to the interpretive traditions they seek to extend? The debate over love in later Middle English literature supplied the test case for that question. Its longstanding interpretive divisions, the different bodies of evidence on which those divisions had been built, and the continuing influence of C. S. Lewis made it possible to test whether computational modelling could enter an established scholarly problem and return results that remained intelligible within its terms. The analysis in Chapter 5 is the proof of concept for that methodological argument.

The maps produced in the preceding chapter are not ends in themselves, nor are they meant to replace literary interpretation with visualization. On the contrary, once the corpus is projected into a common distributional space, the position of each work in relation to others becomes newly available for careful interpretive attention. Proximities, separations, border cases, and outliers all invite precisely the kind of text-sensitive reading that computational models cannot perform on their own. A work's location in the field does not explain itself; it must still be interpreted in relation to genre, rhetoric, narrative situation, and historical context. In this sense, the computational model does not supersede close reading but creates a new occasion for it, making visible relations among texts that a reader may then examine more carefully. A

complete commentary of those positions would require a different project, built around the close reading of the individual texts and clusters identified by the model. The present study uses them instead to test what kind of intervention computational modelling can make in the scholarly debate from which the categories themselves were derived.

Within that methodological trial, the thesis makes two specific contributions to the medieval love debate. Chapter 2 argues that the persistent division among philosophico-theological, sociological, and psychological interpretations can be understood partly through the overlapping but differently delimited bodies of evidence from which each has constructed its object. Chapter 5 then tests one framework that has been unusually consequential in organizing those interpretations, namely, C. S. Lewis's account of love, which provides the benchmark because subsequent scholarship has repeatedly developed from, modified, or defined itself against his synthesis. His categories therefore offer a way of testing an inherited organization of the field rather than selecting an arbitrary taxonomy against which to measure the corpus.

C. S. Lewis's *The Four Loves* treats Affection, Friendship, Eros, and Charity as four categorically distinct forms of love. When this framework is tested against the distributional structure of a corpus of 155 Middle English texts, those four categories do not reproduce themselves as four comparably coherent document-level formations. Affection (storge), in particular, does not gather texts organized around familial tenderness, domestic attachment, or habitual intimacy; instead, it absorbs works whose profiles are insufficiently aligned with the other categories, functioning as a residual rather than an independent discursive formation. Affection itself is nevertheless distributed throughout the corpus. The difficulty lies in treating it as the organizing principle of a document class. The revised three-category scheme addresses this by changing the principle of classification: from a taxonomy of affective kinds to a

distinction among the discursive registers through which literary texts represent and evaluate love. Storge is retained as a relational element but relinquishes the role of an independent document class. The courtly-romantic cluster absorbs the affective features associated with both Eros and Friendship, treating the distinction between friendship and romantic attachment as a difference of emphasis within a shared discursive field rather than a categorical boundary. A secular-erotic register is introduced to name the literary formation Lewis's moral framework structurally excluded—the ribald, the appetitive, the transactional, the desacralizing—because both *The Allegory of Love* and *The Four Loves* presuppose that love is oriented toward some form of the good: courtly love ennobles, charity sanctifies, friendship elevates, affection nurtures. There is no category for love as appetite, deception, bodily comedy, or social transaction. The *fabliaux*, the comic tales, the bawdy narratives treat love as something the Lewisian framework cannot see because it was not conceived with that kind of love in mind. The theologico-philosophical category preserves the distinction corresponding to Lewis's Charity—the one category the corpus sustains without qualification.

That is, on its own, a refinement of an existing framework, not the proposal of a new theory of medieval love. The finding about medieval love, however, is small relative to the apparatus that produced it—four chapters of theoretical, methodological, and epistemological groundwork before a single computational result was reported. That disproportion is the cost of making each stage explicit, and it reflects a condition that will not apply in the same way to subsequent projects. The theoretical apparatus was extensive because the thesis had to address, in sequence, the questions that any such intervention must answer: what is the debate, and where has it left something open? What kind of analysis is appropriate, and why? Does computational intervention replace the interpretive tradition it enters, or can it engage with it on its own terms?

And given the answers to those questions, how should a suitable corpus be constructed, tools applied, and results interpreted? Each question required a chapter because each represents a genuine obstacle to the integration of computational and traditional approaches—an obstacle that, if left unaddressed, would have given any sceptical reader grounds to dismiss the finding before encountering it.

As more work of this kind is undertaken, the theoretical and epistemological groundwork that preceded the analysis here will become less necessary. This is particularly important for the epistemological argument developed through Jockers—that computational modelling adds a descriptive layer rather than replacing what came before, and that the falsification framework through which computational work has sometimes sought legitimacy is a category mistake when applied to humanistic knowledge—is available for other projects to inherit rather than re-derive. The practical issues in corpus construction, the normalization procedure, and the modelling framework are documented in sufficient detail to be replicated and adapted. Others working on different unresolved problems in the analysis of medieval discourse can take the same approach with considerably less preliminary labour than this thesis required, because that labour has been done. The approach also depends on a condition that this thesis had to establish for itself, without the benefit of any prior project having already done so: the alignment of data, method, and research question within a single coherent framework. That alignment was the principal labour of the thesis, and the problems it encountered in achieving it turn out to be recognizable, outside of medieval studies entirely, as instances of a broader entailment, represented by the tacit, underlying issue of data integration.

As Kitchin and Lauriault observe, academic knowledge has been produced for centuries through what can now be recognized as small-data studies, i.e., a kind of research "underpinned

by data produced in tightly controlled ways using sampling techniques that limited their scope, temporality, size and variety" (2015, 463). Small data are characterized by limited volume, non-continuous collection, narrow variety, and—most importantly—the fact that they are generated to answer specific questions rather than harvested as a by-product of other systems. The medieval literary corpus fits this description precisely. It is not large by any computational standard; it is not continuously generated; and it was assembled purposely in response to a research question rather than drawn from an existing data infrastructure. Yet the computational methods applied to it were designed for datasets that are, in Kitchin's terms, "huge in volume, high in velocity, diverse in variety, exhaustive in scope, fine-grained in resolution, and relational in nature" (Kitchin 2014, 1). The gap between the data the methods expect and the data the humanities can provide is not a minor inconvenience.

The process of getting from medieval textual evidence to computational results exposed a set of problems that are worth naming separately from the literary-critical finding, because they are not specific to it. Medieval evidence is scattered across institutions, unstable in its material forms, and available through incompatible editorial representations. Two major projects—the MED and METS—classify the same material under different inclusion criteria, different metadata schemas, and different assumptions about what counts as a Middle English text. The texts themselves survive not as stable data objects but as editorially mediated witnesses whose orthographic forms, chronological identities, and categorical boundaries shift depending on which edition is used and which editorial principles govern its production. None of this resembles the conditions under which computational methods are typically developed and evaluated.

Recent discussions in the field confirm that medieval evidence does not arrive in a form amenable to computational analysis—and that making it amenable is itself a research contribution rather than a technical prerequisite. As Mattison and Ryley observe in their introduction to the *Digital Philology* special issue on big data and manuscript studies, the fundamental questions for medievalists working at scale are not "how many" but "how do we count and why do we count" (2025, 3). Medieval manuscript data is not *Big* by any external standard: even corpora of thousands of items remain small in comparison with the millions of records available for print or the Social Sciences. What makes the data difficult is not its volume but its internal complexity—the fragmentariness of the surviving record, the instability of the material witnesses, the incompatibility of cataloguing traditions, and the foundational questions about manuscripts that remain unanswered (Mattison and Ryley 2025, 7). Loss, as they argue, is not an obstacle to be overcome but a structural condition within which computational work must operate: "big data study does not undo loss, but it provides a set of strategies for thinking with loss" (Mattison and Ryley 2025, 8). The present thesis encountered that condition in Hermanns's distinction between imaginary, virtual, and concrete corpora, which makes clear that the corpus analysed here is a materially bounded construction rather than the discourse itself. Its claims are consequently conditioned by the texts that could be selected, obtained, normalized, and incorporated, as well as by those that remain outside it. Representativeness functions here as a modelling problem whose adequacy must be judged against the research question and the known limits of the concrete corpus.

Warren and Weijer sharpen this point at the level of data design. Their account of the "Imagining History" and "Re-Imagining History" projects for the Middle English prose Brut chronicle demonstrates that digital manuscript corpora possess what they call a "dubious

ontology": their way of being is inseparable from the way we know them, and the metadata elements selected to describe them determine what aspects of the corpus become visible and therefore what the corpus *is* (2021, 115). Data models, in this account, are not neutral containers for pre-existing information; they are objects that produce knowledge—what Warren and Weijer, drawing on Costis Dallas, call "thingformation," taking shape at the intersection of imagination and infrastructure (2021, 116). The present thesis's corpus construction is an instance of exactly this process: each metadata decision shaped what the analysis could see, and the decisions are inseparable from the analytical framework they were designed to support. Camps, Clérice, and Pinche provide the most directly relevant precedent for the practical problem this thesis confronted. Their work on Old French hagiographic manuscripts builds a complete processing chain from digitized manuscript images through handwritten text recognition, word segmentation, abbreviation resolution, normalization, and lemmatization to stylometric analysis—and demonstrates that "each step we pursue might introduce new errors" (2021, 54). The cumulative effect of noise across successive processing stages is not a side effect of computational analysis; it is the central methodological challenge, requiring explicit management at every point.

In the data-integration literature, this is the problem of heterogeneous data integration—the problem of making evidence that is produced by different communities, preserved under different conditions, and classified under incompatible schemas computationally tractable within a common analytical frame. As Putrama and Martinek note in their review of the field, "collecting and integrating such heterogeneous data has been regarded as a significant challenge, mainly due to the unstructured data and variety of data sources," with the heterogeneity of the data creating integration problems across the dimensions of volume, variety, and velocity (2024,

2). Stonebraker and Çetintemel's argument that, for the above reason, "one size fits all" no longer applies to database architectures extends naturally to the methodological domain: the assumption that a single computational pipeline can be applied unchanged across datasets of radically different provenance, scale, and internal structure has, as they put it, "come and gone" (2005, 2). Bensaci et al. (2026) observe that the issue of heterogeneity—especially in relation to recognizing that formally different representations refer to the same underlying semantic entity—remains the principal unresolved challenge.

This thesis does not claim to have solved that general problem. It solved it for one case: the computational discourse analysis of love in Middle English literary texts. Two procedures developed in doing so may be useful beyond that case. The AI-assisted modernizing normalization procedure offers a way of mapping substantial orthographic variation onto a common lexical representation while retaining an auditable record of interventions; its transferability to other historical languages would depend on the capacities of the language model, the character of the textual evidence, and equivalent procedures of verification. The corpus-construction pipeline—from the reconciliation of incompatible editorial traditions through the triple dating system to the treatment of composite works—provides a worked example of small-data integration when historical evidence has been produced under classifications that do not align.

The broader implication follows from the conditions encountered in producing the case study. Scattered evidence, incompatible classification schemas, unstable material witnesses, cumulative noise across processing stages, and the gap between the surviving object and its analytical representation recur wherever historically layered or institutionally fragmented evidence must be brought into a common computational frame. The procedures developed here

are therefore best understood as worked responses to one configuration of that problem. Their value lies in making explicit the decisions by which evidence is selected, normalized, integrated, modelled, and returned to interpretation, so that those decisions can be examined and, where appropriate, adapted elsewhere.

REFERENCES

Abu-Lughod, Lila, and Catherine A. Lutz. 1990. “Introduction: Emotion, Discourse, and the Politics of Everyday Life.” In Language and the Politics of Emotion, edited by Catherine A. Lutz and Lila Abu-Lughod, 1–23. Cambridge: Cambridge University Press.

Akkaya, Aslihan. 2014. “Language, Discourse, and New Media: A Linguistic Anthropological Perspective.” *Language and Linguistics Compass* 8 (7): 285–300. https://doi.org/10.1111/lnc3.12082.

Aledavood, Parham. 2024. “Taking the Middle Road: Reflections on Mixed Methodology within the Digital Humanities.” *Digital Studies/Le champ numérique* 14(1): 1–19. https://doi.org/10.16995/dscn.11069

Baker, Paul. 2010. *Sociolinguistics and Corpus Linguistics*. Edinburgh: Edinburgh University Press.

Bakhtin, M. M. 1981. *The Dialogic Imagination: Four Essays*, translated by Caryl Emerson and Michael Holquist. Austin: University of Texas Press.

Baroni, Marco, Raffaella Bernardi, and Roberto Zamparelli. 2014. “Frege in Space: A Program for Compositional Distributional Semantics.” *Linguistic Issues in Language Technology* 9 (6): 5–110. https://aclanthology.org/2014.lilt-9.5/. Accessed June 23, 2026.

Baxter, Judith A. 2008. “Feminist Post-Structuralist Discourse Analysis: A New Theoretical and Methodological Approach?” In *Gender and Language Research Methodologies*, edited by Kate Harrington, Lia Litosseliti, Helen Sauntson, and Jane Sunderland, 243–255. Basingstoke: Palgrave Macmillan.

Bendel Larcher, Sylvia. 2015. *Linguistische Diskursanalyse. Ein Lehr- und Arbeitsbuch*. Tübingen: Narr Francke Attempto.

Bengio, Yoshua, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. “A Neural Probabilistic Language Model.” *Journal of Machine Learning Research* 3: 1137–55.

Bensaci, Meriem, Mohammed Charaf Eddine Meftah, El-Hadi Meftah, Abdelkader Laouid, and Sajid M. Sheikh. 2026. “Integrating Heterogeneous Data: A Systematic Review of Challenges and Evolution Solution.” In *Proceedings of the 9th International Conference on Future Networks and Distributed Systems (ICFNDS ’25)*, 1186–94. New York: ACM. https://doi.org/10.1145/3789692.3789850.

Benson, C. David. 1986. *Chaucer’s Drama of Style: Poetic Variety and Contrast in the Canterbury Tales*. Chapel Hill: University of North Carolina Press.

Benson, Larry D., ed. 1987. *The Riverside Chaucer*. 3rd ed. Boston: Houghton Mifflin.

Benton, John F. 1968. “Clio and Venus: An Historical View of Medieval Love.” In *The Meaning of Courtly Love*, edited by Francis Xavier Newman, 19–42. New York: State University of New York Press.

Bizzell, Patricia. 1992. *Academic Discourse and Critical Consciousness*. Pittsburgh: University of Pittsburgh Press.

Blake, N. F. 2002. "The Early History of, and Its Impact upon, the Middle English Dictionary." *Dictionaries: Journal of the Dictionary Society of North America* 23: 48–75. https://doi.org/10.1353/dic.2002.0010.

Bloch, Howard R. 1977. *Medieval French Literature and Law*. Berkeley: University of California Press.

———. 1991. "'Mieux vaut jamais que tard': Romance, Philology, and Old French Letters." *Representations* 36: 64–86. https://doi.org/10.2307/2928632.

Bloom, Harold. [1973] 1997. *The Anxiety of Influence: A Theory of Poetry.* 2. ed. Oxford: Oxford University Press.

Boase, Roger. 1977. *The Origin and Meaning of Courtly Love: A Critical Study of European Scholarship*. Manchester; Totowa: Manchester University Press; Rowman and Littlefield.

Bode, Katherine. 2012. *Reading by Numbers: Recalibrating the Literary Field*. London, New York, New Delhi: Anthem Press.

———. 2020. "Why You Can't Model Away Bias." *Modern Language Quarterly* 81 (1): 95–124. https://doi.org/10.1215/00267929-7933102.

Bojanowski, Piotr, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. "Enriching Word Vectors with Subword Information." *Transactions of the Association for Computational Linguistics* 5: 135–46. https://doi.org/10.1162/tacl_a_00051.

Bollmann, Marcel. 2019. "A Large-Scale Comparison of Historical Text Normalization Systems." In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, edited by Jill Burstein, Christy Doran, and Thamar Solorio, 3885–98. Minneapolis: Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1389.

Booth, Alison. 2017. "Mid-Range Reading: Not a Manifesto." *PMLA* 132(3): 620–627. https://doi.org/10.1632/pmla.2017.132.3.620.

———. 2022. "But Why Always the Novel? Midrange Reading Samples of Persons and Texts." *New Literary History* 54 (1): 559–581. https://doi.org/10.1353/nlh.2022.a898321.

Boquet, Damien, and Piroska Nagy. 2018. *Medieval Sensibilities: A History of Emotions in the Middle Ages*, translated by Robert Shaw. Cambridge; Medford: *Polity*.

Bordalejo, Bárbara. 2013. "The Texts We See and the Works We Imagine: The Shift of Focus of Textual Scholarship in the Digital Age." *Ecdotica* 10: 64–75.

Bowker, Lynne, and Jennifer Pearson. 2002. *Working with Specialized Language: A Practical Guide to Using Corpora*. London: Routledge.

Brock, André. 2018. “Critical Technocultural Discourse Analysis.” *New Media & Society* 20 (3): 1012–1030. https://doi.org/10.1177/1461444816677532.

Brooks, Cleanth, and Robert Penn Warren. 1938. *Understanding Poetry*. New York: Henry Holt and Company.

Brooks, Cleanth. [1947] 1975. *The Well Wrought Urn: Studies in the Structure of Poetry*. New York: Harcourt, Brace and Company.

Burrow, J. A. 1971. *Ricardian Poetry: Chaucer, Gower, Langland, and the Gawain Poet*. New Haven, CT: Yale University Press.

Busch, Albert. 2007. “Der Diskurs: Ein linguistischer Proteus und seine Erfassung. Methodologie und empirische Gütekriterien für die sprachwissenschaftliche Erfassung von Diskursen und ihren lexikalischen Inventare.” In *Diskurslinguistik nach Foucault*, edited by Ingo H. Warnke, 141–164. Berlin; Boston: de Gruyter.

Busse, Dietrich, and Wolfgang Teubert. 1994. “Ist Diskurs ein sprachwissenschaftliches Objekt? Zur Methodenfrage der historischen Semantik.“ In *Begriffsgeschichte und Diskursgeschichte: Methodenfragen und Forschungsergebnisse der historischen Semantik*, edited by Dietrich Busse, Fritz Hermanns, and Wolfgang Teubert, 10–28. Opladen: Westdeutscher Verlag.

Calin, William C. 1980. “Defense and Illustration of Fin’ Amor: Some Polemical Comments on the Robertsonian Approach.” *The Expansion and Transformations of Courtly Literature*, edited by Nathaniel B. Smith and Joseph Thomas Snow, 32–48. Athens: University of Georgia Press.

Camps, Jean-Baptiste, Nicolas Baumard, Pierre-Carl Langlais, Olivier Morin, Thibault Clérice, and Jade Norindr. 2023. “Make Love or War? Monitoring the Thematic Evolution of Medieval French Narratives.” Paper presented at *Computational Humanities Research (CHR 2023)*, Paris, December 6–8. https://enc.hal.science/hal-04250657.

Camps, Jean-Baptiste, Thibault Clérice, and Ariane Pinche. 2021. “Noisy Medieval Data, from Digitized Manuscript to Stylometric Analysis: Evaluating Paul Meyer’s Hagiographic Hypothesis.” *Digital Scholarship in the Humanities* 36 (Supplement 2): 49–71. https://doi.org/10.1093/llc/fqab033.

Chun, Jon, and Katherine Elkins. 2022. “What the Rise of AI Means for Narrative Studies: A Response to ‘Why Computers Will Never Read (or Write) Literature’ by Angus Fletcher.” *Narrative* 30(1): 104–113. https://doi.org/10.1353/nar.2022.0005.

Classen, Albrecht. 2002. *Verzweiflung und Hoffnung: Studien zur deutschsprachigen Epik des Mittelalters und zur kommunikativen Gemeinschaft*. Berlin: Walter de Gruyter.

———. 2004. *Discourses on Love, Marriage, and Transgression in Medieval and Early Modern Literature*. Tempe: Arizona Center for Medieval and Renaissance Studies.

———, ed. 2008. *Eroticism and Love in the Middle Ages*. 6th ed. Mason, OH: Cengage Learning.

Cleland, Carol E. 2001. “Historical Science, Experimental Science, and the Scientific Method.” *Geology* 29(11): 987–990.

Collobert, Ronan, and Jason Weston. 2008. “A Unified Architecture for Natural Language Processing: Deep Neural Networks with Multitask Learning.” In *Proceedings of the 25th International Conference on Machine Learning*, edited by William W. Cohen, Andrew McCallum, and Sam T. Roweis, 160–67. New York: Association for Computing Machinery. https://doi.org/10.1145/1390156.1390177.

Craig, Hugh. 2013. “Middle-distance reading with information-theory metrics.” September 24. (Talk abstract). https://www.gcdh.de/en/events/past-events/2013/. Accessed January 22, 2026.

Culler, Jonathan D. 2000. *Literary Theory: A Very Short Introduction*. Oxford: Oxford University Press.

———. 2010. “The Closeness of Close Reading” *ADE Bulletin* 149: 20–25. https://doi.org/10.1632/ade.149.20.

Da, Nan Z. 2019. “The Computational Case against Computational Literary Studies.” *Critical Inquiry* 45 (3): 601–39. https://doi.org/10.1086/702594.

De Man, Paul. 1986. “The Return to Philology.” In *The Resistance to Theory*, 21–26. Minneapolis: University of Minnesota Press.

De Melo Resende, Viviane. 2021. “Decolonizing critical discourse studies: for a Latin American perspective.” *Critical Discourse Studies* 18(1): 26–42. https://doi.org/10.1080/17405904.2018.1490654.

De Rougemont, Denis. 1939. *L’Amour et l’occident.* Paris: Plon.

Dennett, Daniel C. 1991. “Real Patterns.” *Journal of Philosophy* 88 (1): 27–51.

Denomy, Alexander J. 1965. *The Heresy of Courtly Love*. Gloucester: P. Smith.

Donaldson, Talbot E. 1970. “The Myth of Courtly Love.” In *Speaking of Chaucer*, edited by E. Talbot Donaldson, 154–63. London: Athlone.

Dronke, Peter. 1968. *Medieval Latin and the Rise of European Love-Lyric*. 2nd ed. 2 vols. Oxford: Clarendon Press.

Drucker, Johanna. 2017. “Why Distant Reading Isn’t.” *Publications of the Modern Language Association of America* 132 (3): 628–35. https://doi.org/10.1632/pmla.2017.132.3.628.

Eagleton, Terry. 2003. *After Theory*. New York: Basic Books.

Elden, Stuart. 2023. “Foucault as Translator of Binswanger and von Weizsäcker.” *Theory, Culture & Society* 40 (1–2): 91–116. https://doi.org/10.1177/0263276420950459.

Eliot, Thomas Stearns [1920] 1997. *The Sacred Wood: Essays on Poetry and Criticism*. London: Methuen. Reprint, London: Faber & Faber.

Empson, William. 1930. *Seven Types of Ambiguity*. London: Chatto & Windus.

English, James F., and Ted Underwood. 2016. "Shifting Scales: Between Literature and Social Science." *Modern Language Quarterly* 77(3): 277–295. https://doi.org/10.1215/00267929-3570612.

Eve, Martin Paul. 2019. *Close Reading with Computers: Textual Scholarship, Computational Formalism, and David Mitchell's Cloud Atlas*. Stanford, CA: Stanford University Press.

Felder, Ekkehard. 2006. *Semantische Kämpfe: Macht und Sprache in den Wissenschaften*. Berlin; New York: de Gruyter.

———. 2011. "Pragma-Semiotische Textarbeit Und Der Hermeneutische Nutzen von Korpusanalysen Für Die Linguistische Mediendiskursanalyse." In *Korpuspragmatik. Thematische Korpora Als Basis Diskurslinguistischer Analysen von Texten Und Gesprächen*, edited by Ekkehard Felder, Marcus Müller, and Friedemann Vogel, 115–74. https://doi.org/10.1515/9783110269574.115.

———. 2013. "Faktizitätsherstellung mittels handlungsleitender Konzepte und agonaler Zentren: der diskursive Wettkampf um Geltungsansprüche." In *Faktizitätsherstellung in Diskursen*, edited by Ekkehard Felder, 13–28. Berlin: de Gruyter.

Fernández-Vázquez, José-Santiago, and Ángel Sancho-Rodríguez. 2020. "Critical discourse analysis of climate change in IBEX 35 companies." *Technological Forecasting and Social Change* vol. 157 (C): 120063. https://doi.org/10.1016/j.techfore.2020.120063.

Firth, John Rupert. 1957. "A Synopsis of Linguistic Theory, 1930–1955." In *Studies in Linguistic Analysis*, 1–32. Special Volume of the Philological Society. Oxford: Basil Blackwell.

Fish, Stanley. 2011. "The Old Order Changeth." The New York Times (blog). December 26. https://archive.nytimes.com/opinionator.blogs.nytimes.com/2011/12/26/the-old-order-changeth/. Accessed January 22, 2026.

———. 2012a. "Mind Your P's and B's: The Digital Humanities and Interpretation." New York Times (blog). January 23. http://archive.nytimes.com/opinionator.blogs.nytimes.com/2012/01/23/mind-your-ps-and-bs-the-digital-humanities-and-interpretation/. Accessed January 22, 2026.

———. 2012b. "The Digital Humanities and the Transcending of Mortality." The New York Times (blog). January 9. https://archive.nytimes.com/opinionator.blogs.nytimes.com/2012/01/09/the-digital-humanities-and-the-transcending-of-mortality/. Accessed Jan. 21, 2026.

Foucault, Michel. 1969. *L'Archéologie du savoir.* Paris: Éditions Gallimard.

———. 1972. *The Archaeology of Knowledge*, translated by Alan M. Sheridan Smith. London: Tavistock.

———. 1975. *Surveiller et punir*. Paris: Éditions Gallimard.

———. 1976. *La Volonté de savoir.* Paris: Éditions Gallimard.

———. 1977. *Discipline and Punish*. New York: Pantheon Books.

———. 1980. "The Confession of the Flesh." In *Power/Knowledge Selected Interviews and Other Writings*, edited by Colin Gordon, 194-228. New York: Pantheon Books.

———. 1981. "The Order of Discourse." In *Untying the Text: A Post-Structuralist Reader,* edited by Robert Young. Boston, London, Henley: Routledge and Kegan Paul.

———. 1984a. *L'Histoire de la sexualité: l'usage des plaisirs.* Paris: Éditions Gallimard.

———. 1984b. *L'Histoire de la sexualité: le souci de soi.* Paris: Éditions Gallimard.

———. 1994. "Des supplices aux cellules." In *Dits et Ecrits 1970-1975*, edited by Daniel Defert, François Ewald, and Jacques Lagrange, volume II, 716–720. Paris: Éditions Gallimard.

Frappier, Jean. 1973. *Amour courtois et Table Ronde.* Geneva: Droz.

Frelick, Nancy. 2002. "Lacan, courtly love and Anamorphosis." In T*he Court Reconvenes: Courtly Literature across the Disciplines*, edited by Barbara K. Altmann and Carleton W. Carroll, 107–114. Vancouver: University of British Columbia Press.

Frow, John. 2008. "Thinking the Novel. Review of The Novel, edited by Franco Moretti." *New Left Review* 49: 137–145.

Gardt, Andreas. 2007. "Diskursanalyse. Aktueller theoretischer Ort und methodische Möglichkeiten." In *Diskurslinguistik nach Foucault. Theorien und Ge-genstände*, edited by Ingo H. Warnke, 23–48. Berlin, New York: Walter de Gruyter.

Gaunt, Simon. 1995. *Love and death in Medieval French and Occitan Courtly Literature: Martyrs to Love.* Oxford: Oxford University Press.

Gauthier, Laurent. 2025. "Digital Humanities and the Notion of Corpus in Ancient History." *Digital Studies/Le champ numérique* 15 (1): 1–28. https://doi.org/10.16995/dscn.18169.

Gavin, Michael. 2022. *Literary Mathematics: Quantitative Theory for Textual Studies*. Stanford, CA: Stanford University Press.

Girnth, Heiko. 2015. *Sprache Und Sprachverwendung in Der Politik: Eine Einführung in Die Linguistische Analyse Öffentlich-Politischer Kommunikation*. Berlin: De Gruyter.

Glissant, Édouard. [1990] 1997. *Poetics of Relation*, translated by Betsy Wing. Ann Arbor: University of Michigan Press.

Goldstone, Andrew. 2017. "The Doxa of Reading." *PMLA* 132(3): 636–642. https://doi.org/10.1632/pmla.2017.132.3.636.

Guilherme, Manuela. 2019. “The Critical Decolonial Quest for Intercultural Epistemologies and Discourses.” *Journal of Multicultural Discourses* 14 (1): 1–13. https://doi.org/10.1080/17447143.2019.1617294.

Günther, Fritz, Luca Rinaldi, and Marco Marelli. 2019. “Vector-Space Models of Semantic Representation from a Cognitive Perspective: A Discussion of Common Misconceptions.” *Perspectives on Psychological Science* 14 (6): 1006–33. https://doi.org/10.1177/1745691619861372.

Guo, Weijun, Xuedan Zheng, and Mingyue Jiang. 2026. “Investigating the OOV Problem and Its Impacts on Neural Program Repair.” In *Engineering of Complex Computer Systems*, edited by Yuan Zhou, Sin G. Teo, Xiaofei Xie, Zuohua Ding, and Yang Liu, 180–97. Lecture Notes in Computer Science 15746. Cham: Springer. https://doi.org/10.1007/978-3-032-00828-2_11.

Hahn, Thomas, Anna Siebach-Larsen, Susanna Fein, and Alan Lupack, eds. 2024. *The Middle English Texts Series*. Rochester, NY: University of Rochester, Rossell Hope Robbins Library. Accessed March 6, 2024. https://metseditions.org/.

Hall, Gary. 2013. “Toward a Postdigital Humanities: Cultural Analytics and the Computational Turn to Data-Driven Scholarship.” *American Literature* 85 (4): 781–809. https://doi.org/10.1215/00029831-2367337.

Harris, Zellig. 1951. *Methods in Structural Linguistics*. Chicago: University of Chicago Press.

———. 1952. “Discourse Analysis.” *Language* 28 (1): 1–30.

———. 1954. “Distributional Structure.” *WORD* 10 (2–3): 146–162. https://doi.org/10.1080/00437956.1954.11659520.

———. 1957. “Co-Occurrence and Transformation in Linguistic Structure.” *Language* 33 (3): 283–340.

———. 1959. “The Transformational Model of Language Structure.” *Anthropological Linguistics* 1 (1): 27–39.

———. 1982. “Discourse and Sublanguage.” In *Sublanguage: Studies of Language in Restricted Semantic Domains*, edited by Richard Kittredge and John Lehrberger, 231–36. New York; Berlin: De Gruyter. https://doi.org/10.1515/9783110844818-012.

Hayot, Eric. 2021. Humanist Reason. A History, An Argument, A Plan. New York: Columbia University Press.

Henriques, Julian, Wendy Hollway, Cathy Urwin, Couze Venn, and Valerie Walkerdine. 1984. *Changing the Subject: Psychology, Social Regulation, and Subjectivity*. New York: Methuen.

Hermanns, Fritz. 1995. “Sprachgeschichte als Mentalitätsgeschichte. Überlegungen zu Sinn und Form und Gegenstand historischer Semantik.” In *Sprachgeschichte des*

Neuhochdeutschen: Gegenstände, Methoden, Theorien, edited by Andreas Gardt, Klaus Mattheier, and Oskar Reichmann, 69–102. Tübingen: Niemeyer.

Hickey, Raymond. 1993. "Corpus Data Processing with LEXA." *ICAME Journal* 17: 73–95.

Hockey, Susan. 2004. "History of Humanities Computing." In *A Companion to Digital Humanities*, edited by Susan Schreibman, Ray G. Siemens and John Unsworth. Malden, MA: Blackwell.

Honneth, Axel. 2003. "Foucault und die Humanwissenschaften. Zwischenbilanz einer Rezeption." In *Zwischenbilanz einer Rezeption. Frankfurter Foucault-Konferenz 2001*, edited by Axel Honneth and Martin Saar. Frankfurt a. M.: Suhrkamp.

Hult, David F. 1996. "Gaston Paris and the Invention of Courtly Love." In *Medievalism and the Modernist Temper*, edited by Howard R. Bloch and Stephen G. Nichols, 192–224. Baltimore: Johns Hopkins University Press.

Jackson, Anne-Marie. 2015. "Kaupapa Māori Theory and Critical Discourse Analysis: Transformation and Social Change." *AlterNative: An International Journal of Indigenous Peoples* 11 (3): 256–68. https://doi.org/10.1177/117718011501100304.

Jacobs, Thomas, and Robin Tschötschel. 2019. "Topic Models Meet Discourse Analysis: A Quantitative Tool for a Qualitative Approach." *International Journal of Social Research Methodology* 22 (5): 469–85. https://doi.org/10.1080/13645579.2019.1576317.

Jaeger, Stephen C. 1985. *The Origins of Courtliness. Civilizing Trends and the Formation of Courtly Ideals*. Philadelphia: University of Pennsylvania Press.

———. 1999. *Ennobling Love: In Search of a Lost Sensibility*. Philadelphia: University of Pennsylvania Press.

Jenkins, Tiffany. 2013. "Don't count on big data for answers." The Scotsman (Newspaper). https://www.scotsman.com/news/opinion/columnists/tiffany-jenkins-dont-count-on-big-data-for-answers-1590688. Accessed January 22, 2026.

Jockers, Matthew Lee. 2013. *Macroanalysis: Digital Methods and Literary History*. Champaign: University of Illinois Press.

———. 2014. Text Analysis with R for Students of Literature. New York: Springer.

Johnstone, Barbara. 2005. "Discourse Analysis and Narrative." In *The Handbook of Discourse Analysis*, edited by Deborah Schiffrin, Deborah Tannen, and Heidi E. Hamilton, 635–649. https://doi.org/10.1002/9780470753460.ch33.

Jolliffe, Ian T. 2002. *Principal Component Analysis*. 2nd ed. New York: Springer.

Joshi, Priya. 2002. "Quantitative Method, Literary History." *Book History* 5 (1): 263–74. https://doi.org/10.1353/bh.2002.0011.

Jung, Matthias. 2001. "Diskurshistorische Analyse: Eine Linguistische Perspektive." In *Handbuch Sozialwissenschaftliche Diskursanalyse*, edited by Reiner Keller, Andreas

Hirseland, Werner Schneider, and Willy Viehöver, 29–51. Wiesbaden: VS Verlag für Sozialwissenschaften.

Kay, Sarah. 2000. "Courts, Clerks, and Courtly Love." In *The Cambridge Companion to Medieval Romance*, edited by Roberta L. Krueger, 81–96. Cambridge University Press. https://doi.org/10.1017/CCOL0521553423.

Kean, Patricia Margaret. 1972. *Chaucer and the Making of English Poetry. Love Vision and Debate*, vol. 1. London: Routledge; Kegan Paul.

Keller, Reiner, Andreas Hirseland, Werner Schneider, and Willy Viehöver, eds. 2006. *Handbuch Sozialwissenschaftliche Diskursanalyse*: *Theorien und Methoden*. Vol. 1. 2nd ed. Opladen: Leske + Budrich.

Kennedy, Graeme. 1998. *An Introduction to Corpus Linguistics*. London: Longman.

Ker, William P. [1897] 1957. *Epic and Romance: Essays on Medieval Literature*. Repr. New York: Dover Publications.

Kerchner, Brigitte, and Silke Schneider. 2006. "'Endlich Ordnung in der Werkzeugkiste.' Zum Potenzial der Foucaultschen Diskursanalyse für die Politikwissenschaft." In *Foucault: Diskursanalyse der Politik. Eine Einführung*, edited by Brigitte Kerchner and Silke Schneider, 9-30. Wiesbaden: Verlag für Sozialwissenschaften.

Kirschenbaum, Matthew G. 2010. "What Is Digital Humanities and What's It Doing in English Departments?" ADE Bulletin 150: 55–61. https://doi.org/10.1632/ade.150.55.

Kitchin, Rob. 2014. "Big Data, New Epistemologies and Paradigm Shifts." *Big Data and Society* 1 (1): 1–12. https://doi.org/10.1177/2053951714528481.

Kitchin, Rob, and Tracey P. Lauriault. 2015. "Small Data in the Era of Big Data." *GeoJournal* 80 (4): 463–475. https://doi.org/10.1007/s10708-014-9601-7.

Kögler, Hans-Herbert. 2004. *Michel Foucault*. 2nd Ed. Stuttgart; Weimar: Metzler Verlag.

Köhler, Erich. 1976. *Sociologia della fin'amor. Saggi trobadotici*, edited by Mario Mancini. Padova: Liviana Editrice.

Koselleck, Reinhart. 1972. "Einleitung." In *Geschichtliche Grundbegriffe: Historisches Lexikon zur politisch-sozialen Sprache in Deutschland*, edited by Otto Brunner, Werner Conze, and Reinhart Koselleck, Band 1, xiii–xxvii. Stuttgart: Klett-Cotta.

Kristeva, Julia. 1986. "Word, Dialogue, and Novel." In *The Kristeva Reader*, edited by Toril Moi, 34–61. New York: Columbia University Press.

Kuhn, Sherman M. 1982. "On the Making of the Middle English Dictionary." *Dictionaries* 4: 14–41.

Lacan, Jacques. [1959–60] 1986. *Le séminaire, Livre VII: L'éthique de la psychanalyse.* Paris: Seuil.

LaCapra, Dominick. 1985. *History and Criticism*. Ithaca, NY: Cornell University Press.

Lazar, Michelle. 2005. *Feminist Critical Discourse Analysis: Gender, Power and Ideology*. New York: Palgrave Macmillan.

Leavis, Frank Raymond. 1948. *The Great Tradition: George Eliot, Henry James, Joseph Conrad*. London: Chatto & Windus.

Lefebvre, Henri. 1991. *The Production of Space*, edited by Donald Nicholson-Smith. Cambridge; Oxford: Blackwell.

Lepojärvi, Jason. 2022. "Brilliance and Blindspots: New Light on C. S. Lewis's *The Four Loves*." *The Heythrop Journal* 63 (6): 1109–23. https://doi.org/10.1111/heyj.14146

Lewis, Clive S. [1936] 2013. *The Allegory of Love: A Study in Medieval Tradition*. Cambridge: Cambridge University Press.

———. 1940. *The Problem of Pain*. London: Geoffrey Bles.

———. 1942. *The Screwtape Letters*. London: Geoffrey Bles.

———. 1950–56. *The Chronicles of Narnia*. 7 vols. London: Geoffrey Bles and the Bodley Head.

———. 1952. *Mere Christianity*. London: Geoffrey Bles.

———. 1960. *The Four Loves*. London: Geoffrey Bles.

Lewis, Robert E. 2002a. "Some Observations on the Use of Manuscripts, Dates, and Preferred Editions in the Middle English Dictionary." In *And Gladly Wolde He Lerne and Gladly Teche: Essays on Medieval English Presented to Professor Matsuji Tajima on His Sixtieth Birthday*, edited by Yoko Iyeiri and Margaret Connolly, 169–79. Tokyo: Kaibunsha.

———. 2002b. "The Middle English Dictionary at 71." *Dictionaries* 23: 76–94. https://doi.org/10.1353/dic.2002.0012.

Lewis, Robert E., and Mary Jane Williams, eds. 2007. *Middle English Dictionary: Plan and Bibliography*. 2nd ed. Ann Arbor: University of Michigan Press.

Liberman, Mark. 2012. "The 'Dance of the P's and B's': Truth or Noise?" *Language Log* (blog). January 26, 2012. languagelog.ldc.upenn.edu/nll/?p=3730. Accessed January 22, 2026.

Lindquist, Hans and Magnus Levin. 2018. *Corpus Linguistics and the Description of English*. 2nd Ed. Edinburgh: Edinburgh University Press.

Maingueneau, Dominique. 2010. "Literature and Discourse Analysis." *Acta Linguistica Hafniensia* 42 (1): 147–58. https://doi.org/10.1080/03740463.2010.482325.

Malory, Thomas. 2017. *Le Morte Darthur: The Original Text Edited from the Winchester Manuscript and Caxton's Morte Darthur*. Edited by P. J. C. Field. Cambridge: D. S. Brewer.

Mann, Jill. 2013. "Falling in Love in the Middle Ages." In *Traditions and Innovations in the Study of Medieval English Literature: The Influence of Derek Brewer*, edited by Charlotte Brewer and Barry Windeatt, 88–110. Cambridge: Cambridge University Press.

Manovich, Lev. 2020. *Cultural Analytics*. Cambridge, MA: The MIT Press.

Marche, Stephen. 2012. "Literature Is Not Data: Against Digital Humanities." Los Angeles Review of Books (blog). October. https://lareviewofbooks.org/article/literature-is-not-data-against-digital-humanities/. Accessed January 22, 2026.

Markus, Manfred. 1997. "Normalization of Middle English Prose: Possibilities and Limits." In *Corpus-Based Studies in English: Papers from the Seventeenth International Conference on English Language Research on Computerized Corpora*, edited by Magnus Ljung, 211–26. Amsterdam: Rodopi.

Matthews, William. 1965. "Inherited Impediments in Medieval Literary History." In *Medieval Secular Literature*, edited by William Mathews. Berkeley: University of California Press.

Mattison, J. R., and Hannah K. Ryley. 2025. "Big Data and Manuscript Studies: An Introduction with Some Reflections on Methods." *Digital Philology: A Journal of Medieval Cultures* 14 (1): 1–14. https://doi.org/10.1353/dph.2025.a959161.

McCarthy, Conor, ed. 2022. *Love, Sex and Marriage in the Middle Ages: A Sourcebook*. Second edition. London; New York: Routledge.

McCarthy, Willard. 2005. *Humanities Computing.* London: Palgrave Macmillan.

———. 2012. "A Telescope for the Mind?" In *Debates in the Digital Humanities*, edited by Matthew K. Gold. Minneapolis: University of Minnesota Press. https://doi.org/10.5749/minnesota/9780816677948.003.0013.

McEnery, Tony, and Vaclav Brezina. 2022. *Fundamental Principles of Corpus Linguistics*. Cambridge: Cambridge University Press.

McEnery, Tony, and Andrew Wilson. 2001. *Corpus Linguistics: An Introduction*. 2nd ed. Edinburgh: Edinburgh University Press.

McGann, Jerome. 2001. *Radiant Textuality: Literature after the World Wide Web*. New York; London: Palgrave Macmillan.

McGillivray, Murray, ed. 1997. *Geoffrey Chaucer's Book of the Duchess: A Hypertext Edition*. CD-ROM. Calgary: University of Calgary Press.

McInnes, Leland, John Healy, and James Melville. 2018. "UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction." arXiv. https://arxiv.org/abs/1802.03426.

Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. “Efficient Estimation of Word Representations in Vector Space.” arXiv. https://arxiv.org/abs/1301.3781.

Mikolov, Tomas, Wen-tau Yih, and Geoffrey Zweig. 2013b. “Linguistic Regularities in Continuous Space Word Representations.” In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 746–51. Atlanta: Association for Computational Linguistics. https://aclanthology.org/N13-1090/.

Mills, Sara. 2004. *Discourse*. London; New York: Routledge.

Mnih, Andriy, and Geoffrey E. Hinton. 2009. “A Scalable Hierarchical Distributed Language Model.” In *Advances in Neural Information Processing Systems 21*, edited by Daphne Koller, Dale Schuurmans, Yoshua Bengio, and Léon Bottou, 1081–88. La Jolla, CA: Neural Information Processing Systems.

Moller, Herbert. 1959. “The Social Causation of the Courtly Love Complex.” *Comparative Studies in Society and History* 1: 137–63.

Monson, Don A. 2022. Eros and Noesis: A Cognitive Approach to the Courtly Love Literature of Medieval France. Leiden and Boston: Brill.

Moore, John C. 1979. “‘Courtly Love’: A Problem of Terminology.” *Journal of the History of Ideas* 40 (4): 621–632.

Moretti, Franco. 1983. *Signs Taken for Wonders: Essays in the Sociology of Literary Forms*. London: Verso.

———. 2000. ‘The Slaughterhouse of Literature’. *Modern Language Quarterly* 61: 207–228.

———. 2005. *Graphs, Maps, Trees: Abstract Models for a Literary History*. London: Verso.

———. 2013. *Distant Reading.* London: Verso.

Mott, Lewis F. 1896. *The System of Courtly Love*. Boston: Ginn and Co. Publishers; Athenaeum Press.

Mueller, Martin. 2014. “Shakespeare His Contemporaries: Collaborative Curation and Exploration of Early Modern Drama in a Digital Environment.” *Digital Humanities Quarterly* 8 (3). http://digitalhumanities.org:8081/dhq/vol/8/3/000183/000183.html. Accessed September 29, 2024.

Mulhern, Francis. [1979] 1981. *The Moment of “Scrutiny.”* London: Verso.

Nagel, Ernest. 1961. *The Structure of Science: Problems in the Logic of Scientific Explanation*. New York: Harcourt, Brace & World.

Neilson, William Allan. 1899. *The Origins and Sources of the Court of Love*. Boston: Ginn and Co. Publishers.

Newman, Francis Xavier, ed. 1968. *The Meaning of Courtly Love*. New York: State University of New York Press.

Niehr, Thomas. 2014. *Einführung in die linguistische Diskursanalyse*. Darmstadt: Wissenschaftliche Buchgesellschaft.

North, Joseph. 2013. "What's 'New Critical' about 'Close Reading'? I. A. Richards and His New Critical Reception." *New Literary History* 44: 141–157. https://doi.org/10.1353/nlh.2013.0002.

O'Donnell, Daniel P. 2018. "The Humanities and New Technologies: Exploring Digital Tools for Innovation and Development." Lagos Summer School in Digital Humanities October 2018 (LSSDH 2018), Lagos, Nigeria. https://doi.org/10.5281/zenodo.1443256.

O'Donoghue, Bernard. 2006. "The Reality of Courtly Love." In *Writings on Love in the English Middle Ages. Studies in Arthurian and Courtly Cultures*, edited by Helen Cooney, 7–24. New York: Palgrave Macmillan. https://doi.org/10.1057/9781403983534_2.

Oprea, Delia Andreea. 2019. "Texts, Language, and Communication in Social Media: A Critical Discourse Analysis." *International Multidisciplinary Scientific Conference on the Dialogue between Sciences & Arts, Religion & Education* 3(3): 315–320.

Paris, Gaston. 1881. "Etudes sur les romans de la table ronde. I." *Romania* 10 (40): 465–496.

———. 1883. "Études sur les Romans de la Table Ronde. II." *Romania* 12 (48): 459–534.

Parry, John J., ed. 1941. *The Art of Courtly Love.* New York: Columbia University Press.

Paxson, James J., and Cynthia Gravlee, eds. 1998. *Desiring Discourse: The Literature of Love, Ovid through Chaucer*. Selinsgrove: Susquehanna University Press.

Perrine, Laurence. 1956. *Sound and Sense: An Introduction to Poetry*. New York: Harcourt, Brace and Company.

Piotrowski, Michael. 2012. *Natural Language Processing for Historical Texts*. San Rafael, CA: Morgan & Claypool.

Piper, Andrew. 2015. "Novel Devotions: Conversional Reading, Computational Modeling, and the Modern Novel." *New Literary History* 46(1): 63–98. https://doi.org/10.1353/nlh.2015.0008.

———. 2016. "There Will Be Numbers." *Journal of Cultural Analytics* 1 (1). https://doi.org/10.22148/16.006.

———. 2018. *Enumerations*: Data and Literary Study. Chicago: University of Chicago Press.

———. 2020. *Can We Be Wrong? The Problem of Textual Evidence in a Time of Data*. Cambridge: Cambridge University Press.

Potter, Rosanne G. 1988. "Literary Criticism and Literary Computing." *Computers in the Humanities* 22 (2).

Primorac, Antonija, Rosario Arias, Roxana Patras, et al. 2023. “Distant Reading Two Decades On: Reflections on the Digital Turn in the Study of Literature.” *Digital Studies/Le Champ Numérique* 13 (1). https://doi.org/10.16995/dscn.8855.

Putrama, I Made, and Péter Martinek. 2024. “Heterogeneous Data Integration: Challenges and Opportunities.” *Data in Brief* 56: 110853. https://doi.org/10.1016/j.dib.2024.110853.

Radway, Janice. 1984. Reading the Romance: Women, Patriarchy, and Popular Literature. Chapel Hill: University of North Carolina Press.

Ramsay, Stephen. 2003. “Special Section: Reconceiving Text Analysis: Toward an Algorithmic Criticism.” *Literary and Linguistic Computing* 18 (2): 167–74. https://doi.org/10.1093/llc/18.2.167.

———. 2007. “Algorithmic Criticism.” In *A Companion to Digital Literary Studies*, edited by Ray G. Siemens and Susan Schreibman. Malden, MA: Blackwell.

———. 2011. *Reading Machines: Toward an Algorithmic Criticism*. Topics in the Digital Humanities. Urbana: University of Illinois Press.

———. 2014. “The Hermeneutics of Screwing Around; or What You Do with a Million Books.” In *Pastplay: Teaching and Learning History with Technology: Teaching and Learning History with Technology*, edited by Kevin Kee, 111–20. University of Michigan Press.

Ransom, John Crowe. 1937. “Criticism, Inc.” *Virginia Quarterly Review* 13: 586–603.

Reddy, William M. 2012. *The Making of Romantic Love*. Chicago; London: The University of Chicago Press.

Reichardt, Rolf. 1985. “Einleitung.” In *Handbuch politisch-sozialer Grundbegriffe in Frankreich 1680–1820*, edited by Rolf Reichardt and Eberhard Schmitt, Heft 1/2, 39–148. Munich: Oldenbourg.

Reimers, Nils, and Iryna Gurevych. 2019. “Sentence-BERT: Sentence Embeddings Using Siamese BERT-Networks.” In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the Ninth International Joint Conference on Natural Language Processing*, 3982–92. Hong Kong: Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1410.

Richards, I. A. [1924] 2017. *Principles of Literary Criticism*. London: Routledge & Kegan Paul. Reprint, New York: Routledge.

———. [1929] 1978. *Practical Criticism: A Study of Literary Judgment*. New York: Harcourt, Brace and Company. Reprint, London: Kegan Paul, Trench, Trubner & Co.

Robertson, Durant W. 1968. “The Concept of Courtly Love as an Impediment to the Understanding of Medieval Texts.” In *The Meaning of Courtly Love*, edited by Francis Xavier Newman. New York: State University of New York Press.

Robinson, F. N., ed. 1957. *The Works of Geoffrey Chaucer*. 2nd ed. Boston: Houghton Mifflin.

Rockwell, Geoffrey. 2003. “What Is Text Analysis, Really?” *Literary and Linguistic Computing* 18 (2): 209–219.

Rockwell, Geoffrey, and Stéfan Sinclair. 2016. *Hermeneutica: Computer-Assisted Interpretation in the Humanities*. Cambridge, MA: MIT Press. https://doi.org/10.7551/mitpress/9522.001.0001.

Roderick, Ian. 2018. “Multimodal Critical Discourse Analysis as Ethical Praxis.” *Critical Discourse Studies*, 15 (2): 154–168. https://doi.org/10.1080/17405904.2017.1418401.

Rosenwein, Barbara H. 2002. “Worrying about Emotions in History.” *The American Historical Review* 107 (3): 821–45. https://doi.org/10.1086/532498.

———. 2006. Emotional Communities in the Early Middle Ages. Ithaca, NY: Cornell University Press.

Rothberg, Michael. 2010. “Quantifying Culture? A Response to Eric Slauter.” *American Literary History* 22(2): 341–346.

Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. 1986. “Learning Representations by Back-Propagating Errors.” *Nature* 323: 533–36. https://doi.org/10.1038/323533a0.

Ruoff, Michel. 2007. *Foucault Lexikon. Entwicklung – Kernbegriffe – Zusammenhänge*. Paderborn: Fink.

Russell, Jeffrey B. 1965. “Courtly Love as Religious Dissent.” *The Catholic Historical Review* 51 (1): 31–44.

Sahlgren, Magnus. 2024. “Distributional Legacy: The Unreasonable Effectiveness of Harris’s Distributional Program.” *WORD* 70 (4): 246–57. https://doi.org/10.1080/00437956.2024.2414515.

Saint-Amour, Paul. 2019. “The Medial Humanities: Toward a Manifesto for Meso-Analysis.” *Modernism/Modernity* 3(4). https://modernismmodernity.org/forums/posts/medial humanities-toward-manifesto-meso-analysis/. Accessed January 22, 2026.

Sample, Mark. 2017. “A Parallax Reading of Roethke’s ‘My Papa’s Waltz.’” @samplereality (blog). Accessed January 22, 2026. https://samplereality.com/2017/05/31/a-parallax-reading-ofroethkes-my-papas-waltz/.

Sarasin, Philipp. 2005. *Michel Foucault zur Einführung*. Hamburg: Junius.

Sawyer, R. Keith. 2002. “A Discourse on Discourse: An Archeological History of an Intellectual Concept.” *Cultural Studies* 16 (3): 433–56. https://doi.org/10.1080/09502380210128324.

Scharloth, Joachim, David Eugster, and Noah Bubenhofer. 2013. “Das Wuchern der Rhizome. Linguistische Diskursanalyse und Data-driven Turn.” In *Linguistische Diskursanalyse: Neue Perspektiven*, edited by Dietrich Busse and Wolfgang Teubert, 345–380. Wiesbaden: Springer VS.

Scherner, Maximilian. 1994. “Textverstehen Als ‘Spurenlesen.’ Zur Texttheoretischen Tragweite Dieser Metapher.” In *Text Und Grammatik: Festschrift Für Roland Harweg Zum 60. Geburtstag*, edited by Roland Harweg, Peter Canisius, Clemens-Peter Herbermann, and Gerhard Tschauder, 317–340. Bochum: Universitätsverlag N. Brockmeyer.

Schöch, Christof. 2017. “Aufbau von Datensammlungen.” In *Digital Humanities: Eine Einführung*, edited by Fotis Jannidis, Hubertus Kohle, and Malte Rehbein, 223–33. Stuttgart: J. B. Metzler.

Schnell, Rüdiger. 1989. “L’amour courtois en tant que discours courtois sur l’amour (I).” *Romania* 110 (437): 72-126.

Shannon, Claude E., and Warren Weaver. 1949. *The Mathematical Theory of Communication*. Urbana: University of Illinois Press.

Shawver, Gary, ed. 1996. *Patience*. In *Using TACT with Electronic Texts: A Guide to Text-Analysis Computing Tools, Version 2.1 for MS-DOS and PC DOS*, by Ian Lancashire, in collaboration with John Bradley, Willard McCarthy, Michael Stairs, and T. R. Wooldridge. CD-ROM. New York: Modern Language Association of America.

Sinclair, John McH. 1991. *Corpus, Concordance, Collocation*. Oxford: Oxford University Press.

Sinclair, Stefan. 2003. “Computer-Assisted Reading: Reconceiving Text Analysis.” *Literary and Linguistic Computing* 18 (2): 175–84. https://doi.org/10.1093/llc/18.2.175.

Smith, Barbara Herrnstein. 2016. “What Was ‘Close Reading’?” *The Minnesota Review* (87): 57–75. https://doi.org/10.1215/00265667-3630844.

Snow, Charles Percy. [1959] 2001. *The Two Cultures.* Cambridge: Cambridge University Press.

Spieß, Constanze. 2011. *Diskurshandlungen. Theorie und Methode linguistischer Diskursanalyse am Beispiel der Bioethikdebatte.* Berlin, Boston: De Gruyter.

Spitzmüller, Jürgen, and Ingo H. Warnke. 2011. *Diskurslinguistik: Eine Einführung in Theorien und Methoden der transtextuellen Sprachanalyse*. Berlin and Boston: De Gruyter.

Spivak, Gayatri. 2003. *Death of a Discipline*. New York: Columbia University Press.

Stearns, Peter N., and Carol Z. Stearns. 1985. “Emotionology: Clarifying the History of Emotions and Emotional Standards,” *The American Historical Review* 90 (4): 813–36.

Stefanowitsch, Anatol. 2020. *Corpus Linguistics: A Guide to the Methodology*. Berlin: Language Science Press.

Stendhal. 1822. *De l’Amour*. Paris: Mongie.

Stevens, Hallam. 2013. *Life Out of Sequence: A Data-Driven History of Bioinformatics.* Chicago: University of Chicago Press.

Stibbe, Arran. 2014. “An Ecolinguistic Approach to Critical Discourse Studies.” *Critical Discourse Studies* 11 (1): 117–128. https://doi.org/10.1080/17405904.2013.845789.

Stonebraker, Michael, and Uğur Çetintemel. 2005. "'One Size Fits All': An Idea Whose Time Has Come and Gone." In *Proceedings of the 21st International Conference on Data Engineering (ICDE '05)*, 2–11. Washington, DC: IEEE Computer Society.

Strang, Barbara M. H. 1970. *A History of English*. London: Methuen.

Swales, John. 1990. *Genre Analysis: English in Academic and Research Settings*. Cambridge: Cambridge University Press.

Tatlock, John S. P. 1941. "The People in Chaucer's Troilus." *PMLA* 56: 85–104.

Thompson, Nancy S. 1996. *Chaucer, Boccaccio, and the Debate of Love: A Comparative Study of the Decameron and the Canterbury Tales*. Oxford: Clarendon Press.

Tissari, Heli. 2002. "An Experiment with Historical Corpora: A Cognitive Approach to the Study of the English Word Love." In *Variation Past and Present: VARIENG Studies on English for Terttu Nevalainen*, edited by Helena Raumolin-Brunberg, Minna Nevala, Arja Nurmi, and Matti Rissanen, 303–19. Helsinki: Société Néophilologique.

———. 2003. *LOVEscapes: Changes in Prototypical Senses and Cognitive Metaphors since 1500*. Mémoires de la Société Néophilologique de Helsinki 62. Helsinki: Société Néophilologique.

Tolkien, J. R. R. 1936. "Beowulf: The Monsters and the Critics." *Proceedings of the British Academy* 22: 245–295.

Trumpener, Katie. 2009. "Critical response I. Paratext and genre system: a response to Franco Moretti." *Critical Inquiry* 36(1): 159–71.

Turian, Joseph, Lev-Arie Ratinov, and Yoshua Bengio. 2010. "Word Representations: A Simple and General Method for Semi-Supervised Learning." In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, 384–94. Uppsala: Association for Computational Linguistics. https://aclanthology.org/P10-1040/. Accessed July 10, 2026.

Unamuno, Enrique Santos. 2017. "GIS and Telescopic Reading: Between Spatial and Digital Humanities." *Neohelicon* 44(1): 65–81. https://doi.org/10.1007/s11059-017-0381-1.

Underwood, Ted. 2016. "The Real Problem with Distant Reading." The Stone and the Shell (blog). May 29. https://tedunderwood.com/2016/05/29/the-real-problem-with-distant-reading/. Accessed January 21, 2026.

———. 2017. "A Genealogy of Distant Reading." *Digital Humanities Quarterly* 11(2). https://dhq-static.digitalhumanities.org/pdf/000317.pdf. Accessed January 22, 2026.

———. 2019. *Distant Horizons: Digital Evidence and Literary Change*. Chicago: The University of Chicago Press.

Utley, Francis L. 1972. “Must We Abandon the Concept of Courtly Love?” *Medievalia et humanistica* 3: 299–324.

Van Dijk, Teun A. 2008. *Discourse and Power*. London; New York: Palgrave Macmillan.

Van Zundert, Joris J. 2015. “Screwmeneutics and Hermenumericals: The Computationality of Hermeneutics.” In A New Companion to Digital Humanities, edited by Susan Schreibman, Ray Siemens, and John Unsworth, 1st ed., 331–47. Hoboken, NJ: Wiley.

Vandevelde, Pol. 2001. “Translator’s Introduction.” In *The Problem of Love in the Middle Ages: A Historical Contribution*, 11–75. Milwaukee: Marquette University Press.

Viehhauser, Gabriel. 2017. “Digitale Gattungsgeschichten. Minnesang Zwischen Generischer Konstanz Und Wende.” *Zeitschrift für digitale Geisteswissenschaften* 2. https://doi.org/10.17175/2017_003.

Vogelsang, Kai. 2012. “Conceptual History: A Short Introduction.” *Oriens Extremus* 51: 9–24. http://www.jstor.org/stable/24047785.

Wagner, Ina, and Ruth Wodak. 2006. “Performing Success: Identifying Strategies of Self-Presentation in Women’s Biographical Narratives.” *Discourse & Society* 17(3): 385–411. https://doi.org/10.1177/0957926506060251.

Wang, Liang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024. “Improving Text Embeddings with Large Language Models.” In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 11897–11916. Bangkok: Association for Computational Linguistics. https://doi.org/10.18653/v1/2024.acl-long.642.

Warnke, Ingo H. 2007. *Diskurslinguistik Nach Foucault: Theorie und Gegenstände*. Berlin; New York: De Gruyter.

———. 2009. “Die sprachliche Konstituierung von geteiltem Wissen in Diskursen.” In *Wissen durch Sprache. Theorie, Praxis und Erkenntnisinteresse des Forschungsnetzwerkes ‘Sprache und Wissen’*, edited by Ekkehard Felder and Marcus Müller, 113–40. Berlin; New York: de Gruyter.

Warren, Michelle R., and Neil Weijer. 2021. “Re-Imagining Digital Things: Sustainable Data in Medieval Manuscript Studies.” *Digital Philology: A Journal of Medieval Cultures* 10 (1): 111–134. https://doi.org/10.1353/dph.2021.0005.

Watt, Ian P. [1957] 1987. *The Rise of the Novel: Studies in Defoe, Richardson, and Fielding*. London: Hogarth.

Wellek, René, and Austin Warren. 1949. Theory of Literature. New York: Harcourt, Brace.

Wengeler, Martin. 2003. *Topos und Diskurs: Begründung Einer Argumentationsanalytischen Methode und Ihre Anwendung Auf den Migrationsdiskurs. 1960-1985*. Berlin: de Gruyter.

White, Curtis. 2013. The Science Delusion: Asking the Big Questions in a Culture of Easy Answers. New York: Melville House.

Widdowson, Henry G. 1995. “Discourse analysis: a critical view.” *Language and Literature* 4(3): 157–172.

Wilcox, John C. 1930. “Defining courtly love.” *Michigan Academy of Science, Arts and Letters* xii: 313–25.

Wilhelm, James J. 1965. *The Cruelest Month: Spring, Nature and Love in Classical and Medieval Lyrics.* New Haven: Yale University Press.

Wilson, Andrew Norman. 1990. *C. S. Lewis: A Biography*. New York: W. W. Norton.

Wimsatt, W. K. 1954. *The Verbal Icon: Studies in the Meaning of Poetry*. Lexington: University of Kentucky Press.

Wimsatt, William C. 2007. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, MA: Harvard University Press.

Wittgenstein, Ludwig. [1953] 1994. *Philosophical Investigations*. Translated by G. E. M. Anscombe. 3rd ed. Oxford: Blackwell.

Wodak, Ruth. 2009. *The Discourse of Politics in Action: Politics as Usual*. London; New York: Palgrave Macmillan.

Wodak, Ruth, and Michael Meyer. 2015. “Critical Discourse Studies: History, Agenda, Theory and Methodology.” In *Methods of Critical Discourse Studies*, edited by Ruth Wodak and Michael Meyer, 1–22. Thousand Oaks, CA: Sage Publications.

Yin, Zi, and Yuanyuan Shen. 2018. “On the Dimensionality of Word Embedding.” In *Advances in Neural Information Processing Systems 31*, 895–906. Red Hook, NY: Curran Associates.

Zipf, George Kingsley. 1935. *The Psycho-Biology of Language: An Introduction to Dynamic Philology*. Boston: Houghton Mifflin.

Žižek, Slavoj. 2005. *The Metastases of Enjoyment: Six Essays on Women and Causality*. London; New York: Verso.

APPENDIX A: DATASET AND METADATA

#	title	author	corpus	traditional composition date	date of source	date of edition	periodisation	genre	love	Bibliographic reference	Notes
1	Boece	Geoffrey Chaucer	MED	c.1380	15th century manuscripts	1868	Ricardian	philosophical translation	Miscellaneous	Morris, Richard, ed. 1868. Chaucer's Translation of Boethius's "De Consolatione Philosophiæ." Early English Text Society, Extra Series 5. London: N. Trübner.	Morris's text is edited from British Museum / British Library Additional MS 10,340, collated with Cambridge University Library MS Ii.3.21.
2	Everyman	Anonymous	MED	c.1495	early 16th century	1961	Post-Ricardian	morality play	Religious	Cawley, A. C., ed. 1961. Everyman and Medieval Miracle Plays. Manchester: Manchester University Press.	The text refers to the Middle English morality play tradition, adapted from Dutch; the text refers to the broader Elckerlijc/Everyman tradition, which survives in early sixteenth-century printed witnesses.
3	Generydes	Anonymous	MED	c.1375–1400	c.1400	1873; 1878	Ricardian	romance	Secular	Wright, W. Aldis, ed. [1873] 1878. Generydes: A Romance in Seven-Line Stanzas. Early English Text Society, Original Series 55, Rev.	Wright's edition is from the unique paper manuscript in Trinity College, Cambridge, about 1440.

										Ed.. London: Trübner. Generydes: A Romance in Seven-Line Stanzas. Early English Text Society, Original Series 70. London: Trübner.	
4	Havelok the Dane	Anonymous	MED	c.1290	late 13th / early 14th century	1868	Pre-Ricardian	romance	Secular	Skeat, Walter W., ed. 1868. The Lay of Havelok the Dane. London: Trübner; Early English Text Society.	Skeat's edition is based on the unique Bodleian Library MS Laud Misc. 108.
5	Henryson's Minor Poems	Robert Henryson	METS	c.1480	15th century	2010	Post-Ricardian	lyric and moral poetry	Miscellaneous	Henryson, Robert. 2010. The Complete Works. Edited by David J. Parkinson. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Because of the dispute between stronger and weaker attribution of those works, I chose only the stronger attributed texts.
6	Letter of Cupid	Thomas Hoccleve	MED	1402	mid-15th c.	1897	Post-Ricardian	love complaint / anti-misogynist debate poem / Chaucerian love poetry	Secular	Skeat, Walter W., ed. 1897. The Complete Works of Geoffrey Chaucer. Vol. 7: Chaucerian and Other Pieces.	It is part of Skeat's mixed Chaucerian miscellany

										Oxford: Clarendon Press.	
7	Regiment of Princes	Thomas Hoccleve	MED	15th c.	1400–1426	1897	Post-Ricardian	advice for princes / didactic political poetry / mirror for princes	Miscellaneous	Furnivall, Frederick J., and I. Gollancz, eds. 1892–1925. Hoccleve's Works. Early English Text Society, Extra Series 61, 72, 73. London: Kegan Paul, Trench, Trübner.	Furnivall's text as The Regement of Princes A.D. 1411–12, from the Harleian MS. 4866; Medieval Scribes dates Harley 4866 to 1411–1412.
8	Iacob and Iosep	Anonymous	MED	mid-13th c.	soon after mid-13th c.	1916	Post-Ricardian	biblical narrative poem	Religious	Napier, Arthur Sampson, ed. 1916. Jacob and Josep: A Middle English Poem of the Thirteenth Century. Oxford: Clarendon Press.	Text based on Oxford, Bodleian Library MS Bodley 652, ff. 1r–10v
9	Joseph of Arimathie	Anonymous	MED	c.1350	late 14th c.	1871	Pre-Ricardian	Arthurian religious romance / holy grail	Religious	Skeat, Walter W., ed. 1871. Joseph of Arimathie, Otherwise Called the Romance of the Seint Graal. Early English Text Society, Original Series 44. London: N. Trübner.	Text based on Vernon MS, Oxford, Bodleian Library MS Eng. poet. a.1

10	King Ponthus and Fair Sidone	Anonymous	MED	c. 1450–1500	16th century.	1897	Post-Ricardian	romance	Secular	Mather, Frank Jewett, ed. 1897. King Ponthus and the Fair Sidone. Publications of the Modern Language Association 12. Baltimore: MLA.	The first date relates to the French original and the second is the Middle English translation. The witness is preserved in a 16th-century print edition.
11	Le Morte Darthur	Thomas Malory	Other Edition	1469–1470	c.1470-1438; 1485.	2017	Post-Ricardian	Arthurian prose romance	Secular	Malory, Thomas. 2017. Le Morte Darthur: The Original Text Edited from the Winchester Manuscript and Caxton's Morte Darthur. Edited by Peter J. C. Field. Cambridge: D. S. Brewer.	Field's edition based on the Winchester manuscript and the Caxton's edition, whose dates are respectively indicated. (The text was not taken from the MED not because it was not part of the corpus, but because of a better version available.)
12	Lovelich, History of the Holy Grail	Henry Lovelich	MED	1370; 1180-1200	c.1450	1905	Post-Ricardian	Arthurian religious romance	Religious	Furnivall, Frederick J., ed. 1905. The History of the Holy Grail. Early English Text Society, Extra Series 95. London: N. Trübner.	From Cambridge, Corpus Christi College MS 80.
13	Melusine	Jean d'Arras; anonymous	MED	1382–1394	c.1500	1895	Post-Ricardian	romance	Secular	Donald, A. K., ed. 1895. Melusine. Early English Text Society, Extra Series 68.	Translated anonymously from Jean d'Arras's French text from the late fourteenth century. Text based on British

										London: Kegan Paul, Trench, Trübner.	Library Royal MS 18.B.II.
14	Middle English Humorous Tales	Various / Anonymous	MED	c.1500	15th century	1913	Post-Ricardian	comic tales / fabliaux	Profane / comic	McKnight, George H., ed. 1913. Middle English Humorous Tales in Verse. Boston: D. C. Heath.	
15	Mirror of the Life of Jesus	Nicholas Love	MED	c.1409–1410	early 15th c.	1908	Ricardian	devotional prose	Religious	Love, Nicholas. 1908. The Mirror of the Blessed Life of Jesus Christ. Edited by Lawrence F. Powell. Oxford: Clarendon Press.	I deleted everything before the prohemium.
16	Orpheus and Eurydice	Robert Henryson	METS	c.1470	late 15th c.	2010	Post-Ricardian	narrative poem / classicizing romance	Secular	Henryson, Robert. 2010. The Complete Works. Edited by David J. Parkinson. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Parkinson uses 1568 Bannatyne MS as base text. Other witnesses are 1508 Chepman and Myllar print and the early-16th-c. Asloan MS collated.
17	Pearl	Gawain-poet	MED	c.1375	c. 1390–1400	1953	Ricardian	dream vision / elegy	Religious	Gordon, E. V., ed. 1953. Pearl. Oxford: Clarendon Press.	

18	Piers Plowman	William Langland	MED	1360–1390	c. 1377-1379	1978	Ricardian	allegorical dream vision	Religious	Langland, William. 1978. The Vision of Piers Plowman: A Critical Edition of the B-Text. Edited by A. V. C. Schmidt. London and New York: J. M. Dent and E. P. Dutton.	Dates relate to the earliest manuscripts for the B-version of the text.
19	Prick of Conscience	Anonymous	MED	c.1340	14th c.	1863	Pre-Ricardian	devotional didactic poetry	Religious	Morris, Richard, ed. 1863. The Pricke of Conscience (Stimulus Conscientiæ): A Northumbrian Poem. Philological Society. Berlin: Asher.	
20	Purity (or Cleanness)	Gawain-poet	MED	c.1375	c. 1390–1400	1920	Ricardian	homiletic alliterative poetry	Religious	Menner, Robert J., ed. 1920. Purity, a Middle English Poem. Yale Studies in English 61. New Haven: Yale University Press.	
21	Reson and Sensuallyte	John Lydgate	MED	early 15th c.		1903	Post-Ricardian	dream vision / allegory	Miscellaneous	Sieper, Ernst, ed. 1901–1903. Lydgate's Reson and Sensuallyte. 2 vols. Early English Text Society, Extra Series 84, 89.	the edition is from Bodleian Fairfax MS 16 and British Museum/British Library Additional MS 29,729; publication is 1901–

										London: Kegan Paul, Trench, Trübner.	03, so do not enter only 1903
22	Ego Dormio et Cor Meum Vigilat	Richard Rolle	MED	c.1340–1360	14th c.	1896	Pre-Ricardian	devotional / mystical prose	Religious	Horstmann, Carl, ed. 1895–1896. English prose works of Richard Rolle: a selection from the edition by Carl Horstmann. London: Sonnenschein.	Horstmann gives Ego Dormio under Canonical Works, with Cambridge Dd V.64 and Rawlinson A.389 witnesses.
23	The Form of Living	Richard Rolle	MED	c.1340–1360	14th c.	1896	Pre-Ricardian	devotional / mystical prose	Religious	Horstmann, Carl, ed. 1895–1896. English prose works of Richard Rolle: a selection from the edition by Carl Horstmann. London: Sonnenschein.	Also known as The Form of Perfect Living. Horstmann's edition in the Corpus of Middle English Prose and Verse gives Cambridge Dd V.64 and Rawlinson C.285; Horstmann describes Dd V.64 as a fourteenth-century witness.
24	The Commandment	Richard Rolle	MED	c.1340–1360	14th c.	1896	Pre-Ricardian	devotional / mystical prose	Religious	Horstmann, Carl, ed. 1895–1896. English prose works of Richard Rolle: a selection from the edition by Carl Horstmann. London: Sonnenschein.	Horstmann gives The Commandment under Canonical Works, with Cambridge Dd V.64 and Rawlinson A.389 witnesses.

25	Desyre and Delit; Gastly Gladnesse	Richard Rolle	MED	c. 1340-1360	14th c.	1896	Pre-Ricardian	devotional / mystical prose	Religious	Horstmann, Carl, ed. 1895–1896. English prose works of Richard Rolle: a selection from the edition by Carl Horstmann. London: Sonnenschein.	Horstmann lists Desyre and Delit and Gastly Gladnesse separately under Short Prose Tracts; your filename combines them.
26	Sir Gawain and the Green Knight	Gawain-poet	MED	c.1375–1400	c. 1390–1400	1967	Ricardian	Arthurian romance	Secular	Gawain-Poet. 1967. Sir Gawain and the Green Knight. Edited by J. R. R. Tolkien and E. V. Gordon. Revised by Norman Davis. Oxford: Clarendon Press.	
27	The Assembly of Gods	Anonymous	MED	second half 15th c.	1498–c.1505	1895	Post-Ricardian	dream vision / allegory	Miscellaneous	Triggs, Oscar Lovell, ed. 1895. The Assembly of Gods; or, The Accord of Reason and Sensuality in the Fear of Death. University of Chicago English Studies 1. Chicago: University of Chicago Press.	Modern treatment is safer as Anonymous, formerly attributed to Lydgate. The base A text is Cambridge, Trinity College Library MS R.3.19, with early printed witnesses by Wynkyn de Worde and others.

28	The Book of Faith	Reginald Pecock	MED	mid-15th c.	15th c.	1909	Post-Ricardian	theological prose	Religious	Morison, J. L., ed. 1909. Reginald Pecock's Book of Faith: A Fifteenth Century Theological Tractate. Glasgow: James MacLehose and Sons.	The text is a fifteenth-century theological tractate edited from the manuscript in Trinity College, Cambridge; the known witness is Trinity College, Cambridge, MS B.14.45.
29	The Book of the Knight of the Tower	Geoffroy de la Tour Landry; anonymous	MED	c.1422–1461	15th c.	1906	Post-Ricardian	conduct literature / courtesy prose	Miscellaneous	Wright, Thomas, ed. 1906. The Book of the Knight of La Tour-Landry: Compiled for the Instruction of His Daughters. Revised ed. Early English Text Society, Original Series 33. London: Kegan Paul, Trench, Trübner.	The item title says the text was anonymously translated from French into English in the reign of Henry VI; catalogue records identify the source as the unique British Museum/British Library manuscript, Harley 1764, with Caxton's 1484 print used in the edition.
30	The Brut Chronicle	Anonymous	MED	late 14th–early 15th c.	c.1400	1908	Post-Ricardian	chronicle	Miscellaneous	Brie, Friedrich W. D., ed. 1906–1908. The Brut, or The Chronicles of England. 2 vols. Early English Text Society, Original Series 131, 136. London: Kegan Paul, Trench, Trübner.	Brie's edition from Bodleian MS Rawlinson B.171, "&c.", meaning the edition is not simply a single-witness diplomatic text. Rawlinson B.171 contains the Brut Chronicle to 1332, continued by another hand to 1377.

31	The Canterbury Tales	Geoffrey Chaucer	MED	1387-1400	1400-1420	1957	Ricardian	narrative poetry / tale collection	Miscellaneous	Robinson, F. N., ed. 1957. The Works of Geoffrey Chaucer. 2nd ed. Boston: Houghton Mifflin.	He uses Ellesmere and Hengwyrt as codici optimi.
32	The Eremyte and the Outelawe	Anonymous	MED	c.1375	second half of 15th c	1890	Ricardian	romance / moral tale	Religious	Kaluza, Max, ed. 1890. “Kleinere Publikationen aus mittelenglischen Handschriften. 1. The Eremyte and the Outelawe. 2. Die Boten des Todes.” Englische Studien 14: 165–182.	The source is Kaluza 1890, not French and Hale 1930. The text is preserved in BL Additional MS 37492, fols. 76v–82v, a fifteenth-century witness; BL Add. 22577 is a later transcript. Composition is probably earlier than the manuscript, often treated as late fourteenth century.
33	The Lyfe of Ipomydon	Anonymous; Hue de Rotelande	MED	second half 15th c.	late 15th c.	1889	Post-Ricardian	romance	Secular	Kölbing, Eugen, ed. 1889. Ipomedon in drei englischen Bearbeitungen. Breslau: Koebner.	The relevant Lyfe of Ipomydon version is the Harley 2252 text; Middle English Romance identifies Kölbing’s edition as edited from Harley 2252.
34	The Romance of Partenay	Anonymous; La Coudrette	MED	late 15th c.	15th c.	1899	Post-Ricardian	romance	Secular	Skeat, Walter W., ed. [1866] 1899. The Romance of Partenay, or of Lusignan. Early English Text Society, Original	From English text as translated from La Coudrette before 1500 and edited from a unique Trinity College, Cambridge manuscript.

										Series 22. Rev. Ed. London: N. Trübner.	
35	The Romance of the Sowdone of Babylone	Anonymous	MED	late 14th c. - early 15th c.	early 15th c.	1881	Pre-Ricardian	romance	Secular	Hausknecht, Emil, ed. 1881. The Romaunce of the Sowdone of Babylone and of Ferumbras His Sone. Early English Text Society, Extra Series 38. London: N. Trübner.	Hausknecht's edition as re-edited from the unique manuscript of Sir Thomas Phillipps.
36	The Storie of Asneth	Anonymous	MED	mid-15th c.	c.1500	1910	Post-Ricardian	biblical romance / conversion narrative	Religious	MacCracken, H. N., ed. 1910. "The Storie of Asneth." Journal of English and Germanic Philology 9.	Unknown Middle English translation of a lost Latin version; external manuscript evidence places the Middle English poem in one of the Ellesmere-Bridgewater manuscript tradition in the fifteenth century.
37	The Testament of Cresseid	Robert Henryson	METS	c.1480	1593	2010	Post-Ricardian	complaint / tragic sequel	Secular	Henryson, Robert. 2010. The Complete Works. Edited by David J. Parkinson. TEAMS Middle English Texts Series.	1593 Charteris print as base text; earlier Thynne 1532 survives; manuscript fragments in Dean of Lismore 1512–29 and Ruthven MS 1520–30

										Kalamazoo, MI: Medieval Institute Publications.	
38	The Three Kings of Cologne	Anonymous; John Hildesheim	MED	late 14th c. - early 15th c.	early 15th c.	1886	Pre-Ricardian	religious narrative / legend	Religious	Horstmann, Carl, ed. 1886. The Three Kings of Cologne. Early English Text Society, Original Series 85. London: N. Trübner.	Text is an early English translation of John of Hildesheim's Historia Trium Regum; Horstmann's introduction says nearly all manuscripts are fifteenth-century and that Royal 18 A X is older, beginning of the fifteenth century.
39	The Three Kings' Sons	Anonymous; David Aubert	MED	1458-1479	c. 1500	1895	Post-Ricardian	romance	Secular	Aubert, David. 1895. The Three Kings' Sons. Part I: The Text. Edited by Frederick James Furnivall. London: Kegan Paul, Trench, Trübner; Early English Text Society.	Text edited from its unique MS, Harleian 326. It is a translation from the French.
40	The Wright's Chaste Wife	Adam of Cobsam	MED	c.1462	15th c.	1869	Post-Ricardian	fabliau / comic romance / marital	Profane	Furnivall, Frederick J., ed. [1865] 1869. The Wright's Chaste Wife. Early English Text Society, Original Series 12. Rev.	Furnivall identifies copies and edits the work as by Adam of Cobsam, from a Lambeth manuscript about 1462.

										Ed. London: N. Trübner.	
41	Thornton Religious Prose	Anonymous	MED	c.1400-1450	c. 1440-1450	1914	Post-Ricardian	devotional prose	Religious	Perry, George G., ed. [1867] 1914. Religious Pieces in Prose and Verse from the Robert Thornton Manuscript. Early English Text Society, Original Series 26. Rev. Ed. London: N. Trübner.	
42	Towneley Plays	Anonymous	MED	c.1450	c. 1460-1500	1897	Post-Ricardian	drama / mystery plays	Religious	England, George, and Alfred W. Pollard, eds. 1897. The Towneley Plays. Early English Text Society, Extra Series 71. London: Kegan Paul, Trench, Trübner.	
43	Troilus and Criseyde	Geoffrey Chaucer	MED	c.1382-1386	early 15th century	1984	Ricardian	courtly romance / tragedy	Secular	Chaucer, Geoffrey. 1984. Troilus and Criseyde: A New Edition of “The Book of Troilus.” Edited by Barry A. Windeatt.	

										London and New York: Longman.	
44	York Plays	Anonymous	MED	1376-1463	c.1460	1982	Post-Ricardian	drama / mystery plays	Religious	Beadle, Richard, ed. 1982. The York Plays. London: Arnold.	
45	A Complaynte of a Lovers Lyfe	John Lydgate	METS	c.1400-1420	c. 1430-1450	2004	Post-Ricardian	Dream vision, Complaint (Poetry)	Secular	Symons, Dana M., ed. 2004. Chaucerian Dream Visions and Complaints. Kalamazoo, MI: Medieval Institute Publications. Middle English Texts Series.	Symons uses Oxford, Bodleian Library MS Fairfax 16, fols. 20v–30r, as base text. The earlier Shirley MS, BL Add. 16165, is c.1420s but is not the base text and lacks lines 610–51. Also note alternate title: The Complaint of the Black Knight.
46	Alliterative Morte Arthure	Anonymous	METS	c.1400	c. 1440	1994	Ricardian	Arthuriana, Matter of England, Nine Worthies, Romance	Secular	Benson, Larry D., and Edward E. Foster, eds. 1994. King Arthur's Death: The Middle English Stanzaic Morte Arthur and Alliterative Morte Arthure. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Benson-Foster uses the unique Lincoln Thornton Manuscript, Lincoln Cathedral MS 91, fols. 53r–98v, copied by Robert Thornton.

47	Amoryus and Cleopes	John Metham	METS	1449	15th c.	1999	Post-Ricardian	Romance, Family romance, Matter of Rome/Troy, Otherworldly	Secular	Methlam, John. 1999. Amoryus and Cleopes by John Metham, edited by Stephen F. Page. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The edition includes the final ten lines recovered through multispectral imaging of the base manuscript
48	Athelston	Anonymous	METS	late 14th c.	Early 15th c.	1999	Ricardian	Family romance, Romance, Tail rhyme	Secular	Herzman, Ronald B., Eve Salisbury, and Graham Drake, eds. 1999. Four Romances of England: King Horn, Havelok the Dane, Bevis of Hampton, Athelston. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Single manuscript witness, Cambridge, Caius College Library MS 175, fols. 120r–31r
49	Bevis of Hampton	Anonymous	METS	c.1300	c. 1330-1340	1999	Pre-Ricardian	Crusade romance, Family romance, Romance, Tail rhyme	Secular	Herzman, Ronald B., Graham Drake, and Eve Salisbury, eds. 1999. Four Romances of England: King Horn, Havelok the Dane, Bevis	This edition follows the text of the Auchinleck MS (A), fols. 176–201. The editors used E = Duke of Sutherland, now BL Egerton 2862, only to fill the missing leaf in

										of Hampton, Athelston. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Auchinleck, for edition lines 2289–2464.
50	Confessio Amantis	John Gower	METS	1386-1393	c.1390–1415	2013	Ricardian	Complaint (Poetry), Debate poetry, Dream vision, Dream vision (Boethian), Dream vision (Love), Exemplum, Matter of Rome/Troy, Romance, Utopian fiction, Apology, Advice for princes, Alexander the Great, Arthuriana, Matter of England, Nine Worthies	Secular	Gower, John. 2000-2013. Confessio Amantis, edited by Russell A. Peck. 3 vols. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The source is Fairfax 3.

51	Duke Roland and Sir Otuel of Spain	Anonymous	METS	late 14th c.	mid-15th c.	2019	Ricardian	Charlemagne, Emperor, Crusade romance, Matter of France, Nine Worthies, Roland, Romance, Tail rhyme	Secular	Fein, Susanna, Elizabeth Melick, and David Raybin, eds. 2019. The Roland and Otuel Romances and the Anglo-Norman Otinel. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	single witness, London, BL MS Additional 31042, the London Thornton MS, fols. 82r–94r, copied in the mid-15th c.
52	Emaré	Anonymous	METS	late 14th c.	first half of 15th c.	1995	Pre-Ricardian	Family romance, Breton lai, Romance, Tail rhyme, Lay (Narrative poetry)	Secular	Laskaya, Anne, and Eve Salisbury, eds. 1995. The Middle English Breton Lays. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	single witness, BL Cotton Caligula A.ii;
53	Erle of Tolous	Anonymous	METS	c.1350	c. 1420-1450	1995	Pre-Ricardian	Breton lai, Family romance, Romance, Lay (Narrative poetry)	Secular	Laskaya, Anne, and Eve Salisbury, eds. 1995. The Middle English Breton Lays. TEAMS Middle English Texts Series. Kalamazoo, MI:	Laskaya-Salisbury uses the Cambridge University Library MS Ff.2.38 tradition for the edited text; source date c.1420–50 is defensible.

										Medieval Institute Publications.	
54	Fabula Duorum Mercatorum	John Lydgate	METS	c. 1420	1460 c.	2016	Post-Ricardian	Romance, Fable, Exemplum	Secular	Lydgate, John. 2016. John Lydgate: Fabula duorum mercatorum and Guy of Warwyk, edited by Pamela Farvolden and Benjamin Garstad. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Farvolden bases the edition on London, BL MS Harley 2255, dated to the 1460s
55	Fortunes Stabilnes	Charles d'Orleans	METS	15th c.	15th c.	2005	Post-Ricardian	Ballad, Chaucer, Geoffrey, Chaucerian rondel, Dream vision, Dream vision; courtly love	Secular	Mooney, Linne R., and Mary-Jo Arn, eds. 2005. The Kingis Quair and Other Prison Poems. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Charles d'Orleans, from Fortunes Stabilnes. The manuscript abbreviation is London, BL MS Harley 682.

56	Guy of Warwyk	John Lydgate	METS	c.1420s	mid-15th c.	2016	Post-Ricardian	Romance, Chronicle, Matter of England, Hagiographic romance, Ballad	Secular	Lydgate, John. 2016. Lydgate's Fabula duorum mercatorum and Guy of Warwyk, edited by Pamela Farvolden. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	This is Lydgate's short Guy of Warwyk, probably 1420s, not the older anonymous Guy tradition. Farvolden uses Oxford, Bodleian Library, MS Laud Misc. 683 as the text for the present edition; Guy appears on fols. 65–78. The manuscript has a fifteenth-century/Bury-St Edmunds context, with a blanket attribution note after 1449, so use c.1450 / mid-15th c. as the source date.
57	Holkham Prayers and Meditations	Anonymous	METS	c.1400-1420	c. 1430-1450	2023	Post-Ricardian	Biblical exegesis, Prayer	Religious	Parsons, Ben, ed. 2023. Two Middle English Prayer Cycles: Holkham Prayers and Meditations and Simon Appulby, The Fruyte of Redempcyon. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The text survives in a single manuscript, Bodleian MS Holkham Misc. 41.

58	Lybeaus Desconus	Anonymous / Thomas Chestre?	METS	mid-14th c.	15th c.	2013	Ricardian	Tail rhyme, Romance	Secular	Salisbury, Eve, and James Weldon, eds. 2013. Lybeaus Desconus. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	This text is from Naples, Biblioteca Nazionale, MS XIII.B.29. Note that there are two manuscript versions. The other is Lambeth Palace MS 306
59	Octavian	Anonymous	METS	c. 1330-40	early 14th c.	2006	Pre-Ricardian	Family romance, Romance, Tail rhyme, Authorship uncertain	Secular	Hudson, Harriet, ed. 2006. Four Middle English Romances: Sir Isumbras, Octavian, Sir Eglamour of Artois, Sir Tryamour. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Edition uses the Northern Octavian from Lincoln Cathedral Library MS 91, 1430–40, as base text, with missing lines supplied from Cambridge Ff.2.38
60	Otuel A Knight	Anonymous	METS	c. 1330-40	early 14th c.	2019	Pre-Ricardian	Crusade romance, Hagiographic romance, Matter of France, Nine Worthies, Romance, Charlemagn	Secular	Melick, Elizabeth, Susanna Fein, and David Raybin, eds. 2019. The Roland and Otuel Romances and the Anglo-Norman Otinel. TEAMS Middle	Edition contains the following texts: Otuel A Knight, Otuel and Roland, Roland and Vernagu.

								e, Emperor, Roland		English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	
61	Otuel and Roland	Anonymous	METS	mid-14th c.	second half of 15th c.	2019	Ricardian	Charlemagne, Emperor, Crusade romance, Matter of France, Nine Worthies, Roland, Romance, Tail rhyme	Secular	Melick, Elizabeth, Susanna Fein, and David Raybin, eds. 2019. The Roland and Otuel Romances and the Anglo-Norman Otinel. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	It uses London, British Library MS Additional 37492, the Fillingham MS, fols. 30v–76r, copied in the second half of the fifteenth century. Note that several folios are missing/damaged, and the plot can partly be filled from Otinel and Ellis's early summary
62	Paris and Vienne	William Caxton; Pierre de la Cépède	METS	15th c.	1485	2023	Post-Ricardian	Family romance, Romance	Secular	Caxton, William. 2023. William Caxton's Paris and Vienne and Blanchardyn and Eglantine, edited by Harriet Hudson. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The edition is based on the sole surviving complete copy: London, British Library, C.10.b.10. This is Caxton's 1485 print.

63	Prose Merlin	Anonymous; Robert de Boron	METS	15th c.	c.1400	1998	Post-Ricardian	Arthuriana, Romance, Matter of England, Otherworldly, Epic (Prose), Nine Worthies	Secular	Conlee, John W., ed. 1998. Prose Merlin. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The Middle English Prose Merlin survives in one manuscript, Cambridge University Library MS Ff.3.11, and was written near the middle of the fifteenth century after Robert de Boron / Old French Vulgate Cycle.
64	Revelations of Divine Love	Julian of Norwich	METS	c.1373-1393	c. 1650	1994	Post-Ricardian	mystical theology / devotional prose	Religious	Julian of Norwich. 1994. The Shewings of Julian of Norwich, edited by Georgia Ronan Crampton. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	This is based on the long version composed in c.1393. The source is BL MS Sloane 2499, which was probably copied c. 1650 by Mother Anne Clementina Cary of the Paris convent founded as the sister house of Cambrai in 1651.
65	Richard Coer de Lyon	Anonymous	METS	early 14th c.	1400–1450	2015	Pre-Ricardian	Crusade romance, Chanson de geste, Romance	Secular	Larkin, Peter, ed. 2015. Richard Coer de Lyon. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Larkin is based on the a-version, specifically C, with W 1509 used for gaps. Do not use Auchinleck 1330–40 as the METS source date here; that is the oldest extant witness, not the base text.

66	Robert of Cisyle	Anonymous	METS	late 14th c. - early 15th c.	c.1390	2007	Ricardian	Exemplum, Hagiographic romance, Romance	Secular	Foster, Edward E., ed. 2007. Amis and Amiloun, Robert of Cisyle, and Sir Amadace. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The text is based on V, the Vernon Manuscript, with some modifications from previous editions.
67	Roland and Vernagu	Anonymous	METS	c. 1130	1330s	2019	Pre-Ricardian	Charlemagne, Emperor, Crusade romance, Matter of France, Nine Worthies, Roland, Romance, Tail rhyme	Secular	Melick, Elizabeth, Susanna Fein, and David Raybin, eds. 2019. The Roland and Otuel Romances and the Anglo-Norman Otinel. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	It has a sole manuscript witness, Auchinleck, copied in the 1330s.
68	Salutations	John the Blind Audelay	METS	c.1425-1431	1425–1431	2009	Post-Ricardian	Prayer	Religious	Fein, Susanna, ed. 2009. Poems and Carols (Oxford, Bodleian Library MS Douce 203). TEAMS Middle English Texts Series.	This is part of the codex edition Poems and Carols (Oxford, Bodleian Library MS Douce 302).

										Kalamazoo, MI: Medieval Institute Publications.	
69	Siege of Jerusalem	Anonymous	METS	1370s–1380s	late 14th c.	2004	Ricardian	Hagiographic romance, Romance, Vengeance of Our Lord literature	Secular	Livingston, Michael, ed. 2004. The Siege of Jerusalem. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The edition uses Laud Misc. 656 as base text, generally emended by P, A, and D.
70	Sir Amadace	Anonymous	METS	late 14th c.	1450–60	2007	Ricardian	Family romance, Romance	Secular	Foster, Edward E., ed. 2007. Amis and Amiloun, Robert of Cisyle, and Sir Amadace. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The edition takes the text from the Ireland Manuscript
71	Sir Cleges	Anonymous	METS	early 15th c.	c.1425–1430	1995	Ricardian	Breton lai, Family romance, Hagiographic romance, Romance, Lay	Secular	Laskaya, Anne, and Eve Salisbury, eds. 1995. The Middle English Breton Lays. TEAMS Middle English Texts	The editor chose the Bodleian MS/Ashmole 61 because it is more complete.

								(Narrative poetry)		Series. Kalamazoo, MI: Medieval Institute Publications.	
72	Sir Degrevant	Anonymous	METS	c.1385-1410	before 1461	2005	Ricardian	Romance	Secular	Kooper, Erik, ed. 2005. Sentimental and Humorous Romances. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	
73	Sir Eglamour of Artois	Anonymous	METS	c. 1350	c. 1450	2006	Pre-Ricardian	Romance, Tail rhyme	Secular	Hudson, Harriet, ed. 2006. Four Middle English Romances: Sir Isumbras, Octavian, Sir Eglamour of Artois, Sir Tryamour. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	

74	Sir Gawain and the Carle of Carlisle	Anonymous	METS	c.1400	c.1460	1995	Ricardian	Arthuriana, Matter of England, Nine Worthies, Romance	Secular	Hahn, Thomas, ed. 1995. Sir Gawain: Eleven Romances and Tales. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Hahn uses the single surviving manuscript: National Library of Wales, Porkington MS 10, also known as Harlech MS 10 or Brogyntyn MS.
75	Sir Gowther	Anonymous	METS	c.1330	c.1480	1995	Ricardian	Hagiographic romance, Crusade romance, Romance, Breton lai, Family romance, Lay (Narrative poetry)	Secular	Laskaya, Anne, and Eve Salisbury, eds. 1995. The Middle English Breton Lays. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The text survives in two late fifteenth-century manuscripts and presents the Advocates version. The editors present the Advocates version, based primarily on National Library of Scotland MS Advocates 19.3.1, fols. 11a–28a. The NLS catalogue dates this manuscript to c.1480. British Library Royal MS 17.B.43 is used supplementally where Advocates is defective, especially at the opening.
76	Sir Isumbras	Anonymous	METS	c.1320	1425-50	2006	Pre-Ricardian	Crusade romance, Hagiographic romance, Romance	Secular	Hudson, Harriet, ed. 2006. Four Middle English Romances: Sir Isumbras, Octavian, Sir Eglamour of	The romance was already circulating before 1320, while the present edition is based on Cambridge, Gonville and Caius

										Artois, Sir Tryamour. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	College MS 175, dated 1425–50.
77	Sir Launfal	Thomas Chestre	METS	c.1380	early 15th c.	1995	Ricardian	Breton lai, Otherworldly, Romance, Arthuriana, Nine Worthies, Matter of England, Lay (Narrative poetry)	Secular	Laskaya, Anne, and Eve Salisbury, eds. 1995. The Middle English Breton Lays. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Text written in the late fourteenth century and is preserved in British Library MS Cotton Caligula A.ii.
78	Sir Orfeo	Anonymous	METS	late 13th / early 14th c.	1330-40	1995	Pre-Ricardian	Breton lai, Romance, Otherworldly, Lay (Narrative poetry)	Secular	Laskaya, Anne, and Eve Salisbury, eds. 1995. The Middle English Breton Lays. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	From the earliest Middle English version: the Auchinleck manuscript

79	Sir Owain	Anonymous	METS	early 14th c.	1330-40	2004	Pre-Ricardian	Otherworldly, Purgatory, Romance	Secular	Foster, Edward E., ed. 2004. Amis and Amiloun, Robert of Cisyle, and Sir Amadace. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	From the stanzaic Auchinleck version
80	Sir Perceval of Galles	Anonymous	METS	early 14th c.	c. 1440	1995	Pre-Ricardian	Matter of England, Arthuriana, Nine Worthies, Romance, Tail rhyme	Secular	Braswell, Mary Flowers, ed. 1995. Sir Perceval of Galles and Ywain and Gawain. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The poet probably worked in the north-east Midlands in the first half of the fourteenth century, but the edition is based on the Thornton MS, Lincoln Cathedral MS 91, c.1440.
81	Sir Torrent of Portingale	Anonymous	METS	c.1400		2017	Ricardian	Family romance, Crusade romance, Romance, Tail rhyme	Secular	Wade, James, ed. 2017. Sir Torrent of Portingale. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Wade says the romance was likely composed near 1400, but the base text is Manchester, Chetham's Library MS 8009, a late fifteenth-century manuscript. The edition is based on Chetham MS 8009, with readings from

											early print fragments where persuasive
82	Sir Tryamour	Anonymous	METS	late 14th c.	1450-70	2006	Ricardian	Romance, Tail rhyme	Secular	Hudson, Harriet, ed. 2006. Four Middle English Romances: Sir Isumbras, Octavian, Sir Eglamour of Artois, Sir Tryamour. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	
83	Stanzaic Guy of Warwick	Anonymous	METS	early 14th c.	c. 1330	2004	Pre-Ricardian	Hagiographic romance, Romance, Tail rhyme, Vengeance of Our Lord literature	Secular	Wiggins, Alison, ed. 2004. Stanzaic Guy of Warwick. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Texts from the fourteenth-century Auchinleck Manuscript

84	Stanzaic Morte Arthur	Anonymous	METS	mid-14th c.	late 14th c.	1994	Pre-Ricardian	Arthuriana, Matter of England, Nine Worthies, Romance	Secular	Benson, Larry D., and Edward E. Foster, eds. 1994. King Arthur's Death: The Middle English Stanzaic Morte Arthur and Alliterative Morte Arthure. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Benson-Foster uses BL Harley 2252, fols. 86a–133b
85	The Assembly of Ladies	Anonymous	METS	c.1470-80	c.1470-80	1990	Post-Ricardian	Dream vision, Dream vision (Love), Complaint (Poetry)	Secular	Pearsall, Derek, ed. 1990. The Floure and the Leafe, the Assembly of Ladies, the Isle of Ladies. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	All three MSS are last quarter 15th c.; BL Additional 34360 is the earliest/best and is followed through Pearsall's text.
86	The Avowyng of Arthur	Anonymous	METS	late 14th c	c.1450–75	1995	Ricardian	Arthuriana, Matter of England, Nine Worthies, Romance	Secular	Hahn, Thomas, ed. 1995. Sir Gawain: Eleven Romances and Tales. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval	Hahn uses the single surviving witness, Ireland Blackburn MS, now Princeton, Robert H. Taylor Collection; the extant version is third quarter 15th c.

										Institute Publications.	
87	The Awntyrs off Arthur	Anonymous	METS	c.1375	15th c.	1995	Ricardian	Arthuriana, Matter of England, Nine Worthies, Romance	Secular	Hahn, Thomas, ed. 1995. Sir Gawain: Eleven Romances and Tales. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Hahn uses the manuscript tradition headed by Oxford, Bodleian MS Douce 324, with other witnesses including Lambeth 491.B, Thornton, and Ireland Blackburn.
88	The Boke of Cupide, God of Love, or The Cuckoo and the Nightingale	John Clanvowe	METS	late 1380s / before 1391	1430–50	1975	Ricardian	Chaucer, Geoffrey, Dream vision, Dream vision (Love), Debate poetry	Secular	Scattergood, V. J., ed. 1975. The Works of Sir John Clanvowe. Cambridge: D. S. Brewer.	The edition is based on Oxford, Bodleian MS Fairfax 16, dated c.1430–50
89	The Book of Margery Kempe	Margery Kempe	METS	late 1430s	mid-15th c.	1996	Post-Ricardian	spiritual autobiography / devotional prose	Religious	Kempe, Margery. 1996. The Book of Margery Kempe. Edited by Lynn Staley. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Source manuscript is the unique BL Additional 61823

90	The Cloud of Unknowing	Anonymous	METS	late 14th c.	15th c.	1997	Ricardian	Epistolary prose, Mysticism, Prayer, Sermon	Religious	Anonymous. 1997. The Cloud of Unknowing. Edited by Patrick J. Gallacher. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	
91	The Fruyte of Redempcyon	Anonymous ; Simon Appulby	METS	1514	1514	2023	Post-Ricardian	Biblical exegesis, Prayer	Religious	Parsons, Ben, ed. 2023. Two Middle English Prayer Cycles: Holkham Prayers and Meditations and Simon Appulby, The Fruyte of Redempcyon. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Wynkyn de Worde, 1514, Cambridge UL Sel 5.31
92	The Isle of Ladies	Anonymous	METS	late 15th c.	mid-16th c.	1990	Post-Ricardian	Dream vision, Dream vision (Love)	Secular	Pearsall, Derek, ed. 1990. The Floure and the Leafe, the Assembly of Ladies, the Isle of Ladies. TEAMS Middle English Texts Series. Kalamazoo, MI:	The best witness is Longleat House MS 256, mid-16th c.; the poem itself probably dates late in the previous century.

										Medieval Institute Publications.	
93	The Jeaste of Sir Gawain	Anonymous	METS	second half 15th c.	1564	1995	Ricardian	Arthuriana, Matter of England, Nine Worthies, Romance	Secular	Hahn, Thomas, ed. 1995. Sir Gawain: Eleven Romances and Tales. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The surviving handwritten version is dated 1564, probably from a lost John Kynge print of 1557/58
94	The King of Tars	Anonymous	METS	early 14th c.	1330–40	2015	Pre-Ricardian	Romance, Tail rhyme, Hagiographic romance, Dream vision	Secular	Chandler, Venetia, ed. 2015. The King of Tars. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Chandler bases the edition on Auchinleck, with Vernon variants separately given
95	The Knightly Tale of Gologras and Gawain	Anonymous	METS	before 1508	1508	1995	Post-Ricardian	Arthuriana, Matter of England, Nine Worthies, Romance	Secular	Hahn, Thomas, ed. 1995. Sir Gawain: Eleven Romances and Tales. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	No manuscript survives. Hahn works from the Chepman and Myllar print, 8 April 1508.

96	The Marriage of Sir Gawain	Anonymous	METS	c.1400	c.1650	1995	Post-Ricardian	Arthuriana, Matter of England, Nine Worthies, Romance	Secular	Hahn, Thomas, ed. 1995. Sir Gawain: Eleven Romances and Tales. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Hahn uses the Percy Folio manuscript, BL Additional MS 27879; this is a post-medieval witness.
97	The Quare of Jelusy	Anonymous	METS	larte 15th c.	late 15th / ealry 16th c.	2004	Post-Ricardian	Chaucer, Geoffrey, Dream vision, Dream vision (Love), Exemplum	Secular	Mooney, Linne R., and Mary-Jo Arn, eds. 2012. The Kingis Quair and Other Prison Poems. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The edition is from single copy in Oxford, Bodleian MS Arch. Selden B.24
98	The Scale of Perfection	Walter Hilton	METS	before 1396	before 1416	2000	Ricardian	Biblical exegesis, Mysticism, Prayer	Religious	Hilton, Walter. 2000. The Scale of Perfection. Edited by Thomas H. Bestul. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Hilton died in 1396, so c.1390–1405 is too broad for composition. Bestul uses London, Lambeth Palace MS 472, made for John Killum, who died in 1416.

99	The Siege of Milan	Anonymous	METS	second half 14th c.	c. 1450	1990	Ricardian	Charlemagne, Emperor, Matter of France, Nine Worthies, Romance, Roland	Secular	Lupack, Alan, ed. 1990. Three Middle English Charlemagne Romances. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	The poem is in BL Additional 31042, the London Thornton MS, dated c.1450, but the poem was probably written in the second half of the fourteenth century.
100	The Siege of Thebes	John Lydgate	METS	c.1421-1422	c. 1430	2001	Post-Ricardian	Epic Poetry, Romance, Matter of Rome/Troy, Chaucer, Geoffrey	Secular	Lydgate, John. 2001. The Siege of Thebes. Edited by Robert R. Edwards. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Edwards uses BL Arundel 119 as base manuscript, copied about 1430.
101	The Squire of Low Degree	Anonymous	METS	late 15th c.	c. 1560	2005	Post-Ricardian	Family romance, Romance	Secular	Kooper, Erik, ed. 2005. Sentimental and Humorous Romances. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Kooper presents the text from William Copland's c.1560 print, with Wynkyn de Worde fragments and Percy Folio as related witnesses.

102	The Sultan of Babylon	Anonymous	METS	late 14th / early 15th c.	mid-15th c.	2005	Ricardian	Charlemagne, Emperor, Matter of France, Nine Worthies, Romance, Roland	Secular	Hardman, Phillipa, ed. 2005. The Siege of Milan, The Sultan of Babylon, Octavian. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	the romance survives in one manuscript, Princeton Garrett MS 140, dated mid-15th c.; the poem itself is late 14th or early 15th c.
103	The Temple of Glas	John Lydgate	METS	first half 15th c.	1440s	2007	Post-Ricardian	Dream vision (Love)	Secular	Lydgate, John. 2007. The Temple of Glas. Edited by J. Allan Mitchell. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Mitchell uses Oxford, Bodleian MS Tanner 346 as base text; Tanner is dated to the second quarter of the fifteenth century, with the manuscript list specifying Tanner 346 as 1440s.
104	The Tournament of Tottenham	Anonymous	METS	c.1450	c. 1455	2005	Post-Ricardian	Romance	Secular	Kooper, Erik, ed. 2005. Sentimental and Humorous Romances. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Kooper uses BL Harley MS 5396 as base text; the same hand includes memoranda dated 1455

105	The Wedding of Sir Gawain and Dame Ragnelle	Anonymous	METS	ate 15th c.	16th c.	1995	Post-Ricardian	Arthuriana, Matter of England, Nine Worthies, Romance	Secular	Hahn, Thomas, ed. 1995. Sir Gawain: Eleven Romances and Tales. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	Hahn says the text survives in a sixteenth-century Bodleian MS 11951.
106	Ywain and Gawain	Gawain-poet	METS	c.1300	c. 1375–1400	1996	Pre-Ricardian	Arthuriana, Matter of England, Nine Worthies, Romance	Secular	Richmond, Velma Bourgeois, ed. 1996. Ywain and Gawain. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	
107	Boke of the Duchess	Geoffrey Chaucer	Other Edition	c. 1368-1372	1425-1500; 1532	1997	Ricardian	Dream vision, elegy	Secular	Chaucer, Geoffrey. 1997. Geoffrey Chaucer's Book of the Duchess: A Hypertext Edition. Edited by Murray McGillivray. CD-ROM. Calgary: University of Calgary Press.	The edition is made from the three manuscript and the sole printed edition. These are, respectively: Bodleian Library, MS Bodley 638, dated c. 1425–1450; Bodleian Library, MS Fairfax 16, dated c. 1450–1475; Glasgow University Library, MS Hunter 197, dated c. 1475–1500; and Thynne's book, printed in 1532.

108	The Parliament of Fowls	Geoffrey Chaucer	Other Edition	c.1381–1382	c.15th century	1987	Ricardian	Dream vision, courtly allegory	Secular	Benson, Larry D., ed. 1987. The Riverside Chaucer. 3rd ed. Boston: Houghton Mifflin.	
109	Romaunt of the Rose	Geoffrey Chaucer; Guillaume de Lorris	MED	c.1230 (original French) / c.1360–1380 (Chaucer translation)	1440–1450	1891	Ricardian	Romance, dream vision	Secular	Chaucer, Geoffrey. 1891. The Romaunt of the Rose. Edited by Max Kaluza. London: Trübner; Chaucer Society.	This is the Middle English translation of Guillaume de Lorris' part. The edition follows the unique Glasgow, University of Glasgow Library, Hunter Ms. 409 (olim ms. V. 3. 7, Hunterian Museum, Glasgow). Eleven leaves are lost so the deficiency was supplied from Thynne's 1532 edition.
110	Patience	Gawain-poet (Pearl Poet)	Other Edition	c.1375–1400	c. 1390–1400	1996	Ricardian	Alliterative religious poem, homily	Religious	Shawver, Gary, ed. 1996. "Using TACT and Electronic Texts." In TACT: Text-Analysis Computing Tools, Version 2.1 for MS-DOS and PC DOS, by Ian Lancashire et al. CD-ROM. New York: Modern Language	

										Association of America.	
111	Legend of Good Women	Geoffrey Chaucer	Other Edition	c.1386–1388	15th century	1987	Ricardian	Dream vision, courtly legend	Secular	Benson, Larry D., ed. 1987. The Riverside Chaucer. 3rd ed. Boston: Houghton Mifflin.	
112	Iak and His Step Dame	Anonymous	MED	late 14th century (c. 1375–1400)	15th century	1893	Ricardian	fabliau / comic romance	Secular	Zupitza, Julius, ed. 1893. Iak and His Step Dame. Braunschweig: George Westermann.	From the single codex: Oxford, Bodleian Library MS. Rawl. C. 86
113	Poems in the courtly tradition	William Dunbar	METS	15th century	15th century	2004	Post-Ricardian	dream vision, courtly love	Secular	Dunbar, William. 2004. William Dunbar: The Complete Works, edited by John Conlee. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	it contains the following poems: Sweet Rose of Virtue, Beauty and the Prisoner, To a Lady, Good Counsel for Lovers, The Golden Targe, The Merle and the Nightingale, Love's Inconstancy, True Love
114	The Thrissil and the Rois	William Dunbar	Other Edition	1503	1568	1907	Post-Ricardian	Allegorical dream vision; nuptial poem	Secular	Baildon, H. Bellyse, ed. 1907. The Poems of William Dunbar.	For the occasion of the marriage of James IV and Margaret Tudor in August 1503.

										Cambridge: Cambridge University Press.	
115	Pierce the Plowman's Crede	Anonymous	MED	1393-1401	14th c.	1867	Ricardian	alliterative satirical poem	Secular	Skeat, Walter W., ed. 1867. Pierce the Ploughmans Crede. Early English Text Society, Original Series 30. London: N. Trübner. Electronic edition in the Corpus of Middle English Prose and Verse, University of Michigan.	Distinct from Langland's Piers Plowman. The Skeat edition is based on Trinity College Cambridge MS R.3.15 and collated with British Library MS Bibl. Reg. 18.B.XVII and the old printed text of 1553
116	The Kingis Quair	James I of Scotland	METS	c. 1423–1424	Late 15th c.	2005	Post-Ricardian	Boethian dream vision, complaint	Secular	Mooney, Linne R., and Mary-Jo Arn. 2005.The Kingis Quair and Other Prison Poems. TEAMS Middle English Texts Series. Kalamazoo, MI: Medieval Institute Publications.	

117	La Belle dame sans merci	Richard Ross / Alain Chartier	METS	c. 1450-1460	1532	1897	Post-Ricardian	complaint, courtly love	Secular	Ros, Richard. 1897. “La Belle Dame sans Mercy.” In The Complete Works of Geoffrey Chaucer. Vol. 7: Chaucerian and Other Pieces, edited by Walter W. Skeat, 299–326. Oxford: Clarendon Press. Electronic text in the Corpus of Middle English Prose and Verse, University of Michigan.	Skeat uses Thynne 1532 as base, controlled by MSS. English translation of Alain Chartier’s French La Belle Dame sans Mercy.
118	The Testament of Love	Thomas Usk	MED	1385-86	1534	1897	Ricardian	dream vision, boethian consolation, philosophical prose	Secular	Usk, Thomas. 1897. “The Testament of Love.” In Chaucerian and Other Pieces, vol. 7 of The Complete Works of Geoffrey Chaucer, edited by Walter W. Skeat, 1–145. Oxford: Clarendon Press. Electronic text in the Corpus of Middle English Prose and Verse, University of Michigan.	

119	Two Tapestry Poems: The Life of St George and The Falls of Seven Princes	John Lydgate	MED	c.1430s	1430–1450	1911	Post-Ricardian	religious legend / de casibus complaint / public poetry	Miscellaneous	Hammond, Eleanor Prescott, ed. 1910–1911. “Two Tapestry Poems by Lydgate: The Life of St George and the Falls of Seven Princes.” Englische Studien 43: 10–26.	The file contains The Life of St George followed by The Falls of Seven Princes. The Falls section is explicitly from Cambridge, Trinity College MS R.3.20; that manuscript is associated with John Shirley.
120	Gest Historiale	Anonymous / Guido delle Colonne	MED	second half 14th c.	15th c.	1874	Ricardian	alliterative Trojan romance / Matter of Troy / romance	Secular	Panton, George A., and David Donaldson, eds. 1869–1874. The “Gest Hystoriale” of the Destruction of Troy: An Alliterative Romance Translated from Guido de Colonna’s “Hystoria Troiana.” Early English Text Society, Original Series 39 and 56. London: N. Trübner.	The text is an alliterative romance translated from Guido delle Colonne’s Latin Historia Troiana, first edited from the unique Glasgow, Glasgow University Library MS Hunterian 388 manuscript. The Middle English translation appears to have been made in the North or North-West Midlands in the second half of the fourteenth century and survives only in this imperfect fifteenth-century Hunterian manuscript.

121	The Epistle of Othea to Hector; or, The Boke of Knyghthode	Christine de Pizan / Stephen Scrope	MED	c.1400	mid-15th c.	1904	Post-Ricardian	chivalric conduct literature / moral allegory / advice for knights	Miscellaneous	Scrope, Stephen. 1904. The Epistle of Othea to Hector; or, The Boke of Knyghthode. Translated from the French of Christine de Pisan, with a Dedication to Sir John Fastolf, K.G., edited by George F. Warner. London: J. B. Nichols.	The edition begins with Scrope's dedication to an aged knight and says he has translated "out of French tongue" a "Book of Knighthood"; the main text then begins with Othea addressing Hector. The Corpus of Middle English Prose and Verse item identifies the text as Christine de Pizan's Othea, translated by Stephen Scrope, dedicated to Sir John Fastolf, and edited by George F. Warner from a manuscript in the Library of the Marquis of Bath. Warner edits from Longleat MS 253; his introduction says the manuscript hand is about the middle of the fifteenth century.

APPENDIX B: ANNOTATED CODE

```
# -*- coding: utf-8 -*-
“““Thesis_Appendix_B.ipynb

Automatically generated by Colab.

Original file is located at
  https://colab.research.google.com/drive/19wZrtv4PyH73RIWE8boT0qNj9-lGhRim

# modernisation procedure
“““

# Imports
import os
import time
import csv
from datetime import datetime, timedelta
from google import genai
from google.genai import types
from google.colab import userdata

# Configuration
WORK_DIR = “/content/raw_corpus”          # folder containing original .txt corpus
OUTPUT_DIR = “/content/normalized_corpus”    # where normalized files will be saved
LOG_PATH = “/content/normalization_log.csv”   # CSV file to track progress/errors

MODEL = “gemini-3.1-flash-lite-preview”
COOLDOWN = 10                    # delay (seconds) between API calls
MAX_RETRIES = 5                  # how many times to retry failed API calls
BACKOFF_BASE = 5                 # wait grows exponentially at each retry

os.makedirs(OUTPUT_DIR, exist_ok=True)

# Gemini client
from google.colab import userdata
from google import genai

GEMINI_API_KEY = userdata.get(“GEMINI_API_KEY”)
client = genai.Client(api_key=GEMINI_API_KEY)

# Collect all .txt files in corpus
all_files = sorted([
  f for f in os.listdir(WORK_DIR)
  if f.endswith(“.txt”) and os.path.isfile(os.path.join(WORK_DIR, f))
])
total_files = len(all_files)
print(f”Found {total_files} text files in corpus.\n”)

# CSV log setup
log_rows = []

def write_log():
  “““Write current log_rows list to CSV file.”””
  with open(LOG_PATH, “w”, newline=“”, encoding=“utf-8”) as f:
```

```
    writer = csv.DictWriter(f, fieldnames=[
      “timestamp”, “filename”, “char_count”, “status”, “retries”, “error”
    ])
    writer.writeheader()
    writer.writerows(log_rows)

# Prompt
def build_prompt(filename, text):
  “““Build the XML-structured prompt sent to Gemini.”””
  return f”””
<context>
 <project>Middle English corpus normalization for comparison</project>
 <goal>
  Convert each token in the source text into a single comparable target form.
  Prefer a modern English reflex when one exists.
  If no direct modern English reflex exists, use a minimal semantic translation.
  The process must remain strictly word-by-word.
 </goal>
 <file_name>{filename}</file_name>
</context>

<task>
 Produce a strictly word-by-word normalized version of the entire input text.
</task>

<instruction>
 For every source token, output exactly one normalized token whenever possible:
 first a modern English reflex, otherwise a one-word semantic equivalent.
</instruction>

<rules>
 <rule>Work strictly word by word.</rule>
 <rule>Preserve the original word order exactly.</rule>
 <rule>Do not summarize, paraphrase, or restructure clauses or sentences.</rule>
 <rule>For each source token, first choose a single modern English reflex if a clear correspondence exists.</rule>
 <rule>If no clear modern English reflex exists, choose the closest semantic translation.</rule>
 <rule>Map each input word to a single output word whenever possible.</rule>
 <rule>Do not omit content.</rule>
 <rule>Do not merge multiple source words into one unless absolutely necessary.</rule>
 <rule>Do not expand one source word into multiple output words unless absolutely necessary.</rule>
 <rule>Keep equivalent forms identical across the whole text.</rule>
 <rule>Normalize dialectal, orthographic, and inflectional variants to the same output form.</rule>
 <rule>Examples: ys -> is; es -> is; luf -> love; luv -> love; loue -> love; sal -> shall; fra -> from; kirk -> church;
bairn -> child.</rule>
 <rule>If a token is ambiguous, choose the most likely equivalent from context, but still output only one word if
possible.</rule>
 <rule>Preserve punctuation and line breaks as much as possible.</rule>
 <rule>Do not add notes, commentary, brackets, XML, JSON, markdown, or explanations.</rule>
 <rule>Return only the transformed text.</rule>
</rules>

<input_text>
{text}
</input_text>
“””.strip()
```

```
# API call
def normalize_file(filename, text):
  """
  Send the text to Gemini for normalization.
  Retry failures with exponential backoff.
  Return (normalized_text, retries_used).
  """
  prompt = build_prompt(filename, text)

  for attempt in range(MAX_RETRIES):
    try:
      response = client.models.generate_content(
        model=MODEL,
        contents=prompt,
        config=types.GenerateContentConfig(
          temperature=0.1
        )
      )
      result = (response.text or "").strip()

      if not result:
        raise ValueError("Empty response from API")

      ratio = len(result) / max(len(text), 1)
      if ratio < 0.5 or ratio > 1.5:
        print(f"        ⚠ Suspicious length ratio: {ratio:.2f} — proceeding anyway")

      return result, attempt

    except Exception as exc:
      wait = BACKOFF_BASE ** attempt
      print(f"        ⚠ Attempt {attempt + 1}/{MAX_RETRIES} failed: {exc}. Retrying in {wait}s…")
      time.sleep(wait)

  raise RuntimeError(f"All {MAX_RETRIES} attempts failed for {filename}")

# Main loop
run_start = datetime.now()
done_count = 0

for file_num, filename in enumerate(all_files, start=1):
  path = os.path.join(WORK_DIR, filename)
  out_path = os.path.join(OUTPUT_DIR, filename.replace(".txt", "_normalized.txt"))

  if os.path.exists(out_path):
    print(f"[{file_num}/{total_files}] Skipping (done): {filename}")
    continue

  elapsed = (datetime.now() - run_start).total_seconds()
  if elapsed > 0 and done_count > 0:
    rate = done_count / elapsed
    remaining = (total_files - file_num + 1) / rate
    eta_str = str(timedelta(seconds=int(remaining)))
  else:
    eta_str = "calculating…"
```

```
  print(f"[{file_num}/{total_files}] {filename} | ETA: {eta_str}")

  with open(path, "r", encoding="utf-8", errors="ignore") as f:
    text = f.read()

  print(f" Characters: {len(text):,}")

  status, error_msg, retries = "success", "", 0
  try:
    normalized, retries = normalize_file(filename, text)

    with open(out_path, "w", encoding="utf-8") as f:
      f.write(normalized)

    print(f" ✓ Saved: {os.path.basename(out_path)}")
    done_count += 1

  except Exception as exc:
    status = "failed"
    error_msg = str(exc)
    print(f" ✗ Failed permanently: {exc}")

  log_rows.append({
    "timestamp": datetime.now().isoformat(timespec="seconds"),
    "filename": filename,
    "char_count": len(text),
    "status": status,
    "retries": retries,
    "error": error_msg,
  })
  write_log()

  if file_num < total_files:
    time.sleep(COOLDOWN)

# Summary
total_time = datetime.now() - run_start
succeeded = sum(1 for r in log_rows if r["status"] == "success")
failed = sum(1 for r in log_rows if r["status"] == "failed")

print(f"\nDone in {total_time}.")
print(f"Processed : {done_count} files this session")
print(f"Succeeded : {succeeded}")
print(f"Failed    : {failed}")
print(f"Log       : {LOG_PATH}")

"""#corpus metrics"""

import os
import shutil

folders_to_delete = [
  "/content/corpus_modernized",
  "/content/corpus_raw",
]
```

```
for folder in folders_to_delete:
  if os.path.exists(folder):
    shutil.rmtree(folder)
    print(f"Deleted: {folder}")
  else:
    print(f"Not found: {folder}")

'''
Standalone corpus metrics
'''
import re
import math
import numpy as np
from collections import Counter
import glob
import os

WORK_DIR = "/content/corpus_modernized"

paths = sorted(glob.glob(os.path.join(WORK_DIR, "*.txt")))

texts = []
labels = []

for p in paths:
  with open(p, encoding="utf-8", errors="replace") as f:
    content = f.read().strip()
  if content:
    texts.append(content)
    labels.append(os.path.basename(p))

_WORD_RE = re.compile(r"[a-zA-Z\u00C0-\u024F'-]+")
_SENT_RE = re.compile(r'(?<=[.!?])\s+|(?<=\n)\n+')


def tokenize_words(text):
  return _WORD_RE.findall(text)


def split_into_sentences(text):
  sents = [s.strip() for s in _SENT_RE.split(text) if s.strip()]
  return sents if sents else [text.strip()]


def doc_metrics(text):
  word_tokens = tokenize_words(text)
  word_lower = [w.lower() for w in word_tokens]
  N          = len(word_tokens)

  if N == 0:
    return {
      "tokens_N"    : 0,
      "types_V"     : 0,
      "ttr"         : 0.0,
      "rttr"        : 0.0,
      "hapax"       : 0,
```

```
    "hapax_pct"  : 0.0,
    "avg_token_len": 0.0,
    "sentences"  : 0,
    "avg_sent_len": 0.0,
  }

 freq      = Counter(word_lower)
 V         = len(freq)
 ttr       = V / N
 rttr      = V / math.sqrt(N)
 hapax     = sum(1 for f in freq.values() if f == 1)
 hapax_pct = hapax / V if V else 0

 avg_token_len = sum(len(t) for t in word_tokens) / N

 sents    = split_into_sentences(text)
 n_sents             = max(len(sents), 1)  # guard against empty-sentence edge case
 avg_sent_len = N / n_sents

 return {
   "tokens_N"      : N,
   "types_V"       : V,
   "ttr"   : round(ttr, 4),
   "rttr"              : round(rttr, 4),
   "hapax"           : hapax,
   "hapax_pct"  : round(hapax_pct, 4),
   "avg_token_len": round(avg_token_len, 4),
   "sentences"  : n_sents,
   "avg_sent_len": round(avg_sent_len, 4),
 }


# ── Run over corpus
─────────────────────────────────────────────────────────────────────────────────────────────

doc_stats = [doc_metrics(t) for t in texts]

def arr(key):
  return np.array([d[key] for d in doc_stats])

# Corpus-level totals
all_words            = [w.lower() for t in texts for w in tokenize_words(t)]
total_tokens = len(all_words)
total_types = len(set(all_words))

# ── Print
─────────────────────────────────────────────────────────────────────────────────────────────
──────

print("──────────────────── METRICS SUMMARY ────────────────────")
print(f"Documents            : {len(texts)}")
print(f"Total tokens  (N) : {total_tokens:,}")
print(f"Total types           (V) : {total_types:,}")
print(f"Corpus TTR          : {total_types / total_tokens:.4f}")

print("\n--- TOKENS / TYPES ---")
print(f"Median tokens (N)  : {np.median(arr('tokens_N')):.0f}")
print(f"Median types (V)  : {np.median(arr('types_V')):.0f}")
```

```
print("\n--- LEXICAL RICHNESS ---")
print(f"Median TTR          : {np.median(arr('ttr')):.4f}")
print(f"Median RTTR         : {np.median(arr('rttr')):.4f}")

print("\n--- HAPAX ---")
print(f"Median hapax        : {np.median(arr('hapax')):.0f}")
print(f"Median hapax %      : {np.median(arr('hapax_pct')):.2%}")

print("\n--- TOKEN FORM ---")
print(f"Median token length : {np.median(arr('avg_token_len')):.2f}")

print("\n--- SENTENCES ---")
print(f"Median sentences    : {np.median(arr('sentences')):.0f}")
print(f"Median sent length  : {np.median(arr('avg_sent_len')):.2f}")

print("\n--- DISPERSION (tokens N) ---")
tok = arr('tokens_N')
print(f"Min          : {np.min(tok):.0f}")
print(f"25th percentile      : {np.percentile(tok, 25):.0f}")
print(f"Median               : {np.median(tok):.0f}")
print(f"75th percentile      : {np.percentile(tok, 75):.0f}")
print(f"Max          : {np.max(tok):.0f}")
print(f"IQR          : {np.percentile(tok, 75) - np.percentile(tok, 25):.0f}")
print(f"Std dev              : {np.std(tok):.0f}")

print("————————————————————————————————————————————————————————————")

import re
import glob
import os
import pandas as pd

RAW_DIR = "/content/corpus_raw"
MOD_DIR = "/content/corpus_modernized"

OUT_FILE = "/content/token_breakage_check.csv"
NOISE_SUMMARY_FILE = "/content/token_noise_summary.csv"
NOISE_OUTLIERS_FILE = "/content/token_noise_outliers.csv"
NOISE_PARAGRAPH_FILE = "/content/token_noise_paragraph.txt"

# Same token definition as your previous code
_WORD_RE = re.compile(r"[a-zA-Z\u00C0-\u024F'-]+")


def tokenize_words(text):
  return _WORD_RE.findall(text)


def read_text(path):
  with open(path, encoding="utf-8", errors="replace") as f:
    return f.read().strip()


def canonical_name(filename):
  """"""
```

```
  Makes raw and modernized filenames comparable.

  Example:
  Anon_Alliterative_Morte_Arthure.txt
  Anon_Alliterative_Morte_Arthure_normalized.txt

  both become:
  Anon_Alliterative_Morte_Arthure.txt
  """
  base = os.path.basename(filename)

  base = base.replace("_normalized.txt", ".txt")
  base = base.replace("_normalised.txt", ".txt")
  base = base.replace("_modernized.txt", ".txt")
  base = base.replace("_modernised.txt", ".txt")

  return base


def get_txt_files(folder):
  paths = sorted(glob.glob(os.path.join(folder, "**", "*.txt"), recursive=True))
  return {
    canonical_name(p): p
    for p in paths
  }


def classify_noise(abs_delta_pct):
  if abs_delta_pct <= 0.5:
    return "negligible"
  elif abs_delta_pct <= 1.0:
    return "minor"
  elif abs_delta_pct <= 2.0:
    return "moderate"
  elif abs_delta_pct <= 5.0:
    return "inspect"
  else:
    return "serious"


raw_files = get_txt_files(RAW_DIR)
mod_files = get_txt_files(MOD_DIR)

shared_files = sorted(set(raw_files) & set(mod_files))
raw_only = sorted(set(raw_files) - set(mod_files))
mod_only = sorted(set(mod_files) - set(raw_files))

rows = []

for filename in shared_files:
  raw_text = read_text(raw_files[filename])
  mod_text = read_text(mod_files[filename])

  raw_n = len(tokenize_words(raw_text))
  mod_n = len(tokenize_words(mod_text))
```

```
  delta = mod_n - raw_n
  delta_pct_exact = (delta / raw_n) * 100 if raw_n else 0.0

  rows.append({
    "file": filename,
    "raw_file": os.path.basename(raw_files[filename]),
    "modernized_file": os.path.basename(mod_files[filename]),
    "raw_tokens": raw_n,
    "modernized_tokens": mod_n,
    "delta": delta,
    "delta_pct": round(delta_pct_exact, 2),
    "delta_pct_exact": delta_pct_exact,
    "problem": "YES" if mod_n > raw_n else "NO"
  })


df = pd.DataFrame(rows)

print("──────────────────── FILE MATCH CHECK ────────────────────")
print(f"Raw .txt files found              : {len(raw_files)}")
print(f"Modernized .txt files found : {len(mod_files)}")
print(f"Matching files compared     : {len(shared_files)}")
print(f"Raw-only files       : {len(raw_only)}")
print(f"Modernized-only files         : {len(mod_only)}")
print("──────────────────────────────────────────────────────────")

if df.empty:
  print("\nNo matching files were found.")

  print("\nFirst raw files:")
  for f in list(raw_files)[:10]:
    print(" -", f)

  print("\nFirst modernized files:")
  for f in list(mod_files)[:10]:
    print(" -", f)

else:
  # Extra noise metrics
  df["abs_delta"] = df["delta"].abs()
  df["abs_delta_pct"] = df["delta_pct_exact"].abs()
  df["noise_severity"] = df["abs_delta_pct"].apply(classify_noise)

  df["direction"] = df["delta"].apply(
    lambda x: "longer_after_modernisation" if x > 0
    else "shorter_after_modernisation" if x < 0
    else "unchanged"
  )

  # Sort largest absolute differences first
  df = df.sort_values(
    by=["abs_delta_pct", "abs_delta"],
    ascending=[False, False]
  ).reset_index(drop=True)

  display(df)
```

```
df.to_csv(OUT_FILE, index=False)

# Corpus-level noise statistics
n_docs = len(df)

raw_total = int(df[“raw_tokens”].sum())
modern_total = int(df[“modernized_tokens”].sum())

net_delta = int(df[“delta”].sum())
net_delta_pct = (net_delta / raw_total) * 100 if raw_total else 0.0

median_delta = df[“delta”].median()
median_delta_pct = df[“delta_pct_exact”].median()

median_abs_delta = df[“abs_delta”].median()
median_abs_delta_pct = df[“abs_delta_pct”].median()

within_05 = int((df[“abs_delta_pct”] <= 0.5).sum())
within_1 = int((df[“abs_delta_pct”] <= 1.0).sum())
within_2 = int((df[“abs_delta_pct”] <= 2.0).sum())

above_2 = int((df[“abs_delta_pct”] > 2.0).sum())
above_5 = int((df[“abs_delta_pct”] > 5.0).sum())
above_10 = int((df[“abs_delta_pct”] > 10.0).sum())

longer_count = int((df[“delta”] > 0).sum())
shorter_count = int((df[“delta”] < 0).sum())
unchanged_count = int((df[“delta”] == 0).sum())

severity_counts = df[“noise_severity”].value_counts().to_dict()

# Save outliers only
outlier_df = df[df[“abs_delta_pct”] > 2.0].copy()
outlier_df.to_csv(NOISE_OUTLIERS_FILE, index=False)

# Summary table
summary_rows = [
  {“metric”: “documents_compared”, “value”: n_docs},
  {“metric”: “raw_total_tokens”, “value”: raw_total},
  {“metric”: “modernized_total_tokens”, “value”: modern_total},
  {“metric”: “net_delta_tokens”, “value”: net_delta},
  {“metric”: “net_delta_pct”, “value”: round(net_delta_pct, 4)},
  {“metric”: “median_delta_tokens”, “value”: median_delta},
  {“metric”: “median_delta_pct”, “value”: round(median_delta_pct, 4)},
  {“metric”: “median_abs_delta_tokens”, “value”: median_abs_delta},
  {“metric”: “median_abs_delta_pct”, “value”: round(median_abs_delta_pct, 4)},
  {“metric”: “within_plus_minus_0_5_pct”, “value”: within_05},
  {“metric”: “within_plus_minus_1_pct”, “value”: within_1},
  {“metric”: “within_plus_minus_2_pct”, “value”: within_2},
  {“metric”: “above_2_pct_abs_delta”, “value”: above_2},
  {“metric”: “above_5_pct_abs_delta”, “value”: above_5},
  {“metric”: “above_10_pct_abs_delta”, “value”: above_10},
  {“metric”: “longer_after_modernisation”, “value”: longer_count},
  {“metric”: “shorter_after_modernisation”, “value”: shorter_count},
  {“metric”: “unchanged”, “value”: unchanged_count},
```

```
]

for severity, count in severity_counts.items():
  summary_rows.append({
    "metric": f"noise_severity_{severity}",
    "value": count
  })

summary_df = pd.DataFrame(summary_rows)
summary_df.to_csv(NOISE_SUMMARY_FILE, index=False)

# Wording for net change
if net_delta < 0:
  net_change_phrase = (
    f"a net reduction of {abs(net_delta):,} tokens, "
    f"or {abs(net_delta_pct):.2f}% of the raw total"
  )
elif net_delta > 0:
  net_change_phrase = (
    f"a net increase of {net_delta:,} tokens, "
    f"or {net_delta_pct:.2f}% of the raw total"
  )
else:
  net_change_phrase = "no net token change"

paragraph = f"""
To assess residual textual noise after modernisation, I compared each raw file against its corresponding modernised file using token-count deltas. Across the {n_docs} texts, the raw corpus contains {raw_total:,} tokens, while the modernised corpus contains {modern_total:,} tokens, producing {net_change_phrase}. This difference is small at corpus scale and is consistent with the removal of editorial residue, paratextual material, repeated fragments, and non-lexical noise rather than evidence of systematic textual loss. At the document level, the median token delta is {median_delta:.0f} tokens ({median_delta_pct:.2f}%), and the median absolute percentage difference is {median_abs_delta_pct:.2f}%, indicating that most files remain closely aligned with their raw sources. In fact, {within_05} of {n_docs} texts fall within ±0.5%, {within_1} texts within ±1%, and {within_2} texts within ±2% of their raw-token totals. Only {above_2} texts exceed a 2% absolute difference, {above_5} texts exceed 5%, and {above_10} texts exceed 10%. The remaining high-delta cases are therefore best treated as individual outliers requiring file-level inspection, while the corpus as a whole shows a low level of residual token-count noise after cleaning.
""".strip()

  with open(NOISE_PARAGRAPH_FILE, "w", encoding="utf-8") as f:
    f.write(paragraph)

  print("\n——————————— TOKEN BREAKAGE CHECK ———————————")
  print(f"Compared files    : {n_docs}")
  print(f"Raw total tokens  : {raw_total:,}")
  print(f"Modernized total tokens : {modern_total:,}")
  print(f"Net delta         : {net_delta:,}")
  print(f"Net delta %       : {net_delta_pct:.2f}%")
  print(f"Median delta      : {median_delta:.0f}")
  print(f"Median delta %    : {median_delta_pct:.2f}%")
  print(f"Median absolute delta % : {median_abs_delta_pct:.2f}%")
  print(f"Within ±0.5%      : {within_05} / {n_docs}")
  print(f"Within ±1%        : {within_1} / {n_docs}")
  print(f"Within ±2%        : {within_2} / {n_docs}")
  print(f"Above 2%          : {above_2}")
```

```
  print(f"Above 5%          : {above_5}")
  print(f"Above 10%         : {above_10}")
  print(f"Longer after modernisation : {longer_count}")
  print(f"Shorter after modernisation : {shorter_count}")
  print(f"Unchanged         : {unchanged_count}")
  print(f"Saved full CSV              : {OUT_FILE}")
  print(f"Saved noise summary         : {NOISE_SUMMARY_FILE}")
  print(f"Saved noise outliers        : {NOISE_OUTLIERS_FILE}")
  print(f"Saved paragraph             : {NOISE_PARAGRAPH_FILE}")

""""# outcomes audit""""

import os
import re
import json
import time
import math
from typing import Any

from google import genai


# ------------------------------------------------------------
# Configuration
# ------------------------------------------------------------

ORIGINAL_DIR = "/content/raw_corpus"
MODERNISED_DIR = "/content/modernized_corpus"
OUTPUT_DIR = "/content/semantic_audit"

MODEL = "gemini-3.1-flash-lite-preview"

COOLDOWN = 10
MAX_RETRIES = 5
BACKOFF_BASE = 5

# Recommended for your corpus.
# Set to None only for small files.
CHUNK_SIZE = 5000
# CHUNK_SIZE = None

WORD_RE = re.compile(
  r"[A-Za-zÞþÐðƿȝÆæ]+(?:['‘][A-Za-zÞþÐðƿȝÆæ]+)?"
)

TRANSLATION_SCHEMA = {
  "type": "object",
  "properties": {
    "translations": {
      "type": "array",
      "items": {"type": "string"},
    }
  },
  "required": ["translations"],
  "additionalProperties": False,
}
```

```
WORD_RE = re.compile(r"[A-Za-zÞþÐðƿƷȝÆæ]+(?:['‘][A-Za-zÞþÐðƿƷȝÆæ]+)?")

from google.colab import userdata
from google import genai

GEMINI_API_KEY = userdata.get("GEMINI_API_KEY")
client = genai.Client(api_key=GEMINI_API_KEY)


# ----------------------------------------------------------------
# Text helpers
# ----------------------------------------------------------------

def get_words(text: str) -> list[str]:
  return WORD_RE.findall(text)


def read_words(path: str) -> list[str]:
  with open(path, encoding="utf-8", errors="ignore") as f:
    return get_words(f.read())


def find_modernised(filename: str) -> str | None:
  stem, ext = os.path.splitext(filename)

  candidates = [
    f"{stem}_normalized{ext}",
    f"{stem}_modernised{ext}",
    f"{stem}_modernized{ext}",
    filename,
  ]

  for candidate in candidates:
    path = os.path.join(MODERNISED_DIR, candidate)
    if os.path.exists(path):
      return path

  return None


def sanitize_tokens(tokens: list[str]) -> list[str]:
  table = str.maketrans({
    "“": "“",
    "\": "/",
  })
  return [t.translate(table) for t in tokens]


# ----------------------------------------------------------------
# JSON helpers
# ----------------------------------------------------------------

def get_response_text(response: Any) -> str | None:
  """
```

```
  response.text can sometimes be None. This tries response.text first,
  then falls back to candidate content parts if available.
  """
  text = getattr(response, "text", None)
  if text:
    return text

  candidates = getattr(response, "candidates", None) or []

  for candidate in candidates:
    content = getattr(candidate, "content", None)
    parts = getattr(content, "parts", None) or []

    part_texts = []
    for part in parts:
      part_text = getattr(part, "text", None)
      if part_text:
        part_texts.append(part_text)

    if part_texts:
      return "".join(part_texts)

  return None


def print_response_debug(response: Any) -> None:
  print("  Gemini returned no usable text.")

  prompt_feedback = getattr(response, "prompt_feedback", None)
  if prompt_feedback is not None:
    print(f"    prompt_feedback: {prompt_feedback}")

  candidates = getattr(response, "candidates", None) or []

  if not candidates:
    print("    candidates: none")
    return

  for idx, candidate in enumerate(candidates):
    finish_reason = getattr(candidate, "finish_reason", None)
    safety_ratings = getattr(candidate, "safety_ratings", None)
    content = getattr(candidate, "content", None)

    print(f"    candidate {idx} finish_reason: {finish_reason}")
    print(f"    candidate {idx} safety_ratings: {safety_ratings}")
    print(f"    candidate {idx} content: {content}")


def repair_json(raw: str | None) -> dict:
  if raw is None:
    raise ValueError("Cannot parse JSON because Gemini returned response.text=None.")

  if not isinstance(raw, str):
    raise TypeError(f"Expected JSON string, got {type(raw).__name__}.")

  raw = raw.strip()
```

```
  if not raw:
    raise ValueError("Cannot parse JSON because Gemini returned an empty string.")

  # Attempt 1: direct JSON parse.
  try:
    return json.loads(raw)
  except json.JSONDecodeError:
    pass

  # Attempt 2: remove markdown fences.
  cleaned = re.sub(r"^```(?:json)?\s*|\s*```$", "", raw, flags=re.MULTILINE).strip()

  try:
    return json.loads(cleaned)
  except json.JSONDecodeError:
    pass

  # Attempt 3: extract the outermost JSON object.
  start = cleaned.find("{")
  end = cleaned.rfind("}")

  if start != -1 and end != -1 and end > start:
    extracted = cleaned[start:end + 1]
    try:
      return json.loads(extracted)
    except json.JSONDecodeError:
      pass

  raise ValueError("Gemini response could not be repaired into valid JSON.")


def normalize_translations(data: dict) -> list[str]:
  translations = data.get("translations", [])

  if translations is None:
    return []

  if not isinstance(translations, list):
    raise ValueError(
      f"Expected 'translations' to be a list, got {type(translations).__name__}."
    )

  clean = []

  for item in translations:
    if not isinstance(item, str):
      continue

    item = item.strip()

    if item:
      clean.append(item)

  return clean
```

```
# ------------------------------------------------------------
# Prompt
# ------------------------------------------------------------

def build_prompt(filename: str, orig_chunk: list[str], mod_chunk: list[str]) -> str:
  safe_orig = sanitize_tokens(orig_chunk)
  safe_mod = sanitize_tokens(mod_chunk)

  return f"""
You are auditing a Middle English normalization.

Compare the original word token list and the modernised word token list below.

Identify every case where the modernised version uses a SEMANTIC TRANSLATION
rather than a PRESENT-DAY ENGLISH REFLEX.

Definitions:

PRESENT-DAY ENGLISH REFLEX:
The modernised token is the present-day English descendant, spelling equivalent,
or regularised form of the original token. The lexical identity is basically the
same, even if spelling, morphology, dialectal form, or inflection has changed.
These are NOT semantic translations. Exclude them.

SEMANTIC TRANSLATION:
The modernised token is not simply the present-day English reflex of the original.
A different lexical item has been chosen to express the meaning of an obsolete,
opaque, dialectal, or context-dependent Middle English word.
These ARE semantic translations. Include them.

Rules:
- Return only semantic translations.
- Do not return present-day English reflexes.
- Return every token occurrence, not just unique word-pair types.
- Do not collapse repeated pairs.
- Only include pairs that actually occur in the supplied token lists.
- If the distinction is uncertain, exclude the pair.
- Do not add explanations, notes, markdown, or commentary outside the JSON.

Required JSON shape:
{{"translations": ["original -> modernised"]}}

File: {filename}

Original tokens:
{json.dumps(safe_orig, ensure_ascii=False)}

Modernised tokens:
{json.dumps(safe_mod, ensure_ascii=False)}
""".strip()


# ------------------------------------------------------------
# Gemini API call
# ------------------------------------------------------------
```

```
def call_gemini(
  filename: str,
  orig_chunk: list[str],
  mod_chunk: list[str],
) -> list[str] | None:
  prompt = build_prompt(filename, orig_chunk, mod_chunk)

  for attempt in range(MAX_RETRIES):
    try:
      response = client.models.generate_content(
        model=MODEL,
        contents=prompt,
        config={
          "temperature": 0.0,
          "response_mime_type": "application/json",
          "response_json_schema": TRANSLATION_SCHEMA,
        },
      )

      raw = get_response_text(response)

      if raw is None:
        print_response_debug(response)
        raise ValueError("Gemini response text was None.")

      data = repair_json(raw)
      return normalize_translations(data)

    except Exception as exc:
      wait = BACKOFF_BASE * (2 ** attempt)
      print(
        f"   Attempt {attempt + 1}/{MAX_RETRIES} failed: {exc}. "
        f"Retrying in {wait}s..."
      )
      time.sleep(wait)

  print(f"All {MAX_RETRIES} attempts failed. Skipping chunk.")
  return None


# ---------------------------------------------------------------
# Output helpers
# ---------------------------------------------------------------

def safe_stem(filename: str) -> str:
  stem, _ = os.path.splitext(filename)
  return stem


def partial_path(filename: str) -> str:
  return os.path.join(OUTPUT_DIR, f"{safe_stem(filename)}_semantic.partial")


def final_path(filename: str) -> str:
  return os.path.join(OUTPUT_DIR, f"{safe_stem(filename)}_semantic.txt")
```

```
def output_exists(filename: str) -> bool:
  return os.path.exists(final_path(filename))


def load_partial(filename: str) -> tuple[int, list[str]]:
  """
  Partial file format:

    CHUNKS_DONE: N
    translation line
    translation line
    ...

  If you change CHUNK_SIZE while resuming a file, delete the partial first.
  """
  path = partial_path(filename)

  if not os.path.exists(path):
    return 0, []

  with open(path, encoding="utf-8") as f:
    lines = f.read().splitlines()

  if not lines or not lines[0].startswith("CHUNKS_DONE:"):
    return 0, []

  chunks_done = int(lines[0].split(":", 1)[1].strip())
  translations = [line for line in lines[1:] if line.strip()]

  return chunks_done, translations


def write_partial(filename: str, chunks_done: int, translations: list[str]) -> None:
  path = partial_path(filename)

  with open(path, "w", encoding="utf-8") as f:
    f.write(f"CHUNKS_DONE: {chunks_done}\n")

    for translation in translations:
      f.write(f"{translation}\n")


def append_chunk_to_partial(
  filename: str,
  chunk_index: int,
  translations: list[str],
) -> list[str]:
  chunks_done, existing = load_partial(filename)

  expected_chunk_index = chunks_done

  if chunk_index != expected_chunk_index:
    raise RuntimeError(
      f"Partial-file mismatch for {filename}: "
```

```
      f"expected chunk index {expected_chunk_index}, got {chunk_index}."
    )

  all_translations = existing + translations
  write_partial(filename, chunk_index + 1, all_translations)

  return all_translations


def finalize(filename: str, translations: list[str]) -> None:
  out = final_path(filename)

  with open(out, "w", encoding="utf-8") as f:
    f.write(f"File: {filename}\n")
    f.write(f"Count: {len(translations)}\n\n")

    for translation in translations:
      f.write(f"{translation}\n")

  partial = partial_path(filename)

  if os.path.exists(partial):
    os.remove(partial)

  print(f" Semantic translations: {len(translations)}")
  print(f" Saved: {out}")


# ------------------------------------------------------------
# Chunk helpers
# ------------------------------------------------------------

def using_chunking() -> bool:
  return CHUNK_SIZE is not None and CHUNK_SIZE > 0


def count_chunks(original_words: list[str], modernised_words: list[str]) -> int:
  if not using_chunking():
    return 1

  original_chunks = max(1, math.ceil(len(original_words) / CHUNK_SIZE))
  modernised_chunks = max(1, math.ceil(len(modernised_words) / CHUNK_SIZE))

  return max(original_chunks, modernised_chunks)


def get_chunk(
  words: list[str],
  chunk_index: int,
) -> list[str]:
  if not using_chunking():
    return words

  start = chunk_index * CHUNK_SIZE
  end = start + CHUNK_SIZE
```

```
  return words[start:end]


# ---------------------------------------------------------------
# File processing
# ---------------------------------------------------------------

def process_file(
  filename: str,
  original_words: list[str],
  modernised_words: list[str],
) -> list[str] | None:
  total_chunks = count_chunks(original_words, modernised_words)
  chunks_done, all_translations = load_partial(filename)

  if chunks_done > 0:
    print(
      f" Resuming from chunk {chunks_done + 1}/{total_chunks} "
      f"({len(all_translations)} translations recovered)."
    )

  if chunks_done >= total_chunks:
    return all_translations

  for chunk_index in range(chunks_done, total_chunks):
    orig_chunk = get_chunk(original_words, chunk_index)
    mod_chunk = get_chunk(modernised_words, chunk_index)

    print(
      f" Chunk {chunk_index + 1}/{total_chunks} "
      f"({len(orig_chunk)} original / {len(mod_chunk)} modernised tokens)..."
    )

    translations = call_gemini(filename, orig_chunk, mod_chunk)

    if translations is None:
      print(
        f" Chunk {chunk_index + 1} failed permanently. "
        f"Progress saved; rerun to resume here."
      )
      return None

    all_translations = append_chunk_to_partial(
      filename=filename,
      chunk_index=chunk_index,
      translations=translations,
    )

    if using_chunking() and chunk_index + 1 < total_chunks:
      time.sleep(COOLDOWN)

  return all_translations


# ---------------------------------------------------------------
# File discovery
```

```
# ------------------------------------------------------------

def collect_files() -> list[tuple[str, str, list[str], list[str], int]]:
  """
  Returns:

    (filename, modernised_path, original_words, modernised_words, total_tokens)

  Files are sorted ascending by total token count so short files finish first.
  """
  entries = []

  for filename in os.listdir(ORIGINAL_DIR):
    if not filename.endswith(".txt"):
      continue

    if output_exists(filename):
      print(f"Skipping {filename}: already done")
      continue

    modernised_path = find_modernised(filename)

    if not modernised_path:
      print(f"Skipping {filename}: no modernised file found")
      continue

    original_path = os.path.join(ORIGINAL_DIR, filename)

    original_words = read_words(original_path)
    modernised_words = read_words(modernised_path)

    total_tokens = len(original_words) + len(modernised_words)

    entries.append(
      (
        filename,
        modernised_path,
        original_words,
        modernised_words,
        total_tokens,
      )
    )

  entries.sort(key=lambda entry: entry[4])
  return entries


# ------------------------------------------------------------
# Main
# ------------------------------------------------------------

def main() -> None:
  entries = collect_files()

  mode = "no chunking" if not using_chunking() else f"chunks of {CHUNK_SIZE} tokens"
```

```
  print(f"\n{len(entries)} file(s) to process, ordered by token count.")
  print(f"Mode: {mode}\n")

  for filename, modernised_path, original_words, modernised_words, total in entries:
    has_partial = os.path.exists(partial_path(filename))

    print(
      f"{'Resuming' if has_partial else 'Processing'} "
      f"{filename} ({total} total tokens)..."
    )
    print(f"  Original tokens:  {len(original_words)}")
    print(f"  Modernised tokens: {len(modernised_words)}")
    print(f"  Modernised file:  {modernised_path}")

    translations = process_file(
      filename=filename,
      original_words=original_words,
      modernised_words=modernised_words,
    )

    if translations is None:
      print(f"  {filename} paused — rerun to continue from saved progress.")
    else:
      finalize(filename, translations)

    time.sleep(COOLDOWN)

  print("\nAll done.")


if __name__ == "__main__":
  main()

import re
import glob
import os
import csv
from pathlib import Path
import pandas as pd

# ──────────────────────────────────────────────────────────────
# SETTINGS
# ──────────────────────────────────────────────────────────────

SEMANTIC_DIR = "/content/semantic_audit"   # folder with *_semantic.txt files
MOD_DIR = "/content/corpus_modernized"    # folder with modernized texts

OUT_CSV = "/content/semantic_translation_percentages.csv"

_WORD_RE = re.compile(r"[a-zA-Z\u00C0-\u024F'-]+")

NULL_TARGETS = {"", "null", "(null)", "none", "nan"}


# ──────────────────────────────────────────────────────────────
# HELPERS
```

```
# ──────────────────────────────────────────────────────────────

def tokenize_words(text):
  return _WORD_RE.findall(text)


def read_text(path):
  with open(path, encoding="utf-8", errors="replace") as f:
    return f.read()


def canonical_name(filename):
  """
  Converts filenames to a shared base form.

  Example:
  D'arras_Melusine_semantic.txt
  D'arras_Melusine_normalized.txt

  both become:
  D'arras_Melusine.txt
  """
  base = os.path.basename(str(filename))

  base = base.replace("_semantic.txt", ".txt")
  base = base.replace("_normalized.txt", ".txt")
  base = base.replace("_normalised.txt", ".txt")
  base = base.replace("_modernized.txt", ".txt")
  base = base.replace("_modernised.txt", ".txt")

  return base


def find_modernized_file(base_file):
  stem = base_file.replace(".txt", "")

  candidates = [
    f"{stem}_normalized.txt",
    f"{stem}_normalised.txt",
    f"{stem}_modernized.txt",
    f"{stem}_modernised.txt",
    base_file,
  ]

  for candidate in candidates:
    matches = glob.glob(os.path.join(MOD_DIR, "**", candidate), recursive=True)
    if matches:
      return matches[0]

  return None


def clean_field(x):
  return str(x).strip().strip("'").strip('"').lower()
```

```
def count_semantic_instances(path):
  """
  Counts semantic translation instances.

  Supports:
  1. repeated pair lines:
    word -> translation

  2. CSV-style rows:
    source_word,normalized_word,frequency,example_line

  If a frequency column exists, it uses that frequency.
  Otherwise, each row counts as 1.
  """
  count = 0

  text = read_text(path)

  if text.strip() == "NO_SEMANTIC_PAIRS":
    return 0

  for raw_line in text.splitlines():
    line = raw_line.strip()

    if not line:
      continue

    if line.startswith("```") or line == "NO_SEMANTIC_PAIRS":
      continue

    source = None
    target = None
    freq = 1

    if "->" in line:
      source, target = line.split("->", 1)

    else:
      try:
        row = next(csv.reader([line]))
      except Exception:
        continue

      if len(row) < 2:
        continue

      if row[0].strip().lower() in {"source_word", "source"}:
        continue

      source = row[0]
      target = row[1]

      if len(row) >= 3:
        try:
          freq = int(float(row[2]))
        except Exception:
```

```
        freq = 1

    source = clean_field(source)
    target = clean_field(target)

    if not source or target in NULL_TARGETS:
      continue

    count += freq

  return count


# ─────────────────────────────────────────────────────────────
# MAIN
# ─────────────────────────────────────────────────────────────

semantic_files = sorted(
  glob.glob(os.path.join(SEMANTIC_DIR, "**", "*_semantic.txt"), recursive=True)
)

rows = []

for semantic_path in semantic_files:
  semantic_file = os.path.basename(semantic_path)
  base_file = canonical_name(semantic_file)

  mod_path = find_modernized_file(base_file)

  if mod_path is None:
    rows.append({
      "file": base_file,
      "semantic_file": semantic_file,
      "modernized_file": None,
      "semantic_translation_count": count_semantic_instances(semantic_path),
      "modernized_tokens": None,
      "semantic_translation_pct": None,
      "status": "NO_MATCHING_MODERNIZED_FILE"
    })
    continue

  semantic_count = count_semantic_instances(semantic_path)
  modernized_tokens = len(tokenize_words(read_text(mod_path)))

  semantic_pct = (
    (semantic_count / modernized_tokens) * 100
    if modernized_tokens > 0 else 0
  )

  rows.append({
    "file": base_file,
    "semantic_file": semantic_file,
    "modernized_file": os.path.basename(mod_path),
    "semantic_translation_count": semantic_count,
    "modernized_tokens": modernized_tokens,
    "semantic_translation_pct": round(semantic_pct, 4),
```

```
    "status": "OK"
  })

df = pd.DataFrame(rows)

df = df.sort_values(
  by="semantic_translation_pct",
  ascending=False,
  na_position="last"
).reset_index(drop=True)

df.to_csv(OUT_CSV, index=False)

display(df)

valid = df[df["status"] == "OK"].copy()

total_semantic = int(valid["semantic_translation_count"].sum())
total_tokens = int(valid["modernized_tokens"].sum())

weighted_pct = (total_semantic / total_tokens) * 100 if total_tokens else 0
median_pct = valid["semantic_translation_pct"].median()
median_count = valid["semantic_translation_count"].median()

print("\n───────────────── SEMANTIC TRANSLATION PERCENTAGE ─────────────────")
print(f"Audited texts                    : {len(valid)}")
print(f"Total modernized tokens      : {total_tokens:,}")
print(f"Total semantic translation instances: {total_semantic:,}")
print(f"Weighted corpus percentage             : {weighted_pct:.4f}%")
print(f"Median semantic translation count  : {median_count:.0f}")
print(f"Median semantic translation %          : {median_pct:.4f}%")
print(f"Saved to                 : {OUT_CSV}")
print("──────────────────────────────────────────────────────────────────────────")

"""# word level embeddings"""

#imports
!pip install -q fasttext umap-learn

import os
import re
import json
import unicodedata
import numpy as np
import pandas as pd
import fasttext
import umap.umap_ as umap
import matplotlib.pyplot as plt
from adjustText import adjust_text
from collections import Counter
from matplotlib.lines import Line2D

# -----------------------------
# Configuration
# -----------------------------
```

```
TEXT_DIR = “/content/corpus_modernized”
TRAINING_FILE = “/content/fasttext_training.txt”
OUTPUT_DIR = “/content/fasttext_outputs”
os.makedirs(OUTPUT_DIR, exist_ok=True)

DIM = 256
WINDOW = 8
EPOCHS = 20
MIN_N = 3
MAX_N = 6
NEG_SAMPLES = 10

TRAIN_MIN_COUNT = 3
NEIGHBOUR_MIN_FREQ = 5

LOVE_FORMS = [
  “love”, “loved”, “loveth”, “loven”, “loving”,
  “lover”, “lovers”, “beloved”,
]

TOP_N_NEIGHBOURS = 100

# ------------------------------
# Basic checks
# ------------------------------
if not os.path.isdir(TEXT_DIR):
  raise FileNotFoundError(f”TEXT_DIR does not exist: {TEXT_DIR}”)

files = sorted([f for f in os.listdir(TEXT_DIR) if f.endswith(“.txt”)])
if not files:
  raise ValueError(f”No .txt files found in {TEXT_DIR}”)

print(f”Number of text files: {len(files)}”)
print(f”First files: {files[:5]}”)

# ------------------------------
# Tokenisation
# ------------------------------
def clean_text(text):
  text = unicodedata.normalize(“NFKC”, text)
  text = text.casefold()
  text = re.sub(r”\d+”, “ “, text)
  text = re.sub(r”[^\w\s’’\-]”, “ “, text, flags=re.UNICODE)
  text = text.replace(“_”, “ “)
  return text


def tokenize(text):
  text = clean_text(text)
  tokens = []
  for tok in text.split():
    tok = tok.strip(“‘‘-/\u2014\u2013”)
    if tok and not any(ch.isdigit() for ch in tok):
      tokens.append(tok)
  return tokens
```

```
# ----------------------------
# Build training file
# ----------------------------
with open(TRAINING_FILE, “w”, encoding=“utf-8”) as out:
  for fname in files:
    path = os.path.join(TEXT_DIR, fname)
    with open(path, “r”, encoding=“utf-8”, errors=“ignore”) as f:
      raw = f.read()
    units = re.split(r”(?<=[.!?;:])\s+”, raw)
    for unit in units:
      tokens = tokenize(unit)
      if tokens:
        out.write(“ “.join(tokens) + “\n”)

print(f”Training file written: {TRAINING_FILE}”)

# ----------------------------
# Corpus-level frequencies
# ----------------------------
with open(TRAINING_FILE, “r”, encoding=“utf-8”) as f:
  all_tokens = f.read().split()

corpus_freq = Counter(all_tokens)
total_tokens = sum(corpus_freq.values())
vocab_size = len(corpus_freq)

print(f”Total tokens: {total_tokens}”)
print(f”Vocabulary size: {vocab_size}”)

print(f”\nFrequencies of love-family forms:”)
for form in LOVE_FORMS:
  print(f” {form}: {corpus_freq.get(form, 0)}”)

viable_love_forms = [f for f in LOVE_FORMS if corpus_freq.get(f, 0) >= TRAIN_MIN_COUNT]
if not viable_love_forms:
  raise ValueError(
    f”None of the love-family forms meet the TRAIN_MIN_COUNT threshold “
    f”of {TRAIN_MIN_COUNT}. Lower the threshold or revise LOVE_FORMS.”
  )
print(f”\nViable love-family forms (freq >= {TRAIN_MIN_COUNT}): {viable_love_forms}”)

# ----------------------------
# Helpers
# ----------------------------
def safe_unit_vector(vec):
  norm = np.linalg.norm(vec)
  if norm == 0:
    return vec
  return vec / norm


def get_love_centroid(model, love_forms, corpus_freq, min_count):
  vectors, used_forms = [], []
  for form in love_forms:
    if corpus_freq.get(form, 0) < min_count:
      continue
```

```
    vectors.append(safe_unit_vector(model.get_word_vector(form)))
    used_forms.append(form)
  if not vectors:
    raise ValueError("No love-family forms available for centroid.")
  centroid = np.mean(np.vstack(vectors), axis=0)
  return safe_unit_vector(centroid), used_forms


def get_centroid_neighbourhood(model, centroid, exclude_forms, corpus_freq, min_freq):
  exclude_set = set(exclude_forms)
  rows = []
  for word in model.get_words():
    if word in exclude_set:
      continue
    freq = corpus_freq.get(word, 0)
    if freq < min_freq:
      continue
    vec = safe_unit_vector(model.get_word_vector(word))
    sim = float(np.dot(centroid, vec))
    rows.append((word, sim, freq))
  df = pd.DataFrame(rows, columns=["word", "cosine_similarity", "frequency"])
  df = df.sort_values("cosine_similarity", ascending=False).reset_index(drop=True)
  return df

# ------------------------------
# Train headline model (window = 8)
# ------------------------------
print(f"\n--- Training headline model (window={WINDOW}) ---")

headline_model = fasttext.train_unsupervised(
  input=TRAINING_FILE,
  model="skipgram",
  dim=DIM,
  ws=WINDOW,
  epoch=EPOCHS,
  minn=MIN_N,
  maxn=MAX_N,
  neg=NEG_SAMPLES,
  minCount=TRAIN_MIN_COUNT,
  thread=4,
)
print(f"Vocabulary size: {len(headline_model.get_words())}")

headline_model.save_model(f"{OUTPUT_DIR}/fasttext_model_headline.bin")

params = {
  "text_dir": TEXT_DIR,
  "training_file": TRAINING_FILE,
  "love_forms": LOVE_FORMS,
  "viable_love_forms": viable_love_forms,
  "dim": DIM,
  "window": WINDOW,
  "epochs": EPOCHS,
  "min_n": MIN_N,
  "max_n": MAX_N,
  "negative_samples": NEG_SAMPLES,
```

```
  "train_min_count": TRAIN_MIN_COUNT,
  "neighbour_min_freq": NEIGHBOUR_MIN_FREQ,
  "total_tokens": total_tokens,
  "vocabulary_size": vocab_size,
  "number_of_files": len(files),
}
with open(f"{OUTPUT_DIR}/run_parameters.json", "w", encoding="utf-8") as f:
  json.dump(params, f, indent=2)

# ------------------------------
# Headline analysis: love-family centroid
# ------------------------------
headline_centroid, used_forms = get_love_centroid(
  headline_model, LOVE_FORMS, corpus_freq, TRAIN_MIN_COUNT
)
print(f"\nCentroid built from forms: {used_forms}")

headline_neighbours = get_centroid_neighbourhood(
  headline_model, headline_centroid, used_forms,
  corpus_freq, NEIGHBOUR_MIN_FREQ,
)
headline_neighbours.to_csv(
  f"{OUTPUT_DIR}/love_centroid_neighbours_headline.csv", index=False
)

print(f"\nTop {TOP_N_NEIGHBOURS} neighbours of love-family centroid (window={WINDOW}):")
print(headline_neighbours.head(TOP_N_NEIGHBOURS).to_string())

# ------------------------------
# Visualisation
# ------------------------------
print(f"\n--- Generating figure for top-{TOP_N_NEIGHBOURS} ---")

top_words = headline_neighbours["word"].head(TOP_N_NEIGHBOURS).tolist()

neighbour_vectors = np.vstack([
  safe_unit_vector(headline_model.get_word_vector(w)) for w in top_words
])
all_vectors = np.vstack([headline_centroid, neighbour_vectors])

reducer = umap.UMAP(
  n_neighbors=15, min_dist=0.1,
  metric="cosine", random_state=42,
)
coords = reducer.fit_transform(all_vectors)

plt.rcParams.update({
  "font.family": "serif",
  "font.serif": ["Times New Roman", "DejaVu Serif", "serif"],
  "mathtext.fontset": "stix",
})

fig, ax = plt.subplots(figsize=(14, 10))
ax.set_facecolor("white")

ax.scatter(
```

```
  coords[1:, 0], coords[1:, 1],
  s=14, facecolors="white", edgecolors="black",
  linewidths=0.6, zorder=2,
)
ax.scatter(
  coords[0, 0], coords[0, 1],
  s=60, c="black", edgecolors="black",
  linewidths=0.6, zorder=5,
)

texts = [
  ax.text(
    coords[0, 0], coords[0, 1], "love",
    fontsize=10, fontweight="normal", fontstyle="normal",
    color="black", zorder=6,
  )
]
for i, word in enumerate(top_words):
  texts.append(
    ax.text(
      coords[i + 1, 0], coords[i + 1, 1], word,
      fontsize=8, fontstyle="italic",
      color="black", zorder=3,
    )
  )

adjust_text(
  texts, ax=ax,
  expand_points=(1.4, 1.4),
  expand_text=(1.2, 1.2),
  force_text=(0.4, 0.6),
  force_points=(0.3, 0.4),
  arrowprops=dict(arrowstyle="-", color="black", lw=0.3, alpha=0.5),
)

ax.set_xlabel("UMAP 1", fontsize=9)
ax.set_ylabel("UMAP 2", fontsize=9)

for spine in ax.spines.values():
  spine.set_color("black")
  spine.set_linewidth(0.5)
ax.tick_params(axis="both", which="both", length=3, color="black",
       labelsize=8, labelcolor="black")

legend_handles = [
  Line2D([0], [0], marker="o", color="w", markerfacecolor="black",
      markeredgecolor="black", markersize=8, linewidth=0,
      label="Love-family centroid"),
  Line2D([0], [0], marker="o", color="w", markerfacecolor="white",
      markeredgecolor="black", markersize=6, linewidth=0,
      label="Nearest neighbour"),
]
ax.legend(
  handles=legend_handles,
  loc="lower right",
  fontsize=8,
```

```
  frameon=True,
  edgecolor="black",
  fancybox=False,
)

plt.tight_layout()
output_path = f"{OUTPUT_DIR}/love_centroid_neighbourhood_top{TOP_N_NEIGHBOURS}.png"
plt.savefig(output_path, dpi=300, bbox_inches="tight")
plt.show()
print(f"Saved: {output_path}")

# ------------------------------
# Summary
# ------------------------------
print(f"\n--- Summary ---")
print(f"Headline configuration: dim={DIM}, window={WINDOW}, epochs={EPOCHS}")
print(f"Love-family centroid built from: {used_forms}")
print(f"Top-{TOP_N_NEIGHBOURS} neighbours retrieved")
print(f"\nOutputs saved to {OUTPUT_DIR}/")
print(f" - fasttext_model_headline.bin")
print(f" - run_parameters.json")
print(f" - love_centroid_neighbours_headline.csv")
print(f" - love_centroid_neighbourhood_top{TOP_N_NEIGHBOURS}.png")

import os
import re
import json
import unicodedata
import numpy as np
import pandas as pd
import fasttext
import umap.umap_ as umap
import matplotlib.pyplot as plt
from adjustText import adjust_text
from collections import Counter

# ------------------------------
# Configuration
# ------------------------------
TEXT_DIR = "/content/corpus_modernized"
TRAINING_FILE = "/content/fasttext_training.txt"
OUTPUT_DIR = "/content/fasttext_outputs"
os.makedirs(OUTPUT_DIR, exist_ok=True)

DIM = 256
WINDOW = 8
EPOCHS = 20
MIN_N = 3
MAX_N = 6
NEG_SAMPLES = 10

TRAIN_MIN_COUNT = 3
NEIGHBOUR_MIN_FREQ = 5

LOVE_FORMS = [
  "love", "loved", "loveth", "loven", "loving",
```

```
  “lover”, “lovers”, “beloved”,
]

TOP_N_NEIGHBOURS = 100

# ------------------------------
# Basic checks
# ------------------------------
if not os.path.isdir(TEXT_DIR):
  raise FileNotFoundError(f”TEXT_DIR does not exist: {TEXT_DIR}”)

files = sorted([f for f in os.listdir(TEXT_DIR) if f.endswith(“.txt”)])
if not files:
  raise ValueError(f”No .txt files found in {TEXT_DIR}”)

print(f”Number of text files: {len(files)}”)
print(f”First files: {files[:5]}”)

# ------------------------------
# Tokenisation
# ------------------------------
def clean_text(text):
  text = unicodedata.normalize(“NFKC”, text)
  text = text.casefold()
  text = re.sub(r”\d+”, “ “, text)
  text = re.sub(r”[^\w\s’’\-]”, “ “, text, flags=re.UNICODE)
  text = text.replace(“_”, “ “)
  return text


def tokenize(text):
  text = clean_text(text)
  tokens = []
  for tok in text.split():
    tok = tok.strip(“‘‘-/\u2014\u2013”)
    if tok and not any(ch.isdigit() for ch in tok):
      tokens.append(tok)
  return tokens

# ------------------------------
# Build training file
# ------------------------------
with open(TRAINING_FILE, “w”, encoding=“utf-8”) as out:
  for fname in files:
    path = os.path.join(TEXT_DIR, fname)
    with open(path, “r”, encoding=“utf-8”, errors=“ignore”) as f:
      raw = f.read()
    units = re.split(r”(?<=[.!?;:])\s+”, raw)
    for unit in units:
      tokens = tokenize(unit)
      if tokens:
        out.write(“ “.join(tokens) + “\n”)

print(f”Training file written: {TRAINING_FILE}”)

# ------------------------------
```

```
# Corpus-level frequencies
# ------------------------------
with open(TRAINING_FILE, “r”, encoding=“utf-8”) as f:
  all_tokens = f.read().split()

corpus_freq = Counter(all_tokens)
total_tokens = sum(corpus_freq.values())
vocab_size = len(corpus_freq)

print(f”Total tokens: {total_tokens}”)
print(f”Vocabulary size: {vocab_size}”)

print(f”\nFrequencies of love-family forms:”)
for form in LOVE_FORMS:
  print(f” {form}: {corpus_freq.get(form, 0)}”)

viable_love_forms = [f for f in LOVE_FORMS if corpus_freq.get(f, 0) >= TRAIN_MIN_COUNT]
if not viable_love_forms:
  raise ValueError(
    f”None of the love-family forms meet the TRAIN_MIN_COUNT threshold “
    f”of {TRAIN_MIN_COUNT}. Lower the threshold or revise LOVE_FORMS.”
  )
print(f”\nViable love-family forms (freq >= {TRAIN_MIN_COUNT}): {viable_love_forms}”)

# ------------------------------
# Helpers
# ------------------------------
def safe_unit_vector(vec):
  norm = np.linalg.norm(vec)
  if norm == 0:
    return vec
  return vec / norm


def get_love_centroid(model, love_forms, corpus_freq, min_count):
  vectors, used_forms = [], []
  for form in love_forms:
    if corpus_freq.get(form, 0) < min_count:
      continue
    vectors.append(safe_unit_vector(model.get_word_vector(form)))
    used_forms.append(form)
  if not vectors:
    raise ValueError(“No love-family forms available for centroid.”)
  centroid = np.mean(np.vstack(vectors), axis=0)
  return safe_unit_vector(centroid), used_forms


def get_centroid_neighbourhood(model, centroid, exclude_forms, corpus_freq, min_freq):
  exclude_set = set(exclude_forms)
  rows = []
  for word in model.get_words():
    if word in exclude_set:
      continue
    freq = corpus_freq.get(word, 0)
    if freq < min_freq:
      continue
```

```
    vec = safe_unit_vector(model.get_word_vector(word))
    sim = float(np.dot(centroid, vec))
    rows.append((word, sim, freq))
  df = pd.DataFrame(rows, columns=[“word”, “cosine_similarity”, “frequency”])
  df = df.sort_values(“cosine_similarity”, ascending=False).reset_index(drop=True)
  return df

# ------------------------------
# Train headline model (window = 8)
# ------------------------------
print(f”\n--- Training headline model (window={WINDOW}) ---”)

headline_model = fasttext.train_unsupervised(
  input=TRAINING_FILE,
  model=“skipgram”,
  dim=DIM,
  ws=WINDOW,
  epoch=EPOCHS,
  minn=MIN_N,
  maxn=MAX_N,
  neg=NEG_SAMPLES,
  minCount=TRAIN_MIN_COUNT,
  thread=4,
)
print(f”Vocabulary size: {len(headline_model.get_words())}”)

headline_model.save_model(f”{OUTPUT_DIR}/fasttext_model_headline.bin”)

params = {
  “text_dir”: TEXT_DIR,
  “training_file”: TRAINING_FILE,
  “love_forms”: LOVE_FORMS,
  “viable_love_forms”: viable_love_forms,
  “dim”: DIM,
  “window”: WINDOW,
  “epochs”: EPOCHS,
  “min_n”: MIN_N,
  “max_n”: MAX_N,
  “negative_samples”: NEG_SAMPLES,
  “train_min_count”: TRAIN_MIN_COUNT,
  “neighbour_min_freq”: NEIGHBOUR_MIN_FREQ,
  “total_tokens”: total_tokens,
  “vocabulary_size”: vocab_size,
  “number_of_files”: len(files),
}
with open(f”{OUTPUT_DIR}/run_parameters.json”, “w”, encoding=“utf-8”) as f:
  json.dump(params, f, indent=2)

# ------------------------------
# Headline analysis: love-family centroid
# ------------------------------
headline_centroid, used_forms = get_love_centroid(
  headline_model, LOVE_FORMS, corpus_freq, TRAIN_MIN_COUNT
)
print(f”\nCentroid built from forms: {used_forms}”)
```

```
headline_neighbours = get_centroid_neighbourhood(
  headline_model, headline_centroid, used_forms,
  corpus_freq, NEIGHBOUR_MIN_FREQ,
)
headline_neighbours.to_csv(
  f"{OUTPUT_DIR}/love_centroid_neighbours_headline.csv", index=False
)

print(f"\nTop {TOP_N_NEIGHBOURS} neighbours of love-family centroid (window={WINDOW}):")
print(headline_neighbours.head(TOP_N_NEIGHBOURS).to_string())

# ----------------------------
# Visualisation
# ----------------------------
print(f"\n--- Generating figure for top-{TOP_N_NEIGHBOURS} ---")

top_words = headline_neighbours["word"].head(TOP_N_NEIGHBOURS).tolist()

neighbour_vectors = np.vstack([
  safe_unit_vector(headline_model.get_word_vector(w)) for w in top_words
])
all_vectors = np.vstack([headline_centroid, neighbour_vectors])

reducer = umap.UMAP(
  n_neighbors=15, min_dist=0.1,
  metric="cosine", random_state=42,
)
coords = reducer.fit_transform(all_vectors)

plt.rcParams.update({
  "font.family": "serif",
  "font.serif": ["Times New Roman", "DejaVu Serif", "serif"],
  "mathtext.fontset": "stix",
})

fig, ax = plt.subplots(figsize=(14, 10))
ax.set_facecolor("white")

ax.scatter(
  coords[1:, 0], coords[1:, 1],
  s=14, facecolors="white", edgecolors="black",
  linewidths=0.6, zorder=2,
)
ax.scatter(
  coords[0, 0], coords[0, 1],
  s=60, c="black", edgecolors="black",
  linewidths=0.6, zorder=5,
)

texts = [
  ax.text(
    coords[0, 0], coords[0, 1], "love",
    fontsize=10, fontweight="normal", fontstyle="normal",
    color="black", zorder=6,
  )
]
```

```
for i, word in enumerate(top_words):
  texts.append(
    ax.text(
      coords[i + 1, 0], coords[i + 1, 1], word,
      fontsize=8, fontstyle="italic",
      color="black", zorder=3,
    )
  )

adjust_text(
  texts, ax=ax,
  expand_points=(1.4, 1.4),
  expand_text=(1.2, 1.2),
  force_text=(0.4, 0.6),
  force_points=(0.3, 0.4),
  arrowprops=dict(arrowstyle="-", color="black", lw=0.3, alpha=0.5),
)

ax.set_xlabel("UMAP 1", fontsize=9)
ax.set_ylabel("UMAP 2", fontsize=9)

for spine in ax.spines.values():
  spine.set_color("black")
  spine.set_linewidth(0.5)
ax.tick_params(axis="both", which="both", length=3, color="black",
       labelsize=8, labelcolor="black")

plt.tight_layout()
output_path = f"{OUTPUT_DIR}/love_centroid_neighbourhood_top{TOP_N_NEIGHBOURS}.png"
plt.savefig(output_path, dpi=300, bbox_inches="tight")
plt.show()
print(f"Saved: {output_path}")

# ------------------------------
# Summary
# ------------------------------
print(f"\n--- Summary ---")
print(f"Headline configuration: dim={DIM}, window={WINDOW}, epochs={EPOCHS}")
print(f"Love-family centroid built from: {used_forms}")
print(f"Top-{TOP_N_NEIGHBOURS} neighbours retrieved")
print(f"\nOutputs saved to {OUTPUT_DIR}/")
print(f" - fasttext_model_headline.bin")
print(f" - run_parameters.json")
print(f" - love_centroid_neighbours_headline.csv")
print(f" - love_centroid_neighbourhood_top{TOP_N_NEIGHBOURS}.png")

!pip install -q fasttext adjustText

import os
import re
import numpy as np
import pandas as pd
import fasttext
import umap.umap_ as umap
import matplotlib.pyplot as plt
from adjustText import adjust_text
```

```
from collections import Counter
from matplotlib.lines import Line2D

# ------------------------------
# Configuration
# ------------------------------
OUTPUT_DIR = "/content/fasttext_outputs"
MODEL_PATH = f"{OUTPUT_DIR}/fasttext_model_headline.bin"
TRAINING_FILE = f"{OUTPUT_DIR}/fasttext_training.txt"
OP_A_DIR = f"{OUTPUT_DIR}/operation_A_additive"

NEIGHBOUR_MIN_FREQ = 5
N_FUNCTION_WORDS = 100
TOP_N = 50
TOP_N_PER_PANEL = 25

LOVE_FORMS = [
  "love", "loved", "loveth", "loven", "loving",
  "lover", "lovers", "beloved",
]

EXPLICIT_LOVE_VARIANTS = {
  "love", "loves", "loved", "loveth", "loven", "loving", "lover",
  "lovers", "beloved", "lovely", "loveless", "lovelier", "lovesome",
  "lovedrunk", "lovelonging", "loveday", "lovingly", "love's",
  "love\u2019s", "lovede", "lovest", "lovings", "lous", "lyuen",
  "lovell", "lovedays", "lovesome",
}

ANCHORS = [
  "god", "lord", "lady",
  "charity", "war", "death",
  "pain", "soul", "sin",
  "desire", "heart",
  "knight", "service",
]

# ------------------------------
# Load model & corpus
# ------------------------------
print("Loading saved fastText model...")
model = fasttext.load_model(MODEL_PATH)
model_vocab = set(model.get_words())
print(f"Model loaded. Vocabulary size: {len(model_vocab)}")

print("Building corpus frequencies...")
with open(TRAINING_FILE, "r", encoding="utf-8") as f:
  all_tokens = f.read().split()
corpus_freq = Counter(all_tokens)

most_frequent = {w for w, _ in corpus_freq.most_common(N_FUNCTION_WORDS)}
filter_set = most_frequent | EXPLICIT_LOVE_VARIANTS

# ------------------------------
# Helpers
# ------------------------------
```

```
def safe_unit_vector(vec):
  norm = np.linalg.norm(vec)
  if norm == 0:
    return vec
  return vec / norm


def get_anchor_vector(model, word):
  if word not in model_vocab:
    raise ValueError(
      f"Anchor word '{word}' is not in the model vocabulary. "
      f"Frequency in corpus: {corpus_freq.get(word, 0)}"
    )
  return safe_unit_vector(model.get_word_vector(word))


def reconstruct_love_centroid(model, love_forms):
  vectors, used_forms = [], []
  for form in love_forms:
    if form in model_vocab:
      vectors.append(safe_unit_vector(model.get_word_vector(form)))
      used_forms.append(form)
  if not vectors:
    raise ValueError("No love-family forms available for centroid.")
  centroid = np.mean(np.vstack(vectors), axis=0)
  return safe_unit_vector(centroid), used_forms


def get_neighbourhood(model, query_vec, exclude_words, corpus_freq,
          min_freq, filter_set, top_n):
  exclude_set = set(exclude_words) | filter_set
  rows = []
  for word in model.get_words():
    if word in exclude_set:
      continue
    if corpus_freq.get(word, 0) < min_freq:
      continue
    vec = safe_unit_vector(model.get_word_vector(word))
    sim = float(np.dot(query_vec, vec))
    rows.append((word, sim, corpus_freq[word]))
  df = pd.DataFrame(rows, columns=["word", "cosine_similarity", "frequency"])
  df = df.sort_values("cosine_similarity", ascending=False).reset_index(drop=True)
  return df.head(top_n)


def style_axes(ax):
  ax.set_facecolor("white")
  for spine in ax.spines.values():
    spine.set_color("black")
    spine.set_linewidth(0.5)
  ax.tick_params(
    axis="both", which="both",
    length=2, color="black",
    labelsize=6, labelcolor="black",
  )
```

```
def plot_panel(ax, coords, words, composed_pos, anchor):
  ax.scatter(
    coords[:, 0], coords[:, 1],
    s=8, facecolors="white", edgecolors="black",
    linewidths=0.4, zorder=2,
  )
  ax.scatter(
    composed_pos[0], composed_pos[1],
    s=35, c="black", edgecolors="black",
    linewidths=0.4, zorder=5,
  )
  texts = [
    ax.text(
      composed_pos[0], composed_pos[1],
      f"love + {anchor}",
      fontsize=7.5, fontstyle="italic", fontweight="bold",
      color="black", zorder=6,
    )
  ]
  for i, word in enumerate(words):
    texts.append(
      ax.text(
        coords[i, 0], coords[i, 1], word,
        fontsize=6, fontstyle="italic",
        color="black", alpha=0.9, zorder=3,
      )
    )
  adjust_text(
    texts, ax=ax,
    expand_points=(1.3, 1.3),
    expand_text=(1.15, 1.15),
    force_text=(0.3, 0.5),
    force_points=(0.2, 0.3),
    arrowprops=dict(arrowstyle="-", color="black", lw=0.2, alpha=0.5),
  )
  style_axes(ax)

# ------------------------------
# Build love centroid
# ------------------------------
centroid, used_forms = reconstruct_love_centroid(model, LOVE_FORMS)
print(f"\nLove centroid built from: {used_forms}")

# ------------------------------
# Additive composition
# ------------------------------
print(f"\n{'='*60}")
print("Additive composition (love + anchor)")
print(f"{'='*60}")

os.makedirs(OP_A_DIR, exist_ok=True)

baseline_neighbours = get_neighbourhood(
  model, centroid, used_forms, corpus_freq,
  NEIGHBOUR_MIN_FREQ, filter_set, top_n=TOP_N
```

```
)
baseline_neighbours.to_csv(f"{OP_A_DIR}/baseline_love_centroid.csv", index=False)

operation_a_results = {"baseline": baseline_neighbours}
composed_vectors = {}
anchor_neighbours = {}

for anchor in ANCHORS:
  print(f"\n--- love + {anchor} ---")
  try:
    anchor_vec = get_anchor_vector(model, anchor)
  except ValueError as e:
    print(f" SKIPPED: {e}")
    continue

  composed = safe_unit_vector(centroid + anchor_vec)
  composed_vectors[anchor] = composed

  exclude = list(used_forms) + [anchor]
  neighbours = get_neighbourhood(
    model, composed, exclude, corpus_freq,
    NEIGHBOUR_MIN_FREQ, filter_set, top_n=TOP_N
  )
  operation_a_results[anchor] = neighbours
  anchor_neighbours[anchor] = neighbours.head(TOP_N_PER_PANEL)["word"].tolist()

  safe_anchor = re.sub(r"\W+", "_", anchor)
  neighbours.to_csv(f"{OP_A_DIR}/love_plus_{safe_anchor}.csv", index=False)
  print(neighbours.head(20).to_string(index=False))

comparison_data = {
  anchor: df["word"].head(TOP_N).tolist()
  for anchor, df in operation_a_results.items()
}
max_len = max(len(v) for v in comparison_data.values())
for k in comparison_data:
  comparison_data[k] += [""] * (max_len - len(comparison_data[k]))
pd.DataFrame(comparison_data).to_csv(f"{OP_A_DIR}/comparison_table.csv", index=False)
print(f"\nOperation A results saved to: {OP_A_DIR}/")

# ------------------------------
# VISUALISATION (5 rows x 3 cols)
# ------------------------------
active_anchors = [a for a in ANCHORS if a in composed_vectors]
print(f"\nActive anchors for visualisation: {active_anchors}")

plt.rcParams.update({
  "font.family": "serif",
  "font.serif": ["Times New Roman", "DejaVu Serif", "serif"],
  "mathtext.fontset": "stix",
})

print("\n--- Generating per-panel-space figure (5x3) ---")

fig, axes = plt.subplots(5, 3, figsize=(15, 20))
axes_flat = axes.flatten()
```

```
for i, anchor in enumerate(active_anchors):
  ax = axes_flat[i]
  words = anchor_neighbours[anchor]

  neighbour_vectors = np.vstack([
    safe_unit_vector(model.get_word_vector(w)) for w in words
  ])
  panel_vectors = np.vstack([composed_vectors[anchor], neighbour_vectors])

  panel_reducer = umap.UMAP(
    n_neighbors=min(10, len(words)),
    min_dist=0.15,
    metric="cosine",
    random_state=42,
  )
  panel_coords = panel_reducer.fit_transform(panel_vectors)
  plot_panel(ax, panel_coords[1:], words, panel_coords[0], anchor)

for j in range(len(active_anchors), len(axes_flat)):
  axes_flat[j].set_visible(False)

legend_handles = [
  Line2D([0], [0], marker="o", color="w", markerfacecolor="black",
      markeredgecolor="black", markersize=8, linewidth=0,
      label="Love centroid"),
  Line2D([0], [0], marker="o", color="w", markerfacecolor="white",
      markeredgecolor="black", markersize=6, linewidth=0,
      label="Nearest neighbour"),
]
fig.legend(
  handles=legend_handles,
  loc="lower center",
  ncol=2,
  fontsize=8,
  frameon=True,
  edgecolor="black",
  fancybox=False,
  bbox_to_anchor=(0.5, 0.0),
)
plt.tight_layout(rect=[0, 0.02, 1, 1])

perpanel_path = f"{OUTPUT_DIR}/panels_per_panel_space.png"
plt.savefig(perpanel_path, dpi=300, bbox_inches="tight")
plt.show()
print(f"Saved: {perpanel_path}")

""""# context level-embeddings""""

import os
import glob
import numpy as np
import torch
import torch.nn.functional as F
from transformers import AutoTokenizer, AutoModel
from tqdm import tqdm
```

```
import nltk
from typing import List

nltk.download('punkt_tab', quiet=True)

class E5MistralVectorizer:
  def __init__(self, chunk_size=3500, overlap=200, batch_size=1):
    """
    Args:
      chunk_size: E5 Mistral supports up to 4096 tokens.
          Using 3500 leaves room for special tokens.
      overlap: Sentences to carry over between chunks.
      batch_size: Keep at 1 for 7B model on single H100 to avoid OOM.
    """
    self.device = 'cuda' if torch.cuda.is_available() else 'cpu'
    self.model_name = "intfloat/e5-mistral-7b-instruct"

    print(f"Loading {self.model_name} in FP16...")
    self.tokenizer = AutoTokenizer.from_pretrained(self.model_name)
    self.model = AutoModel.from_pretrained(
      self.model_name,
      torch_dtype=torch.float16,
      device_map="auto",    # Automatically places on GPU
      trust_remote_code=True          # Required for E5 models
    )
    self.model.eval()

    self.chunk_size = chunk_size
    self.overlap = overlap
    self.batch_size = batch_size
    self.dim = 4096

    # GPU memory report
    print(f"VRAM used after loading: {torch.cuda.memory_allocated() / 1e9:.2f} GB")

  def _chunk_text(self, text: str) -> List[str]:
    """Token-aware chunking to respect the 4096 limit precisely."""
    sentences = nltk.sent_tokenize(text)
    chunks = []
    current_chunk = []
    current_len = 0

    for sent in sentences:
      # Count tokens using the actual tokenizer for precision
      sent_token_count = len(self.tokenizer.encode(sent, add_special_tokens=False))

      if current_len + sent_token_count > self.chunk_size and current_chunk:
        chunks.append(" ".join(current_chunk))

        # Overlap logic: re-add last few sentences based on token count
        overlap_tokens = 0
        overlap_sents = []
        for s in reversed(current_chunk):
          t = len(self.tokenizer.encode(s, add_special_tokens=False))
          if overlap_tokens + t <= self.overlap:
            overlap_sents.insert(0, s)
```

```
            overlap_tokens += t
          else:
            break
        current_chunk = overlap_sents
        current_len = overlap_tokens

      current_chunk.append(sent)
      current_len += sent_token_count

    if current_chunk:
      chunks.append(" ".join(current_chunk))
    return chunks

  def _mean_pooling(self, model_output, attention_mask):
    """E5 Mistral uses mean pooling over the last hidden state."""
    token_embeddings = model_output[0] # (batch, seq_len, 4096)
    input_mask_expanded = attention_mask.unsqueeze(-1).expand(token_embeddings.size()).float()
    return torch.sum(token_embeddings * input_mask_expanded, 1) / torch.clamp(input_mask_expanded.sum(1),
min=1e-9)

  @torch.no_grad()
  def encode_documents(self, documents: List[str]) -> np.ndarray:
    doc_vectors = []

    for doc_idx, doc in enumerate(tqdm(documents, desc="Encoding with E5 Mistral 7B")):
      if not doc or not doc.strip():
        doc_vectors.append(np.zeros(self.dim))
        continue

      chunks = self._chunk_text(doc)
      if not chunks:
        doc_vectors.append(np.zeros(self.dim))
        continue

      chunk_vecs = []
      for chunk in chunks:
        # E5-Mistral expects instruction prefix for best quality
        input_text = f"Instruct: Retrieve semantically similar text.\nQuery: {chunk}"

        inputs = self.tokenizer(
          input_text,
          max_length=4096,
          padding=True,
          truncation=True,
          return_tensors="pt"
        ).to(self.device)

        outputs = self.model(**inputs)
        embedding = self._mean_pooling(outputs, inputs['attention_mask'])

        # Normalize (E5 models expect cosine similarity, so L2 norm is standard)
        embedding = F.normalize(embedding, p=2, dim=1)
        chunk_vecs.append(embedding.cpu().numpy().flatten())

      # Mean pool over chunks
      doc_vec = np.mean(chunk_vecs, axis=0)
```

```
      doc_vectors.append(doc_vec)

      # Clear GPU cache periodically
      if (doc_idx + 1) % 10 == 0:
        torch.cuda.empty_cache()

    return np.array(doc_vectors)

# ============================================================================
# EXECUTION
# ============================================================================

folder_path = "/content/corpus_modernized"
txt_files = sorted(glob.glob(os.path.join(folder_path, "*.txt")))
print(f"Found {len(txt_files)} text files.")

corpus_texts = []
corpus_filenames = []
for fp in tqdm(txt_files, desc="Reading files"):
  with open(fp, 'r', encoding='utf-8') as f:
    corpus_texts.append(f.read())
  corpus_filenames.append(os.path.basename(fp))

vectorizer = E5MistralVectorizer(
  chunk_size=3500,  # 3.5k tokens per chunk
  overlap=150       # Small overlap for context
)

# This will take ~10-15 minutes for 159 docs (approx 2-3 chunks per doc)
document_embeddings = vectorizer.encode_documents(corpus_texts)

print(f"\nFinal E5 Mistral vector space: {document_embeddings.shape}")
np.save('/content/document_embeddings_e5_mistral.npy', document_embeddings)

""""
concept_similarity_e5_mistral_lewis.py
---------------------------------------

Compare Middle English document embeddings to C. S. Lewis's four loves:

  1. Storge = Affection / familial-natural love
  2. Philia = Friendship / fellowship / shared values
  3. Eros = Romantic-passionate love
  4. Agape  = Charity / divine-sacrificial love

Workflow:
 1. Load precomputed E5-Mistral document embeddings.
 2. Encode Lewis-style concept prompts with the SAME E5-Mistral model.
 3. Average each set of prompts into a concept centroid.
 4. Compare every document to each centroid using cosine similarity.
 5. Save and print ranked results.

Assumptions:
 - document_embeddings_e5_mistral.npy was created with intfloat/e5-mistral-7b-instruct
 - filenames_e5_mistral.npy matches the embedding order
 - document embeddings are shape (N, 4096)
```

```
Install:
  pip install numpy pandas scikit-learn sentence-transformers
“““

import os
import numpy as np
import pandas as pd
from sklearn.metrics.pairwise import cosine_similarity
from sentence_transformers import SentenceTransformer

# ============================================================================
# 1. CONFIG
# ============================================================================

DOCUMENT_EMBEDDINGS_PATH = “/content/document_embeddings_e5_mistral.npy”
FILENAMES_PATH = “/content/filenames_e5_mistral.npy”

OUTPUT_CSV = “/content/concept_similarity_e5_mistral_lewis.csv”

MODEL_NAME = “intfloat/e5-mistral-7b-instruct”

EXPECTED_DIM = 4096

BATCH_SIZE = 2

# ============================================================================
# 2. LOAD DOCUMENT EMBEDDINGS AND FILENAMES
# ============================================================================

if not os.path.exists(DOCUMENT_EMBEDDINGS_PATH):
  raise FileNotFoundError(f”Could not find: {DOCUMENT_EMBEDDINGS_PATH}”)

if not os.path.exists(FILENAMES_PATH):
  raise FileNotFoundError(f”Could not find: {FILENAMES_PATH}”)

document_embeddings = np.load(DOCUMENT_EMBEDDINGS_PATH)
filenames = np.load(FILENAMES_PATH, allow_pickle=True)

filenames = [
  str(f)
  .replace(“_normalized.txt”, ““)
  .replace(“.txt”, ““)
  for f in filenames
]

print(“Loaded document embeddings:”, document_embeddings.shape)
print(“Loaded filenames:”, len(filenames))

if document_embeddings.ndim != 2:
  raise ValueError(
    f”document_embeddings should be 2D, but got shape {document_embeddings.shape}”
  )

if document_embeddings.shape[1] != EXPECTED_DIM:
  raise ValueError(
```

```
    f"Expected {EXPECTED_DIM}-dimensional E5-Mistral embeddings, "
    f"but got {document_embeddings.shape[1]} dimensions."
  )

if len(filenames) != document_embeddings.shape[0]:
  raise ValueError(
    f"Number of filenames ({len(filenames)}) does not match "
    f"number of document embeddings ({document_embeddings.shape[0]})."
  )


# ====================================================================================
# 3. SAFE NORMALIZATION FUNCTION
# ====================================================================================

def l2_normalise(matrix):
  """
  L2-normalise vectors safely.
  Works for both 1D and 2D arrays.
  """

  matrix = np.asarray(matrix, dtype=np.float32)

  if matrix.ndim == 1:
    norm = np.linalg.norm(matrix)
    if norm == 0:
      raise ValueError("Cannot normalise a zero vector.")
    return matrix / norm

  norms = np.linalg.norm(matrix, axis=1, keepdims=True)
  norms[norms == 0] = 1.0

  return matrix / norms


document_embeddings = l2_normalise(document_embeddings)

# ====================================================================================
# 4. DEFINE C. S. LEWIS CONCEPTS
# ====================================================================================

concepts = {
  "storge_affection": [
    "Storge, or affection, in C. S. Lewis's sense: natural love, familial attachment, parent and child, kinship, familiar
comfort, and household bonds.",
    "Natural affection between family members, parents, children, siblings, relatives, and kin.",
    "Familiar love rooted in habit, closeness, comfort, tenderness, and everyday companionship.",
    "Love as domestic attachment, nurture, care, protection, and belonging within the household.",
    "Affection expressed through family, home, familiarity, ordinary loyalty, and long-standing attachment.",
  ],

  "philia_friendship": [
    "Philia, or friendship, in C. S. Lewis's sense: deep friendship, fellowship, shared values, intellectual kinship,
loyalty, and common purpose.",
    "Friendship as companionship, fellowship, trust, counsel, loyalty, and shared moral commitment.",
```

```
    "Love between friends founded on mutual respect, common interests, shared virtues, and chosen companionship.",
    "Bonded friendship, sworn fellowship, brotherhood, counsel, trust, loyalty, and equality between companions.",
    "Companionship and intellectual, ethical, or spiritual kinship between people who share a common vision.",
  ],

  "eros_romantic_love": [
    "Eros, or romantic love, in C. S. Lewis's sense: passionate love focused on the beloved, desire, longing, devotion, and romantic attachment.",
    "Romantic love, courtship, desire, longing, beloved, lover, passion, admiration, and amorous pursuit.",
    "Love directed toward the beloved, involving desire, service, emotional suffering, devotion, and idealisation.",
    "Courtly love, erotic longing, romantic pursuit, chivalric devotion, and passionate attachment to a beloved.",
    "Passionate love that may become selfless commitment but may also become idolatry, obsession, or dangerous desire.",
  ],

  "agape_charity_divine_love": [
    "Agape, or charity, in C. S. Lewis's sense: divine love, unconditional love, selfless sacrifice, grace, mercy, and spiritual charity.",
    "Charity as divine love, sacrificial love, mercy, forgiveness, grace, compassion, and unconditional goodness.",
    "Love of God, love from God, Christian charity, spiritual devotion, redemptive love, and divine grace.",
    "Selfless love that orders, purifies, and guides natural loves toward goodness, sacrifice, and holiness.",
    "Divine charity, mercy, forgiveness, self-giving love, spiritual transformation, and sacrificial compassion.",
  ],
}

# ============================================================================
# 5. LOAD E5-MISTRAL MODEL
# ============================================================================

print(f"\nLoading model: {MODEL_NAME}")
print("This may require substantial RAM/GPU memory.")

model = SentenceTransformer(MODEL_NAME)

# For short prompts this does not matter much, but it keeps the setting explicit.
model.max_seq_length = 4096

print("Model loaded.")

# ============================================================================
# 6. BUILD CONCEPT CENTROIDS
# ============================================================================

def encode_concept_prompts(seed_phrases):
  """
  Encode concept prompts using E5-Mistral's instruction-style query format.
  """

  instruction = (
    "Instruct: Given a concept from literary theory and medieval literary studies, "
    "retrieve Middle English documents that are semantically related to that concept.\n"
    "Query: "
  )

  embeddings = model.encode(
```

```
    seed_phrases,
    prompt=instruction,
    normalize_embeddings=True,
    batch_size=BATCH_SIZE,
    show_progress_bar=True,
  )

  embeddings = np.asarray(embeddings, dtype=np.float32)

  if embeddings.ndim != 2:
    raise ValueError(
      f"Expected 2D concept embeddings, but got shape {embeddings.shape}"
    )

  if embeddings.shape[1] != document_embeddings.shape[1]:
    raise ValueError(
      f"Dimension mismatch. Concept embeddings have {embeddings.shape[1]} dimensions, "
      f"but document embeddings have {document_embeddings.shape[1]} dimensions. "
      "You must use the same embedding model for both."
    )

  return embeddings

def build_centroid(seed_phrases):
  """
  Build a single normalised centroid from several seed prompt embeddings.
  """

  seed_embeddings = encode_concept_prompts(seed_phrases)
  seed_embeddings = l2_normalise(seed_embeddings)

  centroid = seed_embeddings.mean(axis=0)
  centroid = l2_normalise(centroid)

  return centroid


centroids = {}

for concept_name, seed_phrases in concepts.items():
  print(f"\nBuilding centroid for: {concept_name}")
  centroids[concept_name] = build_centroid(seed_phrases)

print("\nAll concept centroids built.")

# ============================================================================
# 7. SCORE DOCUMENTS AGAINST EACH LEWIS CATEGORY
# ============================================================================

results = pd.DataFrame({
  "filename": filenames
})

for concept_name, centroid in centroids.items():
  scores = cosine_similarity(
    document_embeddings,
```

```
    centroid.reshape(1, -1)
  ).ravel()

  results[f"{concept_name}_similarity"] = scores

# ================================================================================
# 8. ADD DOMINANT CATEGORY AND CONTRAST SCORES
# ================================================================================

love_columns = [
  "storge_affection_similarity",
  "philia_friendship_similarity",
  "eros_romantic_love_similarity",
  "agape_charity_divine_love_similarity",
]

missing_columns = [col for col in love_columns if col not in results.columns]

if missing_columns:
  raise ValueError(f"Missing expected similarity columns: {missing_columns}")

results["dominant_love_type"] = (
  results[love_columns]
  .idxmax(axis=1)
  .str.replace("_similarity", "", regex=False)
)

results["eros_minus_agape"] = (
  results["eros_romantic_love_similarity"]
  - results["agape_charity_divine_love_similarity"]
)

results["agape_minus_eros"] = (
  results["agape_charity_divine_love_similarity"]
  - results["eros_romantic_love_similarity"]
)

results["philia_minus_eros"] = (
  results["philia_friendship_similarity"]
  - results["eros_romantic_love_similarity"]
)

results["storge_minus_eros"] = (
  results["storge_affection_similarity"]
  - results["eros_romantic_love_similarity"]
)

results["eros_minus_philia"] = (
  results["eros_romantic_love_similarity"]
  - results["philia_friendship_similarity"]
)

results["agape_minus_storge"] = (
  results["agape_charity_divine_love_similarity"]
  - results["storge_affection_similarity"]
)
```

```
# ============================================================================
# 9. SORT MAIN RESULTS
# ============================================================================

# Main ranking: closest to Eros, since this is probably your most useful
# literary-love category for romance/courtly material.
results = results.sort_values(
  “eros_romantic_love_similarity”,
  ascending=False
).reset_index(drop=True)

results.index += 1

# ============================================================================
# 10. PRINT USEFUL TABLES
# ============================================================================

display_columns = [
  “filename”,
  “dominant_love_type”,
  “storge_affection_similarity”,
  “philia_friendship_similarity”,
  “eros_romantic_love_similarity”,
  “agape_charity_divine_love_similarity”,
]

print(“\n” + “=“ * 80)
print(“TOP 30 DOCUMENTS BY EROS / ROMANTIC LOVE”)
print(“=“ * 80)
print(results[display_columns].head(30).to_string())

print(“\n” + “=“ * 80)
print(“TOP 20 DOCUMENTS WHERE EROS > AGAPE”)
print(“=“ * 80)
print(
  results.sort_values(“eros_minus_agape”, ascending=False)[
    [
      “filename”,
      “eros_romantic_love_similarity”,
      “agape_charity_divine_love_similarity”,
      “eros_minus_agape”,
      “dominant_love_type”,
    ]
  ].head(20).to_string(index=False)
)

print(“\n” + “=“ * 80)
print(“TOP 20 DOCUMENTS WHERE AGAPE > EROS”)
print(“=“ * 80)
print(
  results.sort_values(“agape_minus_eros”, ascending=False)[
    [
      “filename”,
      “agape_charity_divine_love_similarity”,
      “eros_romantic_love_similarity”,
```

```
    "agape_minus_eros",
    "dominant_love_type",
   ]
  ].head(20).to_string(index=False)
)

print("\n" + "=" * 80)
print("TOP 20 DOCUMENTS WHERE PHILIA > EROS")
print("=" * 80)
print(
  results.sort_values("philia_minus_eros", ascending=False)[
   [
    "filename",
    "philia_friendship_similarity",
    "eros_romantic_love_similarity",
    "philia_minus_eros",
    "dominant_love_type",
   ]
  ].head(20).to_string(index=False)
)

print("\n" + "=" * 80)
print("TOP 20 DOCUMENTS WHERE STORGE > EROS")
print("=" * 80)
print(
  results.sort_values("storge_minus_eros", ascending=False)[
   [
    "filename",
    "storge_affection_similarity",
    "eros_romantic_love_similarity",
    "storge_minus_eros",
    "dominant_love_type",
   ]
  ].head(20).to_string(index=False)
)

# ============================================================================
# 11. SAVE RESULTS
# ============================================================================

results.to_csv(OUTPUT_CSV, index_label="rank")

print("\n" + "=" * 80)
print(f"Saved results to: {OUTPUT_CSV}")
print("=" * 80)

# ============================================================
# REVISED THREE-CATEGORY LOVE-DISCOURSE MODEL
# ============================================================
#
# This block follows the E5-Mistral document-embedding and Lewis
# analysis already given above in Appendix B. It reuses:
#    document_embeddings
#    filenames
#    model
#    l2_normalise
```

```
#
# The 4,096-dimensional document embeddings are unchanged.
# Only the category centroids used for comparison are replaced.

import os
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt

from adjustText import adjust_text
from matplotlib.lines import Line2D
from sklearn.decomposition import PCA
from sklearn.metrics.pairwise import cosine_similarity
from sklearn.preprocessing import StandardScaler


# ------------------------------------------------------------
# 1. CONFIGURATION
# ------------------------------------------------------------

OUTPUT_DIR = "/content/love_discourse_figures"
os.makedirs(OUTPUT_DIR, exist_ok=True)

FIG_DPI = 300
BATCH_SIZE = 1

# Thresholds used to distinguish clear from mixed/uncertain signals.
RAW_QUANTILE_THRESHOLD = 0.40
STRONG_Z_MIN = 1.25
STRONG_Z_MARGIN_MIN = 0.50
MODERATE_Z_MIN = 0.75
MODERATE_Z_MARGIN_MIN = 0.35


# ------------------------------------------------------------
# 2. DEFINE THE REVISED THREE-CATEGORY MODEL
# ------------------------------------------------------------

concepts = {
    "secular_erotic": [
        "Medieval English texts centrally concerned with ribald, bawdy, comic, or fabliau-like erotic desire, sexual appetite, adultery, cuckoldry, lust, bodily humour, and sexual deception.",
        "Texts where erotic desire is comic, crude, bodily, deceptive, adulterous, transactional, obscene, or satirical rather than idealised or courtly.",
        "Fabliau-style narratives involving sexual tricks, bawdy jokes, lust, adultery, erotic rivalry, comic punishment, bodily appetite, and anti-romantic sexual comedy.",
        "Desire represented through ribaldry, sexual appetite, bawdy humour, trickery, cuckoldry, obscene comedy, or degraded erotic pursuit.",
        "Comic sexual deception, adultery plots, bawdy erotic humour, lust, bodily appetite, fabliau comedy, and anti-romantic sexuality.",
        "This category includes Chaucerian fabliau-like eroticism, sexual trickery, cuckoldry, obscene humour, and comic bodily desire; it excludes refined courtly devotion and divine charity.",
    ],

    "courtly_romantic": [
```

```
        "Medieval English texts centrally concerned with courtly or romantic love: lovers, beloveds, longing, courtship, desire, service to the beloved, fidelity, jealousy, and amorous suffering.",
        "Texts where love is idealised, refined, elevated, and directed toward a beloved through gentilesse, courtesy, knightlihood, chivalric conduct, service, and emotional devotion.",
        "Courtly love, romance love, fin amour, noble desire, lover-beloved relations, refined speech, beauty, longing, idealisation, fidelity, and the suffering lover.",
        "Narratives or poems about romantic attachment, courtship, amorous complaint, chivalric service, gentilesse, courtesy, honour, knightly identity, and devotion to a beloved.",
        "Serious or elevated love discourse, romance, courtly complaint, dream vision love, chivalric love, and love shaped by gentilesse, courtesy, honour, and knightlihood.",
        "This category includes Troilus-like romantic love, courtly complaint, romance love, noble longing, chivalric devotion, and refined desire; it excludes ribald sexual comedy.",
    ],

    "theologico_philosophical": [
        "Medieval English religious, devotional, mystical, theological, contemplative, or philosophical texts centrally concerned with divine love, caritas, charity, love of God, grace, mercy, salvation, and spiritual transformation.",
        "Texts where love is primarily divine, theological, contemplative, mystical, philosophical, or devotional: love of God, love from God, caritas, grace, mercy, redemption, salvation, and the soul's relation to God.",
        "Affective piety, divine charity, Christian love, contemplative love, philosophical meditation on goodness, grace, wisdom, the soul, virtue, mercy, and spiritual self-giving.",
        "Texts where love is understood through theology, Christian philosophy, contemplation, divine wisdom, moral instruction, caritas, mercy, forgiveness, salvation, or redemptive compassion.",
        "Mystical devotion, contemplative theology, philosophical reflection on love and goodness, divine wisdom, virtue, spiritual transformation, and the ordering of the soul toward God.",
        "This category includes Julian-like divine love, mystical devotion, caritas, theological reflection, philosophical love, grace, mercy, salvation, and wisdom; it does not include texts merely because they mention priests, sin, punishment, or general morality.",
    ],
}


# ------------------------------------------------------------
# 3. ENCODE CATEGORY DESCRIPTIONS AND BUILD CENTROIDS
# ------------------------------------------------------------

INSTRUCTION_PREFIX = (
    "Instruct: Retrieve medieval English literary documents where the following "
    "type of love-discourse is explicit, central, and thematically sustained. "
    "Distinguish carefully between: "
    "1) secular-erotic love expressed through bodily appetite, sexual comedy, "
    "transaction, adultery, deception, or fabliau-like desire; "
    "2) courtly-romantic love shaped by gentilesse, courtesy, knightlihood, "
    "chivalric service, longing, honour, and devotion to a beloved; "
    "and 3) theologico-philosophical love concerned with caritas, God, grace, "
    "salvation, wisdom, virtue, mercy, goodness, contemplation, and spiritual "
    "transformation. Do not retrieve a text merely because it contains general "
    "morality, religion, sexuality, social conflict, or narrative action.\n"
    "Query: "
)


def encode_concept_descriptions(descriptions):
    """Encode category descriptions with E5-Mistral's instruction format."""
    prompted = [INSTRUCTION_PREFIX + text for text in descriptions]
```

```
    embeddings = model.encode(
        prompted,
        normalize_embeddings=True,
        batch_size=BATCH_SIZE,
        show_progress_bar=True,
    )

    embeddings = np.asarray(embeddings, dtype=np.float32)

    if embeddings.shape[1] != document_embeddings.shape[1]:
        raise ValueError(
            "Concept and document embeddings must have the same dimension."
        )

    return embeddings


def build_centroid(descriptions):
    """Average the description embeddings into one normalized centroid."""
    embeddings = encode_concept_descriptions(descriptions)
    embeddings = l2_normalise(embeddings)
    centroid = embeddings.mean(axis=0)
    return l2_normalise(centroid)


centroids = {
    category: build_centroid(descriptions)
    for category, descriptions in concepts.items()
}


# ------------------------------------------------------------
# 4. CALCULATE CATEGORY-SIMILARITY PROFILES
# ------------------------------------------------------------

results = pd.DataFrame({"filename": filenames})

raw_cols = []
for category, centroid in centroids.items():
    col = f"{category}_raw"
    raw_cols.append(col)
    results[col] = cosine_similarity(
        document_embeddings,
        centroid.reshape(1, -1),
    ).ravel()

# Contrastive score = similarity to one category minus the mean
# similarity to the other two categories.
adjusted_cols = []
for category in concepts:
    this_col = f"{category}_raw"
    other_cols = [
        f"{other}_raw"
        for other in concepts
        if other != category
    ]
```

```
    adjusted_col = f"{category}_adjusted"
    adjusted_cols.append(adjusted_col)
    results[adjusted_col] = (
        results[this_col] - results[other_cols].mean(axis=1)
    )

# Standardize each contrastive dimension across the corpus.
z_cols = []
for col in adjusted_cols:
    z_col = f"{col}_z"
    z_cols.append(z_col)

    std = results[col].std(ddof=0)
    if std == 0:
        results[z_col] = 0.0
    else:
        results[z_col] = (
            results[col] - results[col].mean()
        ) / std


# ------------------------------------------------------------
# 5. CATEGORY ASSIGNMENT AND UNCERTAINTY
# ------------------------------------------------------------
#
# Colour and marker shape are assigned separately from PCA position.
# Colour records the category with the strongest adjusted z score.
# Marker shape records whether that assignment is clear or uncertain.

z_matrix = results[z_cols].to_numpy()
category_names = list(concepts.keys())
top_index = np.argmax(z_matrix, axis=1)

results["assigned_category"] = [
    category_names[i] for i in top_index
]
results["top_z"] = z_matrix[np.arange(len(results)), top_index]

sorted_z = np.sort(z_matrix, axis=1)
results["z_margin"] = sorted_z[:, -1] - sorted_z[:, -2]

results["love_discourse_relevance"] = results[raw_cols].max(axis=1)
raw_threshold = results["love_discourse_relevance"].quantile(
    RAW_QUANTILE_THRESHOLD
)

strong_signal = (
    (results["top_z"] >= STRONG_Z_MIN)
    & (results["z_margin"] >= STRONG_Z_MARGIN_MIN)
)

moderate_signal = (
    (results["love_discourse_relevance"] >= raw_threshold)
    & (results["top_z"] >= MODERATE_Z_MIN)
    & (results["z_margin"] >= MODERATE_Z_MARGIN_MIN)
```

```
)

results["signal"] = np.where(
    strong_signal | moderate_signal,
    "clear",
    "mixed_or_uncertain",
)


# ------------------------------------------------------------
# 6. PCA POSITIONING
# ------------------------------------------------------------
#
# PCA is applied to the complete three-score love-discourse profile.
# This operation determines point position; it does not assign colour.

X = results[z_cols].to_numpy()
X_scaled = StandardScaler().fit_transform(X)

pca = PCA(n_components=2)
pca_2d = pca.fit_transform(X_scaled)

results["PC1"] = pca_2d[:, 0]
results["PC2"] = pca_2d[:, 1]

#  Figure 9: secular-erotic at left,
# theologico-philosophical at right,
# and courtly-romantic lower down.
 secular_x = results.loc[
    results["assigned_category"] == "secular_erotic", "PC1"
].mean()
theological_x = results.loc[
    results["assigned_category"] == "theologico_philosophical", "PC1"
].mean()

if theological_x < secular_x:
    results["PC1"] *= -1

secular_y = results.loc[
    results["assigned_category"] == "secular_erotic", "PC2"
].mean()
courtly_y = results.loc[
    results["assigned_category"] == "courtly_romantic", "PC2"
].mean()

if secular_y < courtly_y:
    results["PC2"] *= -1

print(
    "PCA explained variance:",
    f"PC1 = {pca.explained_variance_ratio_[0]:.1%},",
    f"PC2 = {pca.explained_variance_ratio_[1]:.1%}",
)
```

```
category_display = {
    "secular_erotic": "Secular / erotic",
    "courtly_romantic": "Courtly / romantic",
    "theologico_philosophical": "Theologico-philosophical",
}

category_colours = {
    "secular_erotic": "#E69F00",
    "courtly_romantic": "#C84A3A",
    "theologico_philosophical": "#2F7D32",
}


def lighter(hex_colour, amount=0.55):
    """Lighten a category colour for mixed/uncertain assignments."""
    value = hex_colour.lstrip("#")
    rgb = np.array(
        [int(value[i:i + 2], 16) for i in (0, 2, 4)],
        dtype=float,
    )
    rgb = rgb + (255 - rgb) * amount
    rgb = np.clip(rgb, 0, 255).astype(int)
    return "#{:02x}{:02x}{:02x}".format(*rgb)


fig, ax = plt.subplots(figsize=(18, 14))

for category in category_display:
    for signal, marker in [
        ("clear", "o"),
        ("mixed_or_uncertain", "^"),
    ]:
        subset = results[
            (results["assigned_category"] == category)
            & (results["signal"] == signal)
        ]

        if subset.empty:
            continue

        face_colour = category_colours[category]
        if signal == "mixed_or_uncertain":
            face_colour = lighter(face_colour)

        ax.scatter(
            subset["PC1"],
            subset["PC2"],
            s=38,
            marker=marker,
            c=face_colour,
            edgecolors="black",
            linewidths=0.35,
            alpha=0.95,
            zorder=3,
        )
```

```
# Label every text.
texts = []
for _, row in results.iterrows():
    texts.append(
        ax.text(
            row["PC1"],
            row["PC2"],
            str(row["filename"]).replace("_", " "),
            fontsize=5.3,
            alpha=0.90,
        )
    )

adjust_text(
    texts,
    ax=ax,
    expand=(1.03, 1.08),
    force_text=(0.08, 0.12),
    force_static=(0.05, 0.08),
)

# Zero lines create descriptive quadrants only; they do not classify texts.
ax.axhline(0, linewidth=0.55, alpha=0.35)
ax.axvline(0, linewidth=0.55, alpha=0.35)

ax.set_xlabel(
    f"PC1 ({pca.explained_variance_ratio_[0]:.1%} variance)"
)
ax.set_ylabel(
    f"PC2 ({pca.explained_variance_ratio_[1]:.1%} variance)"
)
ax.set_title(
    "PCA map of love-discourse profiles — lighter shades mark uncertain signals",
    fontsize=11,
)

category_handles = [
    Line2D(
        [0], [0],
        marker="o",
        linestyle="",
        markerfacecolor=category_colours[category],
        markeredgecolor="black",
        markeredgewidth=0.4,
        markersize=7,
        label=label,
    )
    for category, label in category_display.items()
]

category_legend = ax.legend(
    handles=category_handles,
    title="Category",
    loc="upper right",
    frameon=True,
    fontsize=8,
```

```
    title_fontsize=8,
)
ax.add_artist(category_legend)

signal_handles = [
    Line2D(
        [0], [0],
        marker="o",
        linestyle="",
        markerfacecolor="0.55",
        markeredgecolor="black",
        markersize=6,
        label="Clear signal",
    ),
    Line2D(
        [0], [0],
        marker="^",
        linestyle="",
        markerfacecolor="0.82",
        markeredgecolor="black",
        markersize=6,
        label="Mixed / uncertain signal",
    ),
]

ax.legend(
    handles=signal_handles,
    title="Signal",
    loc="lower right",
    frameon=True,
    fontsize=8,
    title_fontsize=8,
)

plt.tight_layout()

png_path = os.path.join(OUTPUT_DIR, "figure_9_revised_three_category_pca.png")
pdf_path = os.path.join(OUTPUT_DIR, "figure_9_revised_three_category_pca.pdf")

plt.savefig(png_path, dpi=FIG_DPI, bbox_inches="tight")
plt.savefig(pdf_path, bbox_inches="tight")
plt.show()

print(f"Saved: {png_path}")
print(f"Saved: {pdf_path}")
```