HOW THE ANTERIOR CINGULATE CORTEX AND THE HIPPOCAMPUS REPRESENT A DECISION MAKING TASK: A POPULATION ENCODING APPROACH

DAVID PERE TOMAS CUESTA
Bachelor of Biomedical Engineering, University Pompeu Fabra, 2017

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

MASTER OF SCIENCE

in

NEUROSCIENCE

Department of Neuroscience
University of Lethbridge
LETHBRIDGE, ALBERTA, CANADA

HOW THE ANTERIOR CINGULATE CORTEX AND THE HIPPOCAMPUS REPRESENT A DECISION MAKING TASK: A POPULATION ENCODING APPROACH

DAVID PERE TOMAS CUESTA

Date of Defence: January 04, 2021

Dr. Aaron Gruber Thesis Supervisor	Associate Professor	Ph.D.
Dr. Matthew Tata Thesis Examination Committee	Professor	Ph.D.
Dr. Masami Tatsuno Thesis Examination Committee	Associate Professor	Ph.D.
Dr. Robert Sutherland Chair, Thesis Examination Committee	Professor	Ph.D.

Abstract

The anterior cingulate cortex (ACC) appears to be involved in cognitive control, learning, and memory. Recent proposals integrating these functions suggest that the ACC embodies a cognitive map of task space. In this view, ACC links decision and reward information to particular locations in an abstract map of a task environment, allowing for planning and control. Here, I present novel evidence for task space by showing that, in rats, ACC spatial ensemble activity is preserved when performing a task on a familiar layout, but in a novel spatial context. Furthermore, this encoding similarity is sufficient to decode the animal's position in the new context using the "map" of the familiar layout. I contrast this with hippocampal activity and encoding, which drastically changes when switching spatial contexts. This sheds light into how the brain might encode decision and reinforcement information together with spatial cues to exert top-down control on behavior.

Contents

List of Figures

1 Introduction

In our daily lives, we are barraged with a constant stream of multimodal sensory data. How do we make sense, structure, and act upon these rich amounts of information?

In 1948 Tolman proposed the concept of a cognitive map [1]; a mental representation of an environment that serves as a scaffold to store information about particular locations. Decades later, neurons showing preference to specific locations in an environment, the place cells, were discovered in the rat hippocampus (HPC) [2]. As such, the HPC was proposed to contain a cognitive map [3], with further studies suggesting that it also stores information associated with these locations [4][5][6][7].

Cognitive maps are spatial models of the world. Besides being an effective way to link spatial locations with non-spatial information, they allow the bearer to reason about the environment and to plan strategies [8]. Once a cognitive map has formed, it might be used to predict outcomes and seek actions that are conducive to desirable results. This kind of prospective activity has been seen in HPC forward sweeps. In a maze choice point, the rat HPC produces a series of patterned activity representing a 'virtual walk-through' of the candidate future locations [9][10].

The HPC provides a compelling example of a cognitive map, but it might not be the only area in the brain that embodies a spatial map of the environment. The medial prefrontal cortex (mPFC) is an area known for encoding decision-related reward variables and to mediate both learning and decision making [11][12]. An area in the mPFC, called the anterior cingulate cortex (ACC), combines spatial representation [13][14][15], together with a variety of task features related to reinforcement and decision making [16][17][18][19]. Furthermore, the rat ACC appears to generate predictive information related to these decision variables

[15], in striking similarity to the forward sweeps observed in HPC.

Not all cognitive maps are created equal. The HPC is well known for its role in encoding episodic memories [20][21][22]. The cognitive map in the rat HPC supports the encoding of these memories through a phenomenon called remapping. In remapping, changes in cue configurations in an environment result in substantial change in the firing rates of place cells (rate remapping) [6][23]. If these changes are large enough, or the same task is moved to a different location, place cells will often completely change their spatial tuning over the environment (global remapping) [24][25]. Thus, remapping in HPC modulates or directly changes the cognitive map layout in response to contextual changes in the environment that signal particular events, forming the basis for episodic memory processing [26].

On the other hand, the ACC has been implicated in a wide range of cognitive functions, including conflict monitoring [27], performance and reward evaluation [28][29], emotional processing [30], task switching [31][32], cognitive control [33][34][35], persistence [36][37], and others. In an attempt to reconcile this diversity of function, the ACC has been proposed to form schematic knowledge of the world from the episodic memories encoded in the HPC [38][39]. The formation of this schematic knowledge is a memory consolidation process [40][41] that is mediated by hippocampal - ACC connections [42].

The schematic knowledge crystallizes in the form of associative schemas, mental models of the world (not necessarily spatial) that allow to efficiently integrate information and use it to make deductive inferences and control future behavior [43]. The ACC has been proposed to contain a cognitive map that embodies a schema into task space, an abstract and dynamic spatial representation of an environment [44] [45]. Similar arguments have also been made for the orbitofrontal cortex [46].

In summary, while the HPC cognitive map links episodic and contextual information to locations in the physical space, the ACC cognitive map appears to link abstract spatial location with task-relevant information for control and decision making.

The aim of my thesis work is to better understand the similarities and differences among

the ACC and HPC maps in order to elucidate how they contribute to behavioural control. Specifically, I sought to test the hypothesis that the ACC map is based in 'task-space', whereas the HPC map is based in physical space. If so, I predict that moving a familiar task layout to a novel spatial context will produce qualitatively different changes in the ACC and HPC cognitive maps. Insofar as ACC encodes task space, switching the spatial context, while preserving the same task structure and layout, will preserve its map. To my knowledge, this has never been shown previously. In contrast, the HPC map will change when the spatial context varies, via global remapping, even if the task remains unchanged.

To test this hypothesis, we built a double figure-8-maze with shared reward feeders (figure 1A). Rats learn the task on one side of the figure-8, and then are tested on the novel alternate side, which shares the layout but differs in position (figure 1B). Using recordings of ACC and HPC ensembles containing activity of tens of individual neurons, I use several methods to quantify the similarity of neural encoding on the two sides of the maze. In particular, first I look at whether individual neurons preserve their spatial firing patterns. Then I check whether the whole population level dynamics across the track, within and between trials, are shared between the sides. Finally, I investigate whether the map for the known side can be used to decode the position of the animal in the novel side.

2 Methods

Note: This is a computational and data analysis project. The data I am using has been collected by other members in the lab.

2.1 Subjects

Subjects are two Fisher-Brown Norway rats (4-9 months old), born on site in our colony. Animals were food restricted for the duration of the behavioural testing and their weight monitored so as to not drop below 85% of their baseline weight. All procedures were approved by the university's animal welfare committee (Protocol #1512) in accordance with the Canadian Council on Animal Care.

2.2 Ensemble recording

We used established techniques for recording ensemble activity with high-density electrophysiology [15]. Briefly, we used aseptic surgical techniques to implant onto the rat's head a 'hyperdrive' device that houses several probes that are independently lowered into the ACC (rat #1), or HPC (rat #2). These probes consist of four twisted wires, or tetrodes, that are electroplated with gold to an impedance of 200-400 kΩ. The position of the tetrodes was verified by histological methods post-mortem. Neural signals and subject position were recorded with a digital acquisition system (Cheetah SX, Neuralynx). Neural signals were amplified with a gain of 1000, digitized at 30kHz, and recorded to a hard drive. The waveforms of action potentials were detected offline using semi-automated techniques. The recording data was preprocessed and cluster-cut using MClust [47].

The ACC implanted rat yielded 69 neurons. The HPC implanted rat yielded 65 neurons. Neurons in ACC usually had firing activity all across the track, while the activity of most HPC neurons was concentrated in particular locations, as expected of place cells (supplementary figure 1).

2.3 Behavioural apparatus

The behavioral setup consists of a mirrored figure-8-maze designed and built in-house (figure 1A). The apparatus is 1m off the ground with the dimensions of 173 x 123 cm and a track width of 8 cm. The maze has 4 reward feeders and 12 doors that are controlled independently by a custom-made Matlab script (The MathWorks Inc.) through an Arduino Mega board. Each door is opened and closed through a servo motor and each reward well has nose-poke and lick sensors. The rewards are delivered automatically with the help of a pressure system and the amounts are controlled through a solenoid valve. The position of the rat is established by a thermal camera mounted on the ceiling which determines the appropriate sequence of events (doors open/close, rewards are delivered/withheld). A speaker mounted on the center of the maze delivers an auditory cue when the rat reaches a reward well.

2.4 Experimental design

We recorded ensemble neural activity in ACC (n=1) or HPC (n=1) as rats performed a 2-choice task between 2 choice feeders (north, south). The layout of the task is mirrored across two sides, A and B, which share choice feeders (figure 1B). The rats were trained on side A for 3 weeks before the surgery until they performed the task with minimum bias, and retrained for 3-4 days unplugged and then 4-7 days plugged.

On the recording day, the rats were placed, as in training, on side A. Halfway trough the session duration, they were directed to side B for the first time (with the use of mechanical doors blocking the other paths), for the other half of the session.

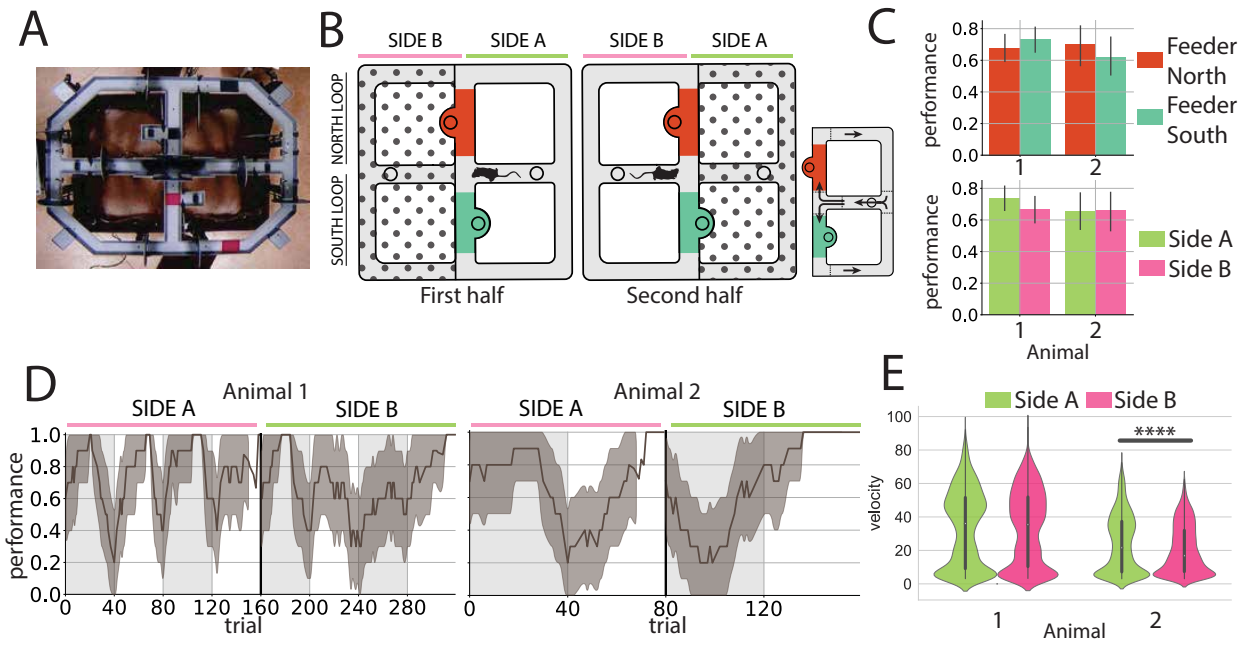

Figure 1: **A**. Top view of the double figure-8-maze. **B**. In the first half of the session the rat is placed on the side A of the maze, where it was trained in previous sessions. A trial starts with the animal in the (always rewarded) central feeder. After arriving at the bifurcation, the animal has to decide whether to go through the north loop or the bottom loop. Each one of the loops has a feeder which yields reward with a different probability. Both loops end at the center feeder, which finalizes the trial. On the second half of the trials the animal is switched to the novel side B, which has not been seen in training, and the task is mirrored. Circles indicate feeders. Inset: Arrows indicate the path of the rat in a given side. **C**. Top: The animal's performance as a function of the most frequently rewarding feeder shows similar preference between the feeders. Bottom: The performance of the animal between sides is also similar. Error bars indicate 95% confidence intervals. **D**. Performance across trials. The different blocks are shaded grey in the background. As expected, the performance tends to rise up at the beginning of each block change. Errors show 95% confidence intervals in the moving average **E**. Velocity distributions of the animals, between the sides.

In a trial (Figure 1B, inset), the rat starts at the center feeder, which delivers a small reward on every trial (chocolate ensure). This motivates rats to return to the start position. The animals run down the middle segment of the maze, and then may turn either north or south to access one of two choice feeders. Finally, rats then run along the lateral arms of the track to return to the center feeder and finish the trial. A set of mechanical doors block the rat from backtracking.

The first rat session lasted 66min, and consisted of 319 trials. The second rat session

lasted 68min, and consisted of 159 trials.

The two choice feeders deliver reward with a probability of 0.8 and 0.2 respectively. These percentages are interchanged between the choice feeders in blocks of 40 trials within each session. Therefore, rats must shift choice preference within the session to maximize reward. The center feeder always delivers reward.

A trial can be either free (with 0.7 probability) or forced (with 0.3 probability). On forced trials, the gates direct the animal to one of the two choice feeders, selected at random ($p = 0.5$), while on free trials, the animal can freely choose whether to go to the north or south choice feeders. I define performance as the ratio of free trials in which the animal chooses the optimal feeder (the feeder that gives reward with the highest probability for that block).

The overall performance was 0.7 for the first animal (ACC recorded animal) and 0.66 for the second (HPC recorded animal) (supplementary figure 2A). The performance was largely independent of which feeder was the optimal one (figure 1C, top), with $pval = 0.38$ and 0.34 for the first and second animal respectively (two-sided proportions z-test, z-val: 0.88, -0.94, dof: 173, 92). The performance was also largely independent of the side (figure 1C, bottom), with $pval = 0.26$ and 0.95 for the first and second animal respectively (two-sided proportions z-test, z-val: 1.13, -0.06, dof: 174.5, 91.0,). Generally, the performance for each block started low, and increased as the animal found the optimal feeder (figure 1D, 10 trial moving average).

I compute the speed of the animals by taking the numerical derivative of the position across time with a time window (dt = 0.17sec). To remove jumps due to inaccuracies of the tracker, I remove the outliers at 3σ. When comparing between sides, I also remove the first 5 trials of the new side, as they are considerably longer. The velocity distributions were largely bi-modal, indicating a stereotypical running speed (supplementary figure 2B). The second animal velocity was consistently lower than the first, and completed roughly as half as many trials. The distribution of speeds between sides (figure 1E, supplementary 2C) was largely

similar for the first animal ($pval = 0.07$), but differed significantly ($pval = 2.2 \times 10^{-36}$) for the second (two-sided Welch's t-test, tval: 0.31, 12.7, dof: 16888, 12513). In the track, speed was highest on the horizontal central arm and at the four distant corners (supplementary figure 2C).

2.5 Statistical analysis

All analysis and statistical tests were done in Python using numerical and statistical libraries [48]. In order to show the results of statistical tests in the figures, I assign a p value of < 0.05 with "$*$", p value of < 0.01 with "$**$", p value of < 0.001 with "$***$", and a p value of < 0.0001 with "$****$".

In the box plots, outliers are identified as points falling outside 50% of the interquartile range.

To show errors, unless otherwise specified, I compute the 95% confidence interval using bootstrapping with 1000 randomly drawn samples from the data (with replacement) [49].

2.6 Correlational analysis

First, I divide each side of the track into 12 spatial bins. For each bin and side I count the number of spikes for each neuron and divide them by the time spent in that bin to get a rate (spikes/sec). I then map the rates per bin between the sides, ending up, for each cell, with two vectors of 12 values. To pair the rates per bin between the sides, I use two different mappings between the bins of the two sides, corresponding to mirroring or rotation of the space (see figure 2.A). In rotation, the bins labels in side B are rotated 180 degrees with respect to side A. In mirroring, the bin labels in side B are mirrored across the center of the track with respect to side A. I then compute the Pearson correlation between the rate vectors of the mapped sides, for each neuron. To remove spurious correlations, if a given neuron only fires in 3 or less bins, I exclude it from the analysis. To compute the histograms for other subsets of bins, I compute the correlation for these bins exclusively. As a control, I split the

trials in side A in two halves of equal duration, and compare the spatial correlations between the two halves on that side.

2.7 Trial averaged PCA

First, I remove neurons that are tonically spiking by removing neurons with an average firing rate higher than 2σ over the mean. This removes 3 neurons in each animal. I also remove unusually long trials by removing trials with a duration larger than 3σ over the mean duration. This has the effect of removing 3 trials in each animal.

To average over trials I need all trials to have the same length. To achieve this without losing the spatial location of the activity, I first split each trial into 6 mutually exclusive spatial segments (figure 3A). The first part comprises the time the animal spends in the center feeder. The second part starts at the center feeder and ends at the vertical arm of the maze. The third part starts at the vertical arm and ends at the side feeder. The fourth part is the time spent in the side feeder. The fifth part comprises from the side feeder exit to the exit of the central vertical arm. The sixth part contains the track marking the return to the central feeder. For each part in each trial I time-wrap the spike times to the median length of that part over all trials. To do this, I obtain the ratio of the length of the segment in that trial to the median length of the segment over all trials. Then, I multiply the spike times with that ratio, in order to expand or contract the spike train in the segment. Thus, the relative spike-times in the segment are not modified, but time dilated/contracted to mach the median segment duration.

Next, I create a grid of discrete time-steps ($dt = 0.01sec$). For each neuron, I sum the spikes that occurred inside every time-step. I end up with a matrix of dimensions cells x time for each trial. I then smooth the resulting matrix across time with a Gaussian kernel with $\sigma = 2sec$ and normalize each neuron activity so it peaks at 1. Finally, I bin the activity in time with $dt = 0.05sec$.

Finally, I group the matrices based on the trial's side and the lateral loop (choice feeder),

and I average the matrices in the same group. I use this data for principal component analysis (PCA).

2.8 TCA

As with the trial averaged PCA, by time-wrapping the spikes in 6 segments around the track I can form a matrix of the same size for each trial. To improve the convergence of TCA, I use a kernel of $sigma = 4sec$ to smooth the spike data, and bin it in time with a $dt = 0.2sec$.

Instead of averaging over trials (as in PCA), I stack the trials to form a tensor of size cells x time x trials. I then use the tensorly python package [50] to decompose that tensor into a set of time, cell and trial factors. Briefly, the time factors show activity levels across the trial duration. The neuronal activities across the trial can be reconstructed by taking weighted sums of these time factors. These weights are given by the cell factors. In addition, the time factors are amplitude modulated in each trial by some amount that is given by the trial factors. See figure 4A for a more complete explanation. To obtain factors that are easier to interpret, I use a non negative parafac decomposition [51].

2.9 Decoders

First, I create a grid of discrete time-steps ($dt = 0.001sec$). For each neuron, I sum the spikes that occurred inside a certain time-step. I then smooth the resulting tensor across time with a Gaussian kernel with $\sigma = 0.15sec$, following previous work [15]. Next, I average the neural activity in blocks of $dt = 0.05sec$. I divide each side of the track into 12 spatial bins, as in 2.6, and assign to each block the spatial bin corresponding to the median animal position in the block. Finally, I z-score the block activities. These pairs of a z-scored neural activity averaged over the $0.05sec$ block and their corresponding spatial bin are respectively the input and targets of my decoder.

As I will have a different number of these blocks for each animal and side, I first equalize

the data by using only as many blocks as the side with less blocks across animals (this turns out to be the side A in the HPC recorded animal). This leaves us with 27 755 blocks per each side per each animal (1388 sec). Unless otherwise specified, I grab these blocks from the end of the time in each side to avoid the habituation period when entering the side.

I use side A of the maze for training and validation and side B for testing. In side A I use 80% of the data for training and 20% for validation, without shuffling (as the neural activity was smoothed and contiguous samples usually share the same spatial bin, shuffling here would artificially boost my validation score by effectively "copying" data from the training to the validation set). For decoding, I use a 4 layer artificial neural network with $(300, 150, 150, 12)$ units. After each layer I apply heavy dropout ($p = 0.5$) and batch-normalize the activities to avoid over-fitting. See supplementary figure 12A for the training and validation curves as a function of the epoch. I trained the network until the validation loss stabilized and then took the weights of the best epoch.

2.10 Nonlinear dimensionality reduction

I use Isomap [52], a nonlinear dimensionality reduction algorithm, to separately analyze the input data that I feed to my decoders, and the activities of the second to last decoder layer (before the soft-max layer) (see 2.9).

As Isomap is heavy on computation time, before applying the dimensionality reduction I downsample the data by a factor of 2 with a order 8 Chebyshev type I filter. To make the results easier to interpret I equalize the data between the classes, so every bin has the same number of data points.

Thus, my data consists of a set of points in a dimensional space of size equal to the number of neurons. Isomap computes a low dimensional embedding of this data that tries to preserve its intrinsic geometry.

Isomap is a variant of multidimensional scaling (MDS). Classical MDS computes a lower-dimensional embedding; each point in the original space is projected to a lower dimensional

space such that the euclidean distances between the data-points in the high-dimensional space are maintained in the lower-dimensional space. Isomap uses MDS, but tries instead to preserve the geodesic distance between the points. To compute the geodesic distance, Isomap builds a graph of the data by connecting every datapoint with its closest neighbors. For every two datapoints, the geodesic distance is computed as the shortest path between the points, i.e., the minimum number of connections traversed to get from one point to the other. Embedding the data in a lower dimensional space that preserves geodesic distances instead of euclidean distances allows to better maintain the shape of the high-dimensional manifold in the lower-dimensional projection.

Besides giving visually convincing results on our data, Isomap allows to transform new data onto the embedding learned in previously embedded data. I make use of this by learning an embedding in side A and using it to transform the data from sides A and B of the maze.

3 Results

3.1 Neurons in ACC preserve their spatial firing patterns between sides

I first sought out to quantify the similarity of the individual neuron spatial responses between the two sides of the track. For this, I divide the familiar side of the maze in 12 spatial bins. I then map these bins to the new side of the maze following two hypothesis: In the first, the cognitive map is orientation dependent (egocentric), and changing sides produces a rotation in the maze layout representation (figure 2A, left). In the second, the cognitive map is orientation independent (allocentric), and changing sides maintains (mirrors) the same layout representation (figure 2A, right).

Once I have the spike rate for each bin and neuron in the two sides, I match the bins of side A with their counterparts in side B according to the desired mapping (rotation or mirroring), and compute the correlation across the bins, between the sides (see 2.6 for details).

To avoid direct comparison between the two areas, which come from different animals and therefore from different behaviors, I also compute, as a control, the correlation in side A by splitting the first half of the session.

In both ACC and HPC (figure 2B), spatial correlations across the two halves of side A is high (ACC 0.89 median, HPC 0.83 median). When going from side A to side B, however, the correlation decreases significantly for both rotation and mirroring in HPC (pval 1.50×10^{-14} and pval 6.52×10^{-13}, respectively, Welch ttest, tval: -8.73, -8.01, dof: 122, 125), and for mirroring in ACC (pval 1.11×10^{-5}, pval 0.17 for mirroring and rotation respectively, Welch ttest, tval: -4.56, -1.37, dof: 131, 134). There is not a significant drop in rotation in ACC,

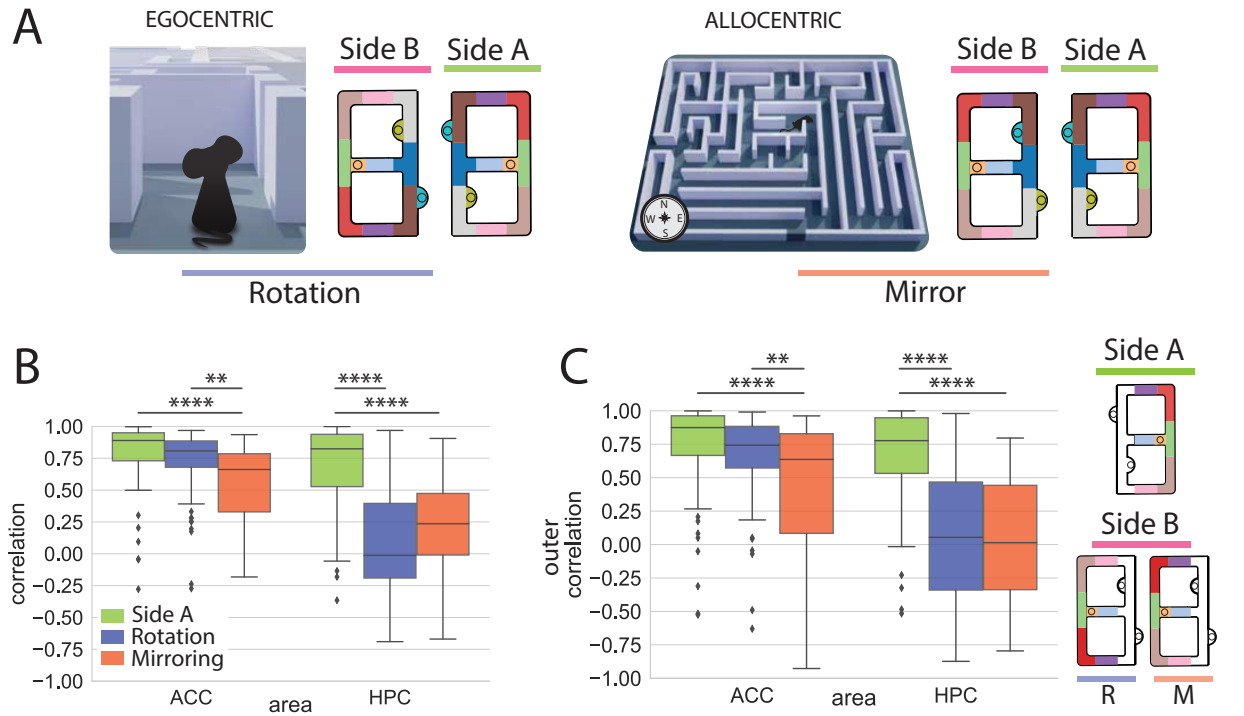

Figure 2: **A**. When going from the familiar to the novel side of the maze, the cognitive map might rotate, following the relative position of the animal (egocentric view, left), or mirror keeping the same orientation in both sides (allocentric view, right). This corresponds to a rotation and a mirroring of the bin labels, respectively. **B**. A higher proportion of neurons preserve their spatial firing patterns in ACC when going from the familiar to the novel side if the track layouts of the sides are related by a rotation. **C**. The same results also hold on the outer part of the track, showing that the similarity of spatial firing patterns between the sides is not exclusively due to their shared space.

which indicates that the spatial firing patterns are largely conserved between the two sides, and depend on the relative orientation of the animal (pval 1.00×10^{-3}, pval 0.12 between rotation and mapping in ACC and HPC, respectively, Welch ttest, tval: 3.36, -1.54, dof: 130, 123).

In our track, both sides share the central vertical arm, which might over-represent the spatial similarity between sides. To make sure the above results hold in the portion of the maze that is novel, I compute the correlations excluding the bins that belong to this arm (figure 2C). The results are largely similar (pval 6.15×10^{-5} and 6.54×10^{-13} between side A and mirror for ACC and HPC (Welch ttest, tval: -5.92, -5.73, dof: 131, 121), pval 0.24 and 3.63×10^{-11} between side A and rotation for ACC and HPC (Welch ttest, tval: -1.79,

-7.03, dof: 133, 120), pval 1.28×10^{-3} and 0.79 between rotation and mirroring for ACC and HPC (Welch ttest, tval: 4.28, -1.80, dof: 131, 124), indicating that the spatial firing patterns are conserved in ACC for rotation in the novel space of the track. Looking at the central vertical arm instead (supplementary figure 3C) shows more similarity in the HPC responses, specially for mirroring, while still significantly different from same side correlations.

Another possible source of overestimated correlations is shared reward encoding in the feeders. Alas, removing the feeders from the analysis yields similar results (supplementary figure 3A). Interestingly, removing the feeders increases the separation between rotation and mirroring in the ACC, which might hint that the feeder locations are mixed in the neural encoding.

Because some bins don't change between the mirror and rotation mappings, the differences between between rotation and mirroring are attenuated. I can highlight these differences by only considering spatial correlations across the bins that change between mappings (supplementary figure 3B). Doing the opposite removes rotation mirroring differences (supplementary figure 3D).

In summary, as expected, I see that more neurons in ACC seem to preserve their firing patterns between the sides, and this is greater with a rotation mapping rather than with a mirroring mapping.

3.2 ACC preserves trial dynamics

In the previous section I have shown preservation of the spatial firing patterns by looking at spike rate averages over bins. In here, I take a more nuanced approach and look at whether the neural activity across side A trials is preserved on side B trials.

Furthermore, instead of looking at individual neurons I focus on the overall population dynamics. As a measure of the dynamics of the brain area, I use principal component analysis (PCA) to find a low dimensional summary of the ensemble neural dynamics that preserve as much variance in the neuronal activities as possible. This allows for some wiggle room in

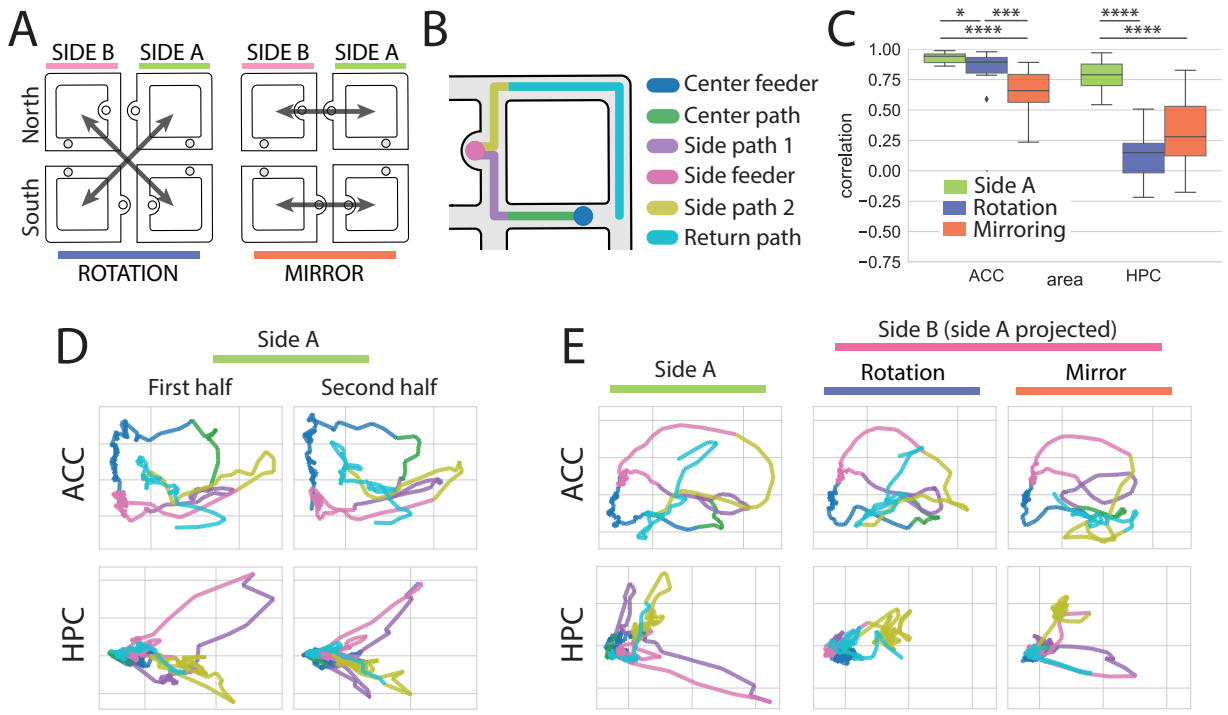

Figure 3: **A**. Notion of rotation and mirroring in trials by mapping opposite or parallel loops between the sides, respectively. **B**. To preserve activity location in the track when comparing different trials, I divide the track into 6 parts and wrap each one to the median length across trials. **C**. Comparison of the correlation of the first 8 principal components between sides A, B (for rotation and mirroring), and for the first and second halves of side A. ACC shows more preservation of the dynamics, specially when comparing contralateral loops (rotation). **D**. Activity of the two principal components while the animal visits the track. The topology of the dynamics is preserved in the first and second halves of side A. **E**. The topology of the dynamics is preserved for rotation in ACC. Notice the shared encoding of the feeders.

the neuron's individual responses; the population activity will still be similar between sides as long as a certain kind of weighted averages (which maximize the variance of activity) of the neuron activities are similar between the sides.

Thus, instead of looking at whether the proportions of the spike rates of individual neurons at different locations are similar, I look at whether the sequence of population activity across the track is preserved between the sides.

As described in 2.4, a trial consists of a loop trough the track, starting from center feeder, trough one of the choice feeders (either north or south), and back to the center feeder. Thus, depending on the side and north / south loop, I can classify trials in 4 groups. To have

a notion of rotation and mirroring mappings, I will compare, between sides, the activities across the track in the trials of different loops (north or south) for rotation, and the same loop on mirroring (figure 3A). To increase signal to noise, I average the trials in the same group together.

The goal of this analysis is to compare activities across the track. However, because the animal takes a different amount of time to complete each trial, I need to make sure that I am comparing the activities of corresponding spatial locations. To accomplish this, I divide each trial into 6 sections (figure 3B, see 2.7 for details), and wrap each trial so all the sections are consistent between all the trials.

To compare these low dimensional summary of population dynamics between sides I first compute the principal component decomposition for side A (see 2.7 for details). Then, I project the activities of side B to the component space computed in side A. This has the effect of assuming that all the neurons have the same loading's (i.e. each neuron draws from the same components in both sides). I can now focus on the similarity between the components (scores). A high component similarity implies that the dynamics of the population across the track are maintained between sides. As before, I contrast the side A / B differences with the component similarities between the first and second halves of side A.

First, I look at the aggregated correlations between sides for the first 8 components (aggregating both lateral loops for a total of 16 data-points), enough to explain $> 70\%$ of variance (supplementary figure 4A, supplementary figure 5). While the component correlations of both areas are significantly different from the same side controls, pval 1.91×10^{-5} and 9.06×10^{-6} between side A and mirror for ACC and HPC (Welch ttest, tval: -2.15, -5.42, dof: 256, 232), pval 0.03 and 8.57×10^{-11} between side A and rotation for ACC and HPC, (Welch ttest, tval: 1.53, -8.41, dof: 254, 206), the drop in correlation is much larger for HPC than for ACC. Furthermore, rotation in ACC is, again, significantly better at explaining the relationship between the sides, pval 3.71×10^{-4} and 0.05 between rotation and mirroring for ACC and HPC, (Welch ttest, tval: -4.01, 3.02, dof: 260, 234).

To investigate where this similarity comes from, I compute the component correlations for specific parts of the track. As before, the ACC similarity is not due to shared space in the vertical center arm of the track (supplementary figure 6A). This analyses also allows us to look at the dynamics at the feeders, which are highly preserved between sides and indistinguishable from mirroring or rotation for ACC, but not for HPC (supplementary figure 6B). Individual correlations for the different parts of the track show a similar picture across the track (supplementary figure 7).

While these analyses only include the first 8 components (supplementary figure 8A), which is somewhat arbitrary, the same findings hold when averaging over more components (supplementary figure 8B).

Similarly, plotting the two first components in the 2D plane (figure 3D-E, also see supplementary figure 8C for the other lateral loop) shows that the trajectory of the neural dynamics is better preserved between sides on the ACC (with the rotation mapping specifically).

In summary, as expected, the trial dynamics are better preserved in ACC than HPC. As in the previous analysis, rotation gives a closer match between the sides than mirroring. In HPC, the pattern is reversed, and the mirroring shows more correlation than rotation on the central arm. Looking at individual neurons instead of at PCA components gives a similar interpretation (supplementary figure 8).

3.3 ACC preserves trial by trial dynamics

A disadvantage of trial-averaging methods such as the one used in 3.2 is that we ignore the trial to trial variability, effectively treating it as noise. Instead, we might want to take into account how the dynamics change across trials or depend on the trial characteristics. For example, in rewarded trials, the response of some neurons might be drastically different than on non-rewarded. While it is possible to account for this using methods such as trial-concatenated PCA, a more elegant solution is to extend the principles of PCA to find factors that directly reconstruct the full neurons x time x trials tensor [53]. This approach is called

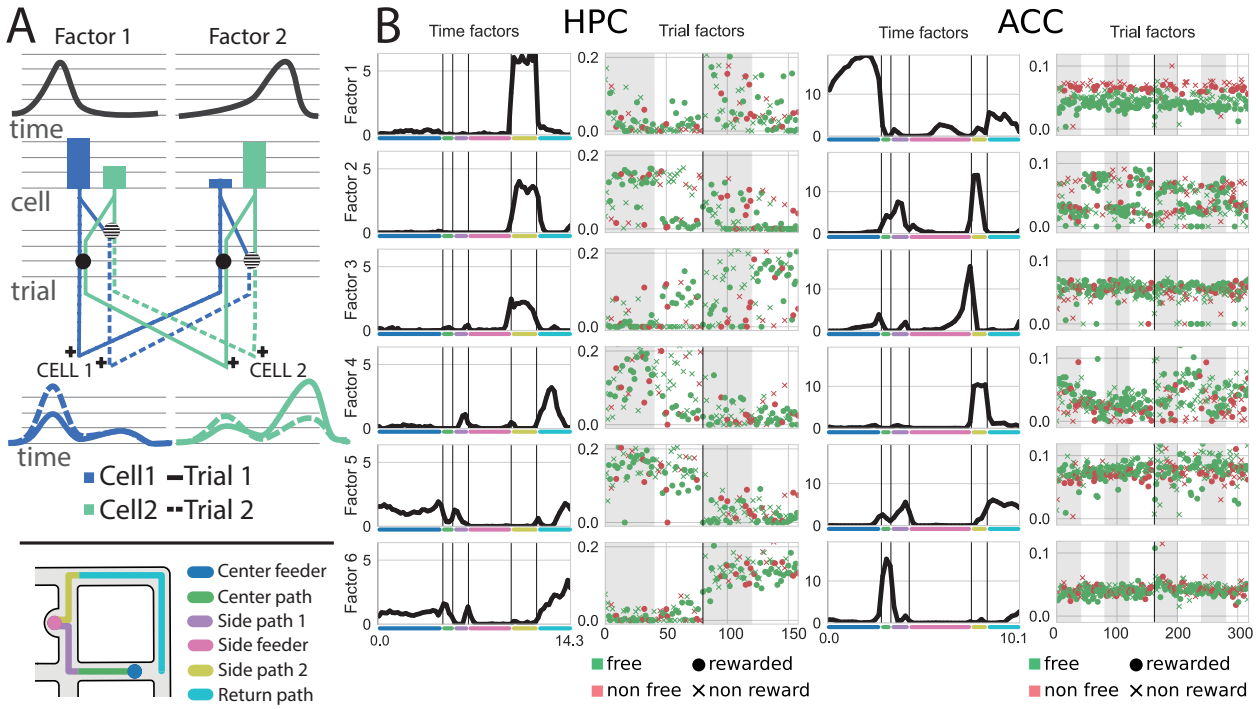

Figure 4: **A**. Schematic of tensor component analysis (TCA) as a gain modulated linear network model. The neural activity of two neurons in two different trials is decomposed into a set of two time, cell and trial factors. In TCA, cell activities are approximated by a linear combination of gain modulated time factors. The trial factors determine the amplitude of the time factors on that particular trial, while the cell factors determine the linear combination weights among the modulated time factors. **B**. Time and trial factor for 6 of the factors of a 8 component non-negative decomposition. Time factors show demixed dynamics on the trial, while trial factors, labeled with the trial type and reward, show variable dependent encoding. Some of the trial factors disassociate between free and forced or rewarded and non-rewarded trials. Notice the large irregularities between sides present in the trial factors for some HPC factors, implying a radical shift in the trial by trial dynamics.

tensor component analysis (TCA)[54].

TCA has a nice neuroscientific interpretation as fitting a gain-modulated linear network model to neural data [53]. As an example, let's say that we are trying to fit the responses of two neurons with two factors as in figure 4A. TCA finds a set of time factors, cell factors and trial factors that reconstruct the data. The time factors provide the time dependent basis from which the neuronal responses will be reconstructed. As in PCA, each neuron response will be approximated as a weighted sum of time factors, with the weights given by the cell factors. However, in TCA the time factors are further multiplied, i.e., gain modulated, by

the trial factor before the weighted sum of components, allowing to take into account trial to trial variability by modifying the amplitude of the corresponding basis. Figure 4A provides an example schematic showcasing this interpretation.

Another advantage of TCA is that it does not require the cell factors or the time factors to be orthogonal. This has the benefit of demixing the data, beyond reducing its dimensionality [53]. However, in contrast to PCA, solving TCA does not guarantee to find the optimal solution, and each solution is particular to the number of components used.

For ease of interpretability, I used a non-negative tensor decomposition (see 2.8). I computed the variance explained per number of components (supplementary Figure 4B), which turns out to be lower in HPC. This might be in part due the higher variability between the two sides. As I found no clear inflexion point in the curves, I decided to choose 8 components to match the PCA analysis.

The temporal and trial factors of 6 of these components are shown in figure 4C (see supplementary figures 10 and 11 for complete list). As with PCA, the time factors are highly localized in the trial. The trial factors are displayed according to the trial type (free, forced) and reward (rewarded, non-rewarded). Some of the trial factors (1 in ACC) seem selective to whether the trial is free or not, whether others (3 in ACC, 8 in HPC) seem selective to whether the trial was rewarded, others multiplex feeder and reward (2 in ACC, 2 in HPC).

Interestingly, a large percentage of the trial factors in HPC (2,4,5,6,7) have clear discontinuities between the sides, which might indicate that the dynamics across trials drastically change when switching sides. This follows the direction of the previous results. On the contrary, the ACC components seem to be preserved across sides. When the factors depend on the trial characteristics (free / not free, rewarded / non-rewarded), this trial factor structure is also largely preserved between sides, which hints at the ACC preserving the encoding of task relevant variables.

In summary, the TCA analysis seem to indicate that the neural trial-to-trial dynamics, together with the neural encoding of trial-dependent variables, is largely preserved between

sides in ACC, but not in HPC.

3.4 ACC preserves position encoding

Previous research has been successful in extracting the position of the animal from its ACC ensemble activity [13][15]. Therefore, I set out to test whether the similarities in the spatial encoding observed in the previous section are sufficient to coarsely decode the position of the animal. In other words, I investigate whether learning the map that the rat uses in side A can be used to find the animal in side B.

To that end, I trained a decoder to extract the spatial information encoded in the neural activity on the side A, and see if the learned structure is useful to decode position in side B. For ease of interpretation and to provide an easier task for the decoder, I convert the regression problem of detecting the animal position into a classification problem of deciding in which bin is the animal currently located. Akin to 3.1, I divide each side of the track into 12 bins and map the bins either by a rotation or mirroring (figure 5A).

I trained and validated a 4 layer neural network in side A and tested it in side B (see 2.9 for details). As a control, I use the validation set of side A, which was not used to train the network. The network outputs the probability of the mouse being in each bin (supplementary figure 13). In the following, I take the highest of these probabilities as the bin that was predicted by the network.

The predicted bins at each position of the rat (figure 5B) show good matches to the true bins for side A in both ACC and HPC. This holds for side B in the ACC, which clearly favors a rotation mapping, even if this is not forced in any way by the decoder. In HPC, however, only the center segment seems to be predicted reliably.

To get a better idea of where the network makes a mistake I plot, for each position of the rat, whether the predicted bin was correct or not (Figure 5C). This reveals an even prediction error spread across the track on side A, with lower accuracy at bin boundaries for both areas and for side B rotation in the ACC. Again, for HPC only the bin at the center of the maze

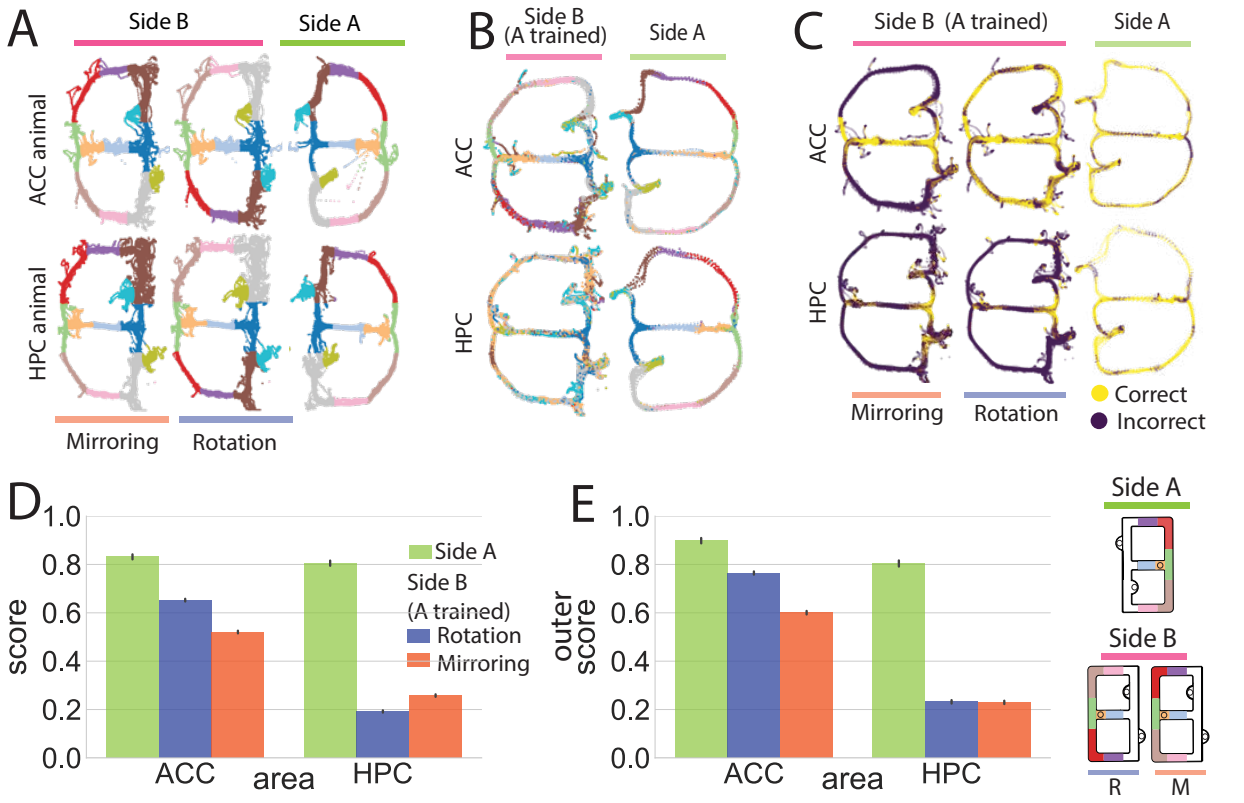

Figure 5: **A**. I spatially bin the position of the animal, and map these bin labels for side A to side B, either with mirroring or rotation mappings. I then ask a neural network decoder to predict the bin in which the animal is located in 0.05 second intervals given the neural activity on that interval. **B**. Predicted bins for the position of every data-point in sides A (validation set) and B. Side B in ACC shows predicted bins that match with the rotation mapping. **C**. Correct or incorrect predictions for the position of every data-point in sides A (validation set) and B. While side A is similar between the areas, in side B only ACC for the rotation mapping shows correct predictions across the track. **D**. Accuracy of the network for ACC and HPC on all the bins. The network is trained and validated on side A, and tested on side B. Side A shows similar performance between the areas. However, in side B HPC is substantially worse than ACC. Errors are 95% confidence intervals **E**. Accuracy of the network for ACC and HPC on all the outer part, showing that the prediction holds for the novel space in the side.

seems to be reliably predicted, with a hint of better prediction on the central (vertical) arm for mirroring.

Quantifying this results (Figure 5D), on side A the decoder achieved similar accuracy (ratio of correct versus total bins predicted) for both areas (0.82 for ACC and 0.80 for HPC). On side B, the ACC accuracy was markedly better (0.64 and 0.50 for rotation and mirroring, respectively) than HPC accuracy (0.18 and 0.25 for rotation and mirroring, respectively).

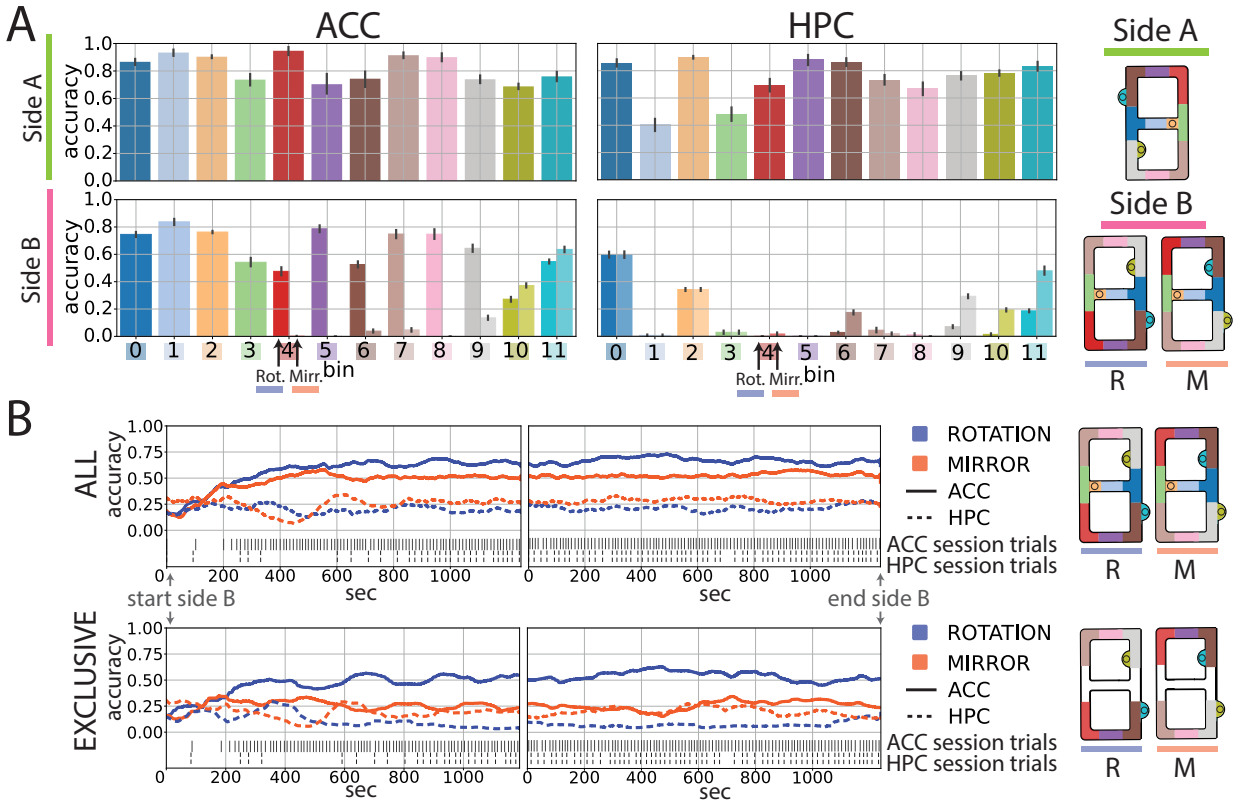

Figure 6: **A**. The predicted bin distribution for each bin. Elements on the diagonal were correctly predicted. Most of the errors in the ACC for side B rotation are in the choice feeders. In HPC side B, the network becomes biased to the feeder bins. **B**. Prediction accuracy in time of side B, for data aligned to the start and end of the time spent in the side. ALL indicates all bins, while EXCLUSIVE indicates that I only consider bins that differ between the rotation and mirroring rotation. For the start aligned accuracy there is a build up of accuracy in ACC, which might indicate that it takes some time for the animal to transfer learn the encoding for side A.

These results hold when only looking at the predicted bin errors on the outer part of the track, which has novel space that the network had not seen in training. The prediction errors for individual bins (Figure 6A, supplementary figure 13B) reveal that, in ACC, the accuracy is similar between the mappings for the choice feeders.

I also looked at how the accuracy changes across time. For this, I compute the mean accuracy over a rolling window of 139sec (a tenth of the length of the side B decoder data), and smooth the resulting curves with a moving average of 15sec. As mentioned in 2.9, I usually align the decoder data to the end of the time in each side. This skips the first few trials that the animal is in the new side, which are usually slower. Doing it this way (Figure

6B, all, right), gives pretty constant accuracy curves. If I align the decoding data to the start of the side B, however, I see that, while the HPC accuracy remains constant, the ACC accuracy increases during the span of ∼400 sec and then plateaus (figure 6B, all, left). This might be an indication that the animal realizes the similarity of the task and starts making use of the learned encoding on the familiar side. Similar plots for side A (validation set) can be found in supplementary figure 9B. Only considering the accuracy of the bins that differ between mappings accentuates the differences between rotation and mirroring (figure 6B, exclusive).

3.5 ACC preserves spatial cluster locations

To get a better idea of how the spatial position is encoded I applied non-linear dimensionality reduction to the neural activity (for details see 2.10). In particular, I find a 2 dimensional mapping for side A and use it to map the data from side A and from side B. In a similar manner, I also reduce to 2 dimensions the activity of the second to last layer in my neural network decoder, in order to improve bin separability and to gain insight into how the neural network decoder is transforming the data.

Then, I assign a spatial bin to each one of the data-points that are projected in the plane. Figure 7 shows the neuronal activities for the neural data (ORIGINAL) and the neural network processed data (NETWORK) on the reduced dimensional space. For clarity, I am only taking the bins that are exclusive between mirroring and mapping rotations (see supplementary figure 10 for all bins and feeder bins). I can see that the bins are clearly separated in the side A of ACC, and the network is just clustering the data points into more compacted areas. In the side A of HPC there is more mixing of the different bins in the original data, which the network is successfully separating. In side B of ACC the shape of the dynamics is visibly preserved, and the rotation mapping correctly maps the bins to the same relative positions than in side A. This is also the case in the network processed data. The choice feeder bins seem to get intermixed in both original and network processed

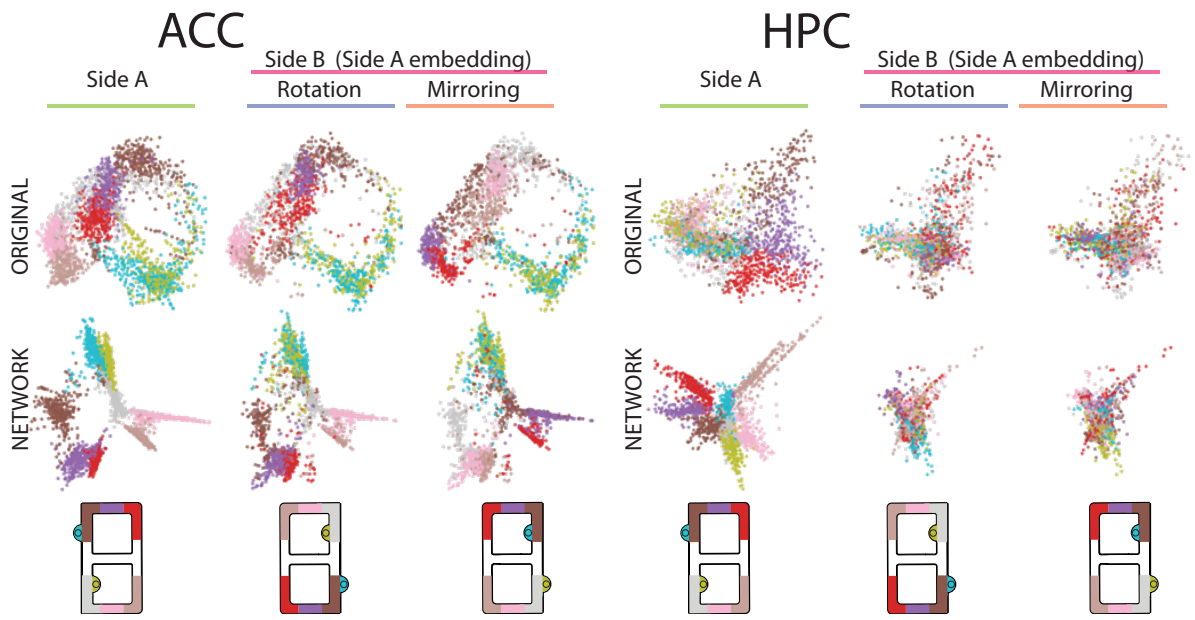

Figure 7: Neural activities plotted on the plane using nonlinear dimensionality reduction (ISOMAP), for the original and network processed activities. Each activity datapoint is labeled with the bin where it was produced. Bin clusters mixed in the original data are better separated after going trough the network. In ACC, the shape of the manifold is preserved between sides, and the bin labels match the rotation mapping. In HPC, the shapes are worse preserved and I lose the cluster locations of the different bins.

data, which aligns with the error profile discussed in figure 5A. In HPC, in both original and network processed data the shape of the dynamics is less recognisable between the sides, and the clusters are mixed in both rotation and mirroring.

This provides further indication that the spatial encoding in ACC is preserved between the sides (in a rotation mapping), while the encoding in HPC gets mixed.

4 Discussion

I have shown that spatial firing patterns and trial dynamics across the track in ACC are largely preserved when going from a familiar to a novel side in a double figure-8 maze. Furthermore, this similarity is sufficient to coarsely decode the position of the animal with a decoder trained on the familiar side. This provides evidence the ACC contains a cognitive map of "task space" [45], an abstraction of the task disassociated from the physical space. I have also shown that in HPC, which is known for having a map of the physical space, the encoding largely changes when going from one side to the other.

A caveat of my results is the limited number of animals and the behavioral differences between them. Ideally, rats with electrodes in both ACC and HPC should be used to control for behavioral divergence. Nevertheless, replicating these results with more rats is necessary to support the soundness of these claims.

This proposal extends others characterizing ACC as integrating sensory with memory aspects [44] to form schemas to map context and events onto actions [55]. In this view, ACC ensemble dynamics live in a high dimensional task space, sitting on a learned manifold, a schema that allows for rapid assimilation of new information. A trajectory through the maze corresponds to a trajectory in this task space, which can act as a model to generate predictions among stimuli, actions, and reward [56] [15].

My results are also in accordance with previous studies that have found spatial or contextual information in ACC [13][14][15] . While in [14] ACC changes encoding following changes in environmental contexts, which might conflict with these results, the author's experimental paradigm was not task focused. This would provide support to the view of task space as effectively shared between instances of the same (or similar) tasks, and not between random

environments.

Most of the measures used here were spatial in nature. Only with the TCA analysis (3.3), I showed some preservation of a limited range of task variables across sides (reward/non reward and forced/free). Therefore, it will be interesting to see if the shared cognitive map can be used to decode other task relevant variables, such as the current behavioral block or the choice of the animal in free trials. Unfortunately, the limited number of trials precluded reaching conclusions about whether the encoding of these variables is preserved between the sides.

Is ACC encoding space in any meaningful way? Because of the stereotyped nature of the task, what I identified as spatial encoding might instead be a series of patterned activity signaling task related events. To the extent in which this description is isomorphic to an abstract map of the task, the rat might be still navigating in "task space", but task relevant information might be tied to cues, timing and events instead of locations in the task layout. Further research is needed to elucidate the "spatialness" of this task space.

That ACC preserves the spatial encoding while HPC remaps might seem surprising, giving their strong connectivity [42]. While I have not shown this disassociation explicitly, as I did not have concurrent ACC and HPC recordings, it seems plausible given what is known about HPC and remapping. It has been suggested that interplay between ACC and HPC seems necessary when new learning overlaps or contradicts associations in the schemas [57], so this suggestive disassociation might be due the similarity of the familiar and novel tasks. Including a third phase in which the layout or task strategy change might be a way to flesh out the mechanisms by which ACC and HPC interact.

These results are also related to transfer of learning. Transfer of learning is the transfer of domain knowledge to novel tasks and contexts, and is a hallmark of animal learning [58]. In the brain, transfer of learning is associated with the prefrontal cortex. Studies involving lesions of the PFC in rats show a decline in animals' performance when faced with novel variants of a known maze [59][60]. In humans, administration of a dopamine

D_2 receptor agonist led to an increase of activity in the ventromedial PFC (vmPFC) and concordant better performance in a transfer of learning task [61]. Another human study found that individual differences in transfer of learning were related to differences in vmPFC-hippocampal connectivity [62].

Is the ACC doing transfer of learning by transferring the knowledge from the familiar side to the novel side in our task? Concordant with this hypothesis, the ACC ensemble activity similarity is present when the animal moves to the novel side. However, the similarity of encoding might just be a consequence of the similarity between the task on both sides. Because of this similarity and the lack of behavioral correlates showing leverage of the knowledge learnt on the familiar side I cannot claim that the animal is performing actual transfer of learning.

Computationally, cognitive maps suggest a form of model-based reinforcement learning, where the animal learns a model of the world that allows for planning future actions. Here, a model is anything that the animal can use to predict how the environment will respond to its actions [63]. Formally, it is usually represented as an estimator of the transition probabilities, i.e., what is the probability that doing an action leads me to a given state. Thus, while cognitive maps are clear examples of a model of the environment, they are not the only possible way to embody models in the brain. Model based reinforcement learning is very sample efficient, and allows for rapid strategy switching when the environment changes [63]. Models of the environment have been linked to motivated and goal directed behavior, and associated to frontal brain areas, mainly OFC [64]. On the other hand, habitual behavior, in which the animal responds to a series of stimulus action associations, has traditionally been linked to dopamine signals [65]. However, there is now evidence that even dopamine reward signals contain model based information [66]. In fact, model-based signals seem to be ubiquitous wherever the brain processes rewards [67], indicating that there may not a clear separation in the brain between habitual and goal directed mechanisms. Showing that the ACC is a full fledged model of the environment tips the scale towards model based learning

in the brain.

Understanding how the brain learns a task, how the task knowledge is represented and how it might be transferred to other contexts or variants might shed new light into the neural mechanisms of animal cognition. Artificial deep neural networks have shown incredible potential in multiple domains such as image classification, language modeling or reinforcement learning [68]. However, they are often restricted to particular domains, perform poorly on data distribution shifts [69][70] and require massive amounts of data to train [71]. Designing systems that surpass these limitations is a challenging problem [72]. On the opposite side, humans (and animals in some cases) seem to effortlessly generalize to novel cases, and can learn concepts from one single example case. Thus, understanding how the brain can flexibly represent a task in new contexts and situations will not only help discover fundamentally important processes of cognition in animals, it may very well provide a blueprint (or at least inspiration) to develop artificial systems with vastly superior capabilities to deal with situations outside of the training set.

Bibliography

[1] Edward C Tolman. Cognitive maps in rats and men. *Psychological review*, 55(4):189, 1948.

[2] John O'Keefe and Jonathan Dostrovsky. The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain research*, 1971.

[3] John O'keefe and Lynn Nadel. *The hippocampus as a cognitive map*. Oxford: Clarendon Press, 1978.

[4] James B Ranck Jr. Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats: Part i. behavioral correlates and firing repertoires. *Experimental neurology*, 41(2):462–531, 1973.

[5] Robert E Hampson, John D Simeral, and Sam A Deadwyler. Distribution of spatial and nonspatial information in dorsal hippocampus. *Nature*, 402(6762):610–614, 1999.

[6] Emma R Wood, Paul A Dudchenko, and Howard Eichenbaum. The global record of memory in hippocampal neuronal activity. *Nature*, 397(6720):613–616, 1999.

[7] Sarah J Cohen, Alcira H Munchow, Lisa M Rios, Gongliang Zhang, Herborg N Ásgeirsdóttir, and Robert W Stackman Jr. The rodent hippocampus is essential for nonspatial object memory. *Current Biology*, 23(17):1685–1690, 2013.

[8] Philip Nicholas Johnson-Laird. *Mental models: Towards a cognitive science of language, inference, and consciousness*. Number 6. Harvard University Press, 1983.

[9] Adam Johnson and A David Redish. Neural ensembles in ca3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience*, 27(45):12176–12189, 2007.

[10] Matthijs AA van der Meer, Adam Johnson, Neil C Schmitzer-Torbert, and A David Redish. Triple dissociation of information processing in dorsal striatum, ventral striatum, and hippocampus on a learned spatial decision task. *Neuron*, 67(1):25–32, 2010.

[11] Matthew FS Rushworth, MaryAnn P Noonan, Erie D Boorman, Mark E Walton, and Timothy E Behrens. Frontal cortex and reward-guided learning and decision-making. *Neuron*, 70(6):1054–1069, 2011.

[12] Aaron J Gruber, Gwendolyn G Calhoon, Igor Shusterman, Geoffrey Schoenbaum, Matthew R Roesch, and Patricio O'Donnell. More is less: a disinhibited prefrontal cortex impairs cognitive flexibility. *Journal of Neuroscience*, 30(50):17102–17110, 2010.

[13] Shigeyoshi Fujisawa, Asohan Amarasingham, Matthew T Harrison, and György Buzsáki. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature neuroscience*, 11(7):823, 2008.

[14] James M Hyman, Liya Ma, Emili Balaguer-Ballester, Daniel Durstewitz, and Jeremy K Seamans. Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proceedings of the National Academy of Sciences*, 109(13):5086–5091, 2012.

[15] Ali Mashhoori, Saeedeh Hashemnia, Bruce L McNaughton, David R Euston, and Aaron J Gruber. Rat anterior cingulate cortex recalls features of remote reward locations after disfavoured reinforcements. *Elife*, 7:e29793, 2018.

[16] Madoka Matsumoto, Kenji Matsumoto, Hiroshi Abe, and Keiji Tanaka. Medial prefrontal cell activity signaling prediction errors of action values. *Nature neuroscience*, 10(5):647–656, 2007.

[17] David R Euston and Bruce L McNaughton. Apparent encoding of sequential context in rat medial prefrontal cortex is accounted for by behavioral variability. *Journal of Neuroscience*, 26(51):13143–13155, 2006.

[18] Steven W Kennerley, Timothy EJ Behrens, and Jonathan D Wallis. Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nature neuroscience*, 14(12):1581, 2011.

[19] Matthew M Walsh and John R Anderson. Learning from experience: event-related potential correlates of reward processing, neural adaptation, and behavioral choice. *Neuroscience & Biobehavioral Reviews*, 36(8):1870–1884, 2012.

[20] William Beecher Scoville and Brenda Milner. Loss of recent memory after bilateral hippocampal lesions. *Journal of neurology, neurosurgery, and psychiatry*, 20(1):11, 1957.

[21] Larry R Squire. Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychological review*, 99(2):195, 1992.

[22] John Patrick Aggleton and Malcolm W Brown. Episodic memory, amnesia and the hippocampal-anterior thalamic axis. *Behavioral and brain sciences*, 22(3):425–444, 1999.

[23] Robin MA Hayman, Subhojit Chakraborty, Michael I Anderson, and Kathryn J Jeffery. Context-specific acquisition of location discrimination by hippocampal place cells. *European Journal of Neuroscience*, 18(10):2825–2834, 2003.

[24] Robert U Muller, John L Kubie, EM Bostock, JS Taube, and GJ Quirk. Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. *Brain and space*, 1991.

[25] Stefan Leutgeb, Jill K Leutgeb, Carol A Barnes, Edvard I Moser, Bruce L McNaughton, and May-Britt Moser. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science*, 309(5734):619–623, 2005.

[26] Laura Lee Colgin, Edvard I Moser, and May-Britt Moser. Understanding memory through hippocampal remapping. *Trends in neurosciences*, 31(9):469–477, 2008.

[27] Massimo Silvetti, William Alexander, Tom Verguts, and Joshua W Brown. From conflict management to reward-based decision making: actors and critics in primate medial frontal cortex. *Neuroscience & Biobehavioral Reviews*, 46:44–57, 2014.

[28] Amitai Shenhav, Matthew M Botvinick, and Jonathan D Cohen. The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron*, 79(2):217–240, 2013.

[29] Tommy C Blanchard and Benjamin Y Hayden. Neurons in dorsal anterior cingulate cortex signal postdecisional variables in a foraging task. *Journal of Neuroscience*, 34(2):646–655, 2014.

[30] Amit Etkin, Tobias Egner, and Raffael Kalisch. Emotional processing in anterior cingulate and medial prefrontal cortex. *Trends in cognitive sciences*, 15(2):85–93, 2011.

[31] Mark E Walton, David M Bannerman, Karin Alterescu, and Matthew FS Rushworth. Functional specialization within medial frontal cortex of the anterior cingulate for evaluating effort-related decisions. *Journal of Neuroscience*, 23(16):6475–6479, 2003.

[32] Ziv M Williams, George Bush, Scott L Rauch, G Rees Cosgrove, and Emad N Eskandar. Human anterior cingulate neurons and the integration of monetary reward with motor responses. *Nature neuroscience*, 7(12):1370–1375, 2004.

[33] Yogita Chudasama, Teresa E Daniels, Daniel P Gorrin, Sarah EV Rhodes, Peter H Rudebeck, and Elisabeth A Murray. The role of the anterior cingulate cortex in choices based on reward value and reward contingency. *Cerebral Cortex*, 23(12):2884–2898, 2013.

[34] Clay B Holroyd and Samuel M McClure. Hierarchical control over effortful behavior by rodent medial frontal cortex: A computational model. *Psychological review*, 122(1):54, 2015.

[35] Steven W Kennerley, Mark E Walton, Timothy EJ Behrens, Mark J Buckley, and Matthew FS Rushworth. Optimal decision making and the anterior cingulate cortex. *Nature neuroscience*, 9(7):940–947, 2006.

[36] Irma Triasih Kurniawan, Ben Seymour, Deborah Talmi, Wako Yoshida, Nick Chater, and Raymond J Dolan. Choosing to make an effort: the role of striatum in signaling physical effort of a chosen action. *Journal of neurophysiology*, 104(1):313–321, 2010.

[37] Josef Parvizi, Vinitha Rangarajan, William R Shirer, Nikita Desai, and Michael D Greicius. The will to persevere induced by electrical stimulation of the human cingulate gyrus. *Neuron*, 80(6):1359–1367, 2013.

[38] Szu-Han Wang, Dorothy Tse, and Richard GM Morris. Anterior cingulate cortex in schema assimilation and expression. *Learning & Memory*, 19(8):315–318, 2012.

[39] David R Euston, Aaron J Gruber, and Bruce L McNaughton. The role of medial prefrontal cortex in memory and decision making. *Neuron*, 76(6):1057–1070, 2012.

[40] James L McClelland, Bruce L McNaughton, and Randall C O'Reilly. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419, 1995.

[41] Gordon Winocur, Morris Moscovitch, and Bruno Bontempi. Memory formation and long-term retention in humans and animals: Convergence towards a transformation account of hippocampal–neocortical interactions. *Neuropsychologia*, 48(8):2339–2356, 2010.

[42] Priyamvada Rajasethupathy, Sethuraman Sankaran, James H Marshel, Christina K Kim, Emily Ferenczi, Soo Yeun Lee, Andre Berndt, Charu Ramakrishnan, Anna Jaffe, Maisie Lo, et al. Projections from neocortex mediate top-down control of memory retrieval. *Nature*, 526(7575):653–659, 2015.

[43] Dorothy Tse, Rosamund F Langston, Masaki Kakeyama, Ingrid Bethus, Patrick A Spooner, Emma R Wood, Menno P Witter, and Richard GM Morris. Schemas and memory consolidation. *Science*, 316(5821):76–82, 2007.

[44] Christopher C Lapish, Daniel Durstewitz, L Judson Chandler, and Jeremy K Seamans. Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, 105(33):11963–11968, 2008.

[45] Sarah R Heilbronner and Benjamin Y Hayden. Dorsal anterior cingulate cortex: a bottom-up view. *Annual review of neuroscience*, 39:149–170, 2016.

[46] Robert C Wilson, Yuji K Takahashi, Geoffrey Schoenbaum, and Yael Niv. Orbitofrontal cortex as a cognitive map of task space. *Neuron*, 81(2):267–279, 2014.

[47] Luca Scrucca, Michael Fop, T. Brendan Murphy, and Adrian E. Raftery. mclust5: clustering, classification and density estimation using Gaussian finite mixture models. *The R Journal*, 8(1):289–317, 2016.

[48] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

[49] Pierre Dragicevic. Fair statistical communication in hci. In *Modern statistical methods for HCI*, pages 291–330. Springer, 2016.

[50] Jean Kossaifi, Yannis Panagakis, Anima Anandkumar, and Maja Pantic. Tensorly: Tensor learning in python. *Journal of Machine Learning Research*, 20(26):1–6, 2019.

[51] Amnon Shashua and Tamir Hazan. Non-negative tensor factorization with applications to statistics and computer vision. In *Proceedings of the 22nd international conference on Machine learning*, pages 792–799, 2005.

[52] Joshua B Tenenbaum, Vin De Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.

[53] Alex H Williams, Tony Hyun Kim, Forea Wang, Saurabh Vyas, Stephen I Ryu, Krishna V Shenoy, Mark Schnitzer, Tamara G Kolda, and Surya Ganguli. Unsupervised discovery of demixed, low-dimensional neural dynamics across multiple timescales through tensor component analysis. *Neuron*, 98(6):1099–1115, 2018.

[54] Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.

[55] William H Alexander and Joshua W Brown. Medial prefrontal cortex as an action-outcome predictor. *Nature neuroscience*, 14(10):1338–1344, 2011.

[56] Erin L Rich and Matthew Shapiro. Rat prefrontal cortical neurons selectively code strategy switches. *Journal of Neuroscience*, 29(22):7208–7219, 2009.

[57] Alison R Preston and Howard Eichenbaum. Interplay of hippocampus and prefrontal cortex in memory. *Current Biology*, 23(17):R764–R773, 2013.

[58] Sebastian Thrun and Lorien Pratt. *Learning to learn*. Springer Science & Business Media, 2012.

[59] Gordon Winocur and Morris Moscovitch. Hippocampal and prefrontal cortex contributions to learning and memory: analysis of lesion and aging effects on maze learning in rats. *Behavioral neuroscience*, 104(4):544, 1990.

[60] Bryan Kolb, Kristin Buhrmann, Robert McDonald, and Robert J. Sutherland. Dissociation of the Medial Prefrontal, Posterior Parietal, and Posterior Temporal Cortex for Spatial Navigation and Recognition Memory in the Rat. *Cerebral Cortex*, 4(6):664–680, 11 1994.

[61] Gerhard Jocham, Tilmann A Klein, and Markus Ullsperger. Dopamine-mediated reinforcement learning signals in the striatum and ventromedial prefrontal cortex underlie value-based choices. *Journal of Neuroscience*, 31(5):1606–1613, 2011.

[62] Raphael T Gerraty, Juliet Y Davidow, G Elliott Wimmer, Itamar Kahn, and Daphna Shohamy. Transfer of learning relates to intrinsic connectivity between hippocampus, ventromedial prefrontal cortex, and large-scale networks. *Journal of Neuroscience*, 34(34):11297–11303, 2014.

[63] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[64] Vivian V Valentin, Anthony Dickinson, and John P O'Doherty. Determining the neural substrates of goal-directed learning in the human brain. *Journal of Neuroscience*, 27(15):4019–4026, 2007.

[65] Nathaniel D Daw and Daphna Shohamy. The cognitive neuroscience of motivation and learning. *Social Cognition*, 26(5):593–620, 2008.

[66] Clara Kwon Starkweather, Benedicte M Babayan, Naoshige Uchida, and Samuel J Gershman. Dopamine reward prediction errors reflect hidden-state inference across time. *Nature neuroscience*, 20(4):581–589, 2017.

[67] Bradley B Doll, Dylan A Simon, and Nathaniel D Daw. The ubiquity of model-based reinforcement learning. *Current opinion in neurobiology*, 22(6):1075–1081, 2012.

[68] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.

[69] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 427–436, 2015.

[70] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.

[71] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40, 2017.

[72] Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dharshan Kumaran, and Matt Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.

A Appendix Supplementary Figures

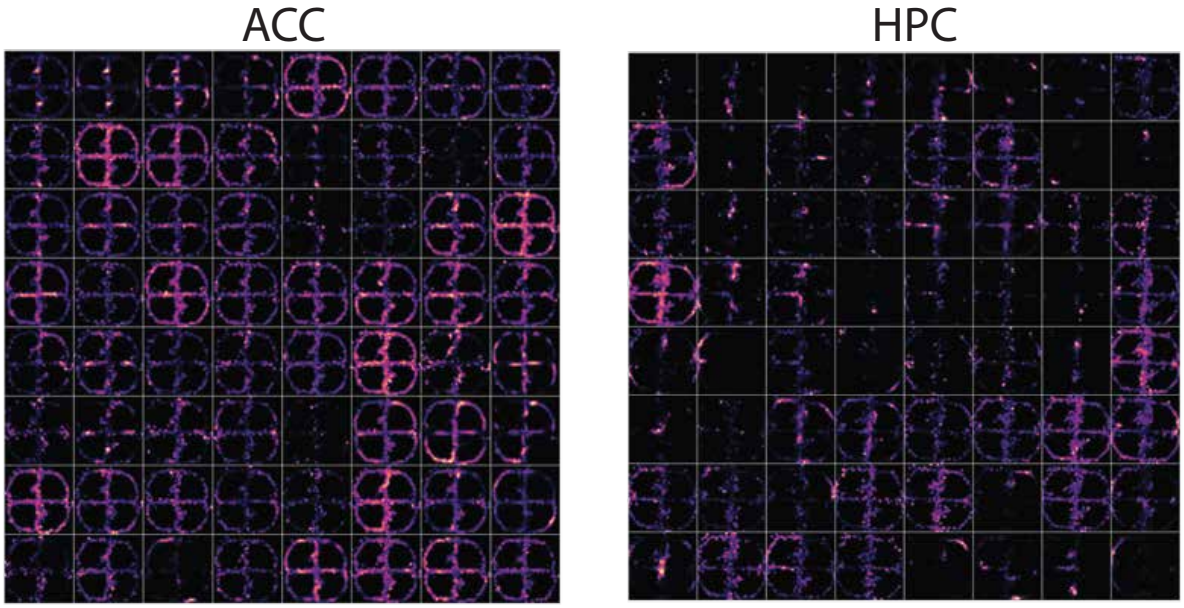

Figure S1: Average firing rates (spikes/sec) of 64 ACC and HPC neurons across the track. Most ACC neurons are active across the track, while HPC, as expected, has a high percentage of neurons with a narrow receptive field.

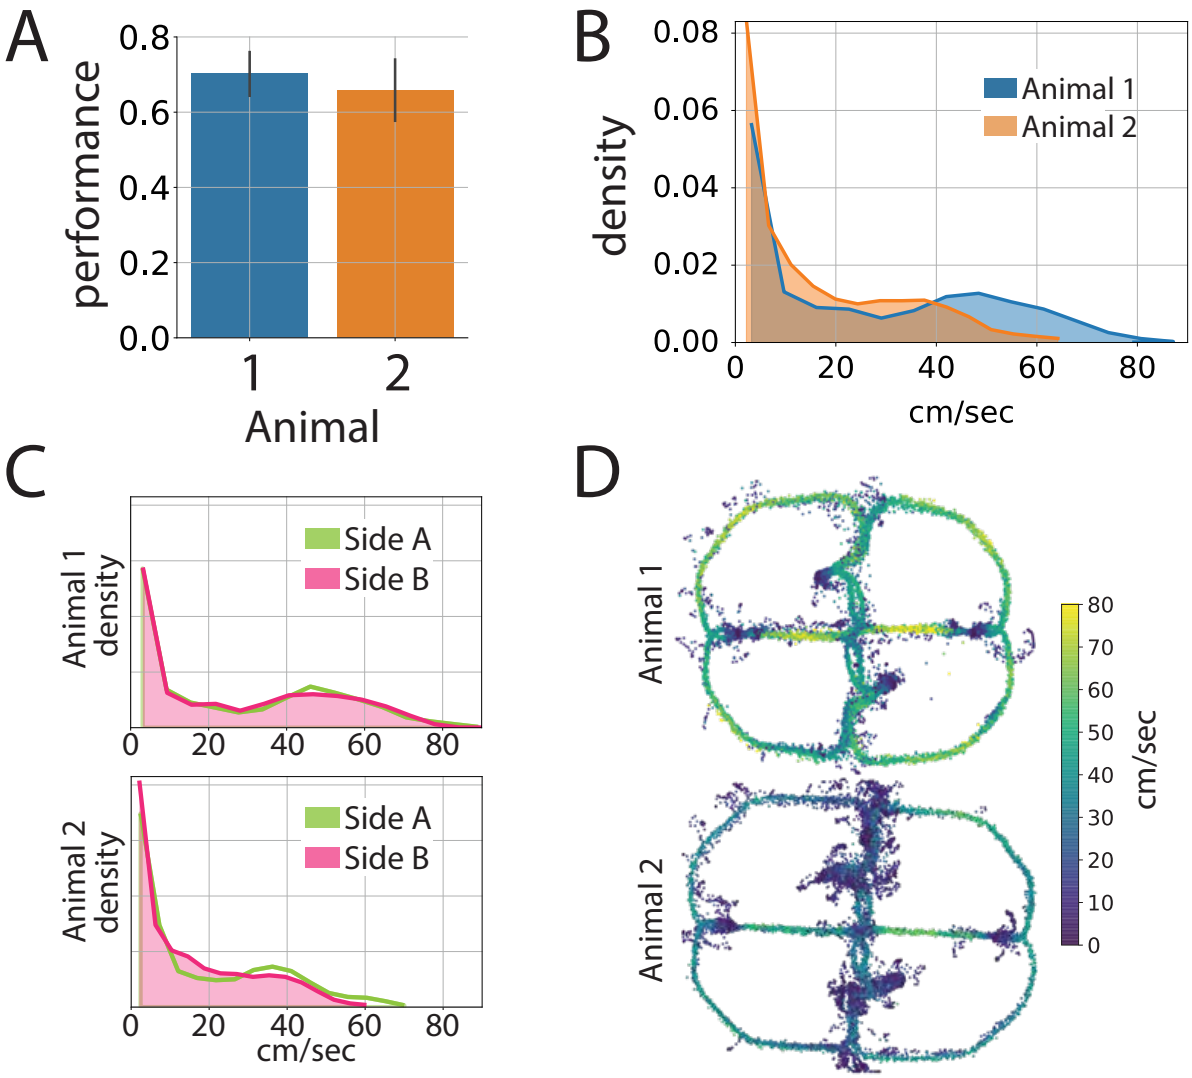

Figure S2: **A**. Performance of the animal, defined as the ratio of free trials where the animal choose the optimal feeder. The two animals show similar performance. Errors are the 95% confidence interval. **B**. Histograms for the speed of the animals. Animal 2 was consistently slower and stopped more frequently than animal 1. **C**. Distribution of speed between the two sides. **D**. The speed of the animals across the track shows that the animals tend to speed up at the corners and slow down at the intersections.

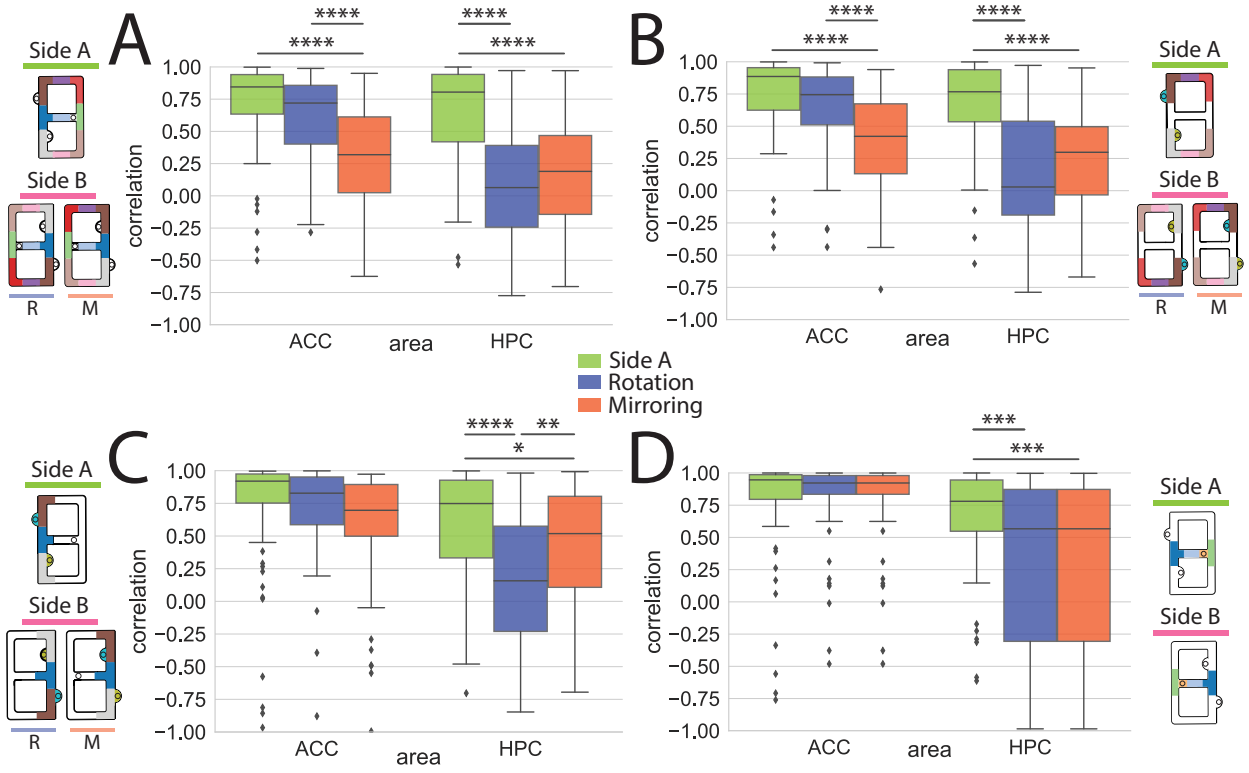

Figure S3: **A**. Spatial firing pattern correlation between sides ignoring feeders (pval 4.92×10^{-9} and 6.61×10^{-9} between side A and mirror for ACC and HPC, pval 0.06 and 5.35×10^{-10} between side A and rotation for ACC and HPC , pval 0.32 and 7.96×10^{-6} between rotation and mirroring for ACC and HPC, Welch's t-test) **B**. Spatial firing pattern correlation between sides taking only bins that change between rotation and mirror mapping. (pval 0.06 and 0.04 between side A and mirror for ACC and HPC, pval 0.83 and 2.20×10^{-6} between side A and rotation for ACC and HPC , pval 1.28×10^{-3} and 2.37×10^{-3} between rotation and mirroring for ACC and HPC, Welch's t-test) **C**. Spatial firing pattern correlation on the shared space between the sides (pval 2.54×10^{-8} and 7.39×10^{-8} between side A and mirror for ACC and HPC, pval 7.54×10^{-3} and 1.31×10^{-10} between side A and rotation for ACC and HPC , pval 3.57×10^{-5} and 7.28×10^{-3} between rotation and mirroring for ACC and HPC, Welch's t-test). **D**. Spatial firing pattern correlation between sides taking only bins that do not change between rotation and mirror mapping. (pval 0.56 and 6.20×10^{-4} between side A and mirror/rotation for ACC and HPC, Welch's t-test).

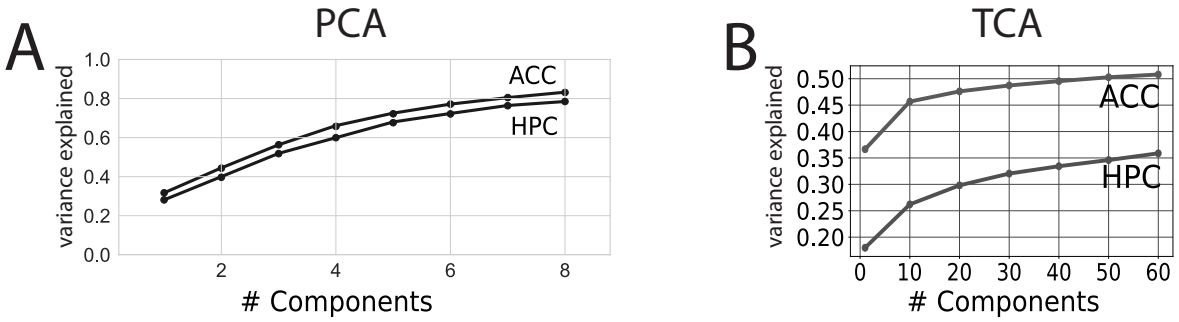

Figure S4: **A**. Typical cumulative variance explained of PCA as a function of the number of components. **B**. Cumulative variance explained for different number of components when decomposing the neural data into a non negative tensor decomposition with that number of components. The ACC variance explained is substantially higher than HPC's. As this is not the case in the PCA, which only has to explain one side variance, the higher gap in TCA might be due to the HPC change in encoding between the sides.

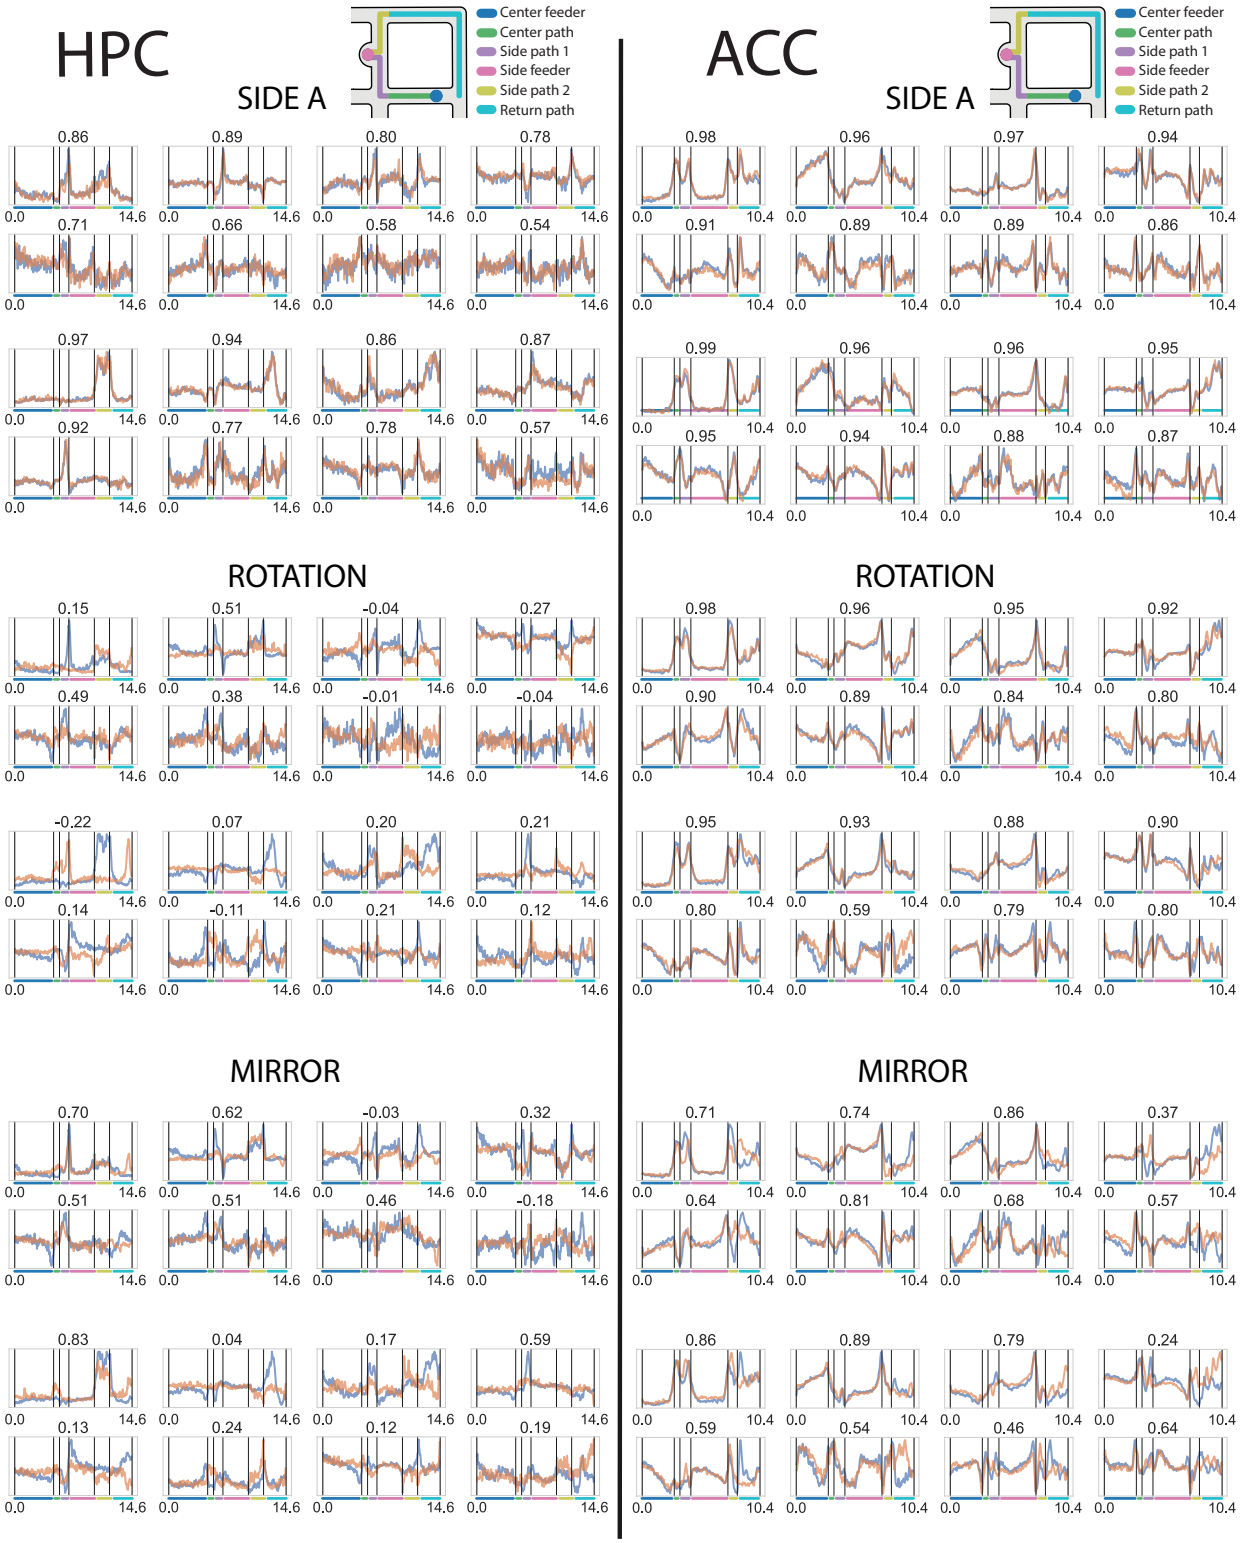

Figure S5: PCA components between the first and second half of side A (side A), between the track sides for contralateral loops (rotation) and parallel loops (mirror). Plots show specificity to certain parts of the track. The number above each component pair indicates the correlation between the components of side A and the activity of side B projected using side A's loadings.

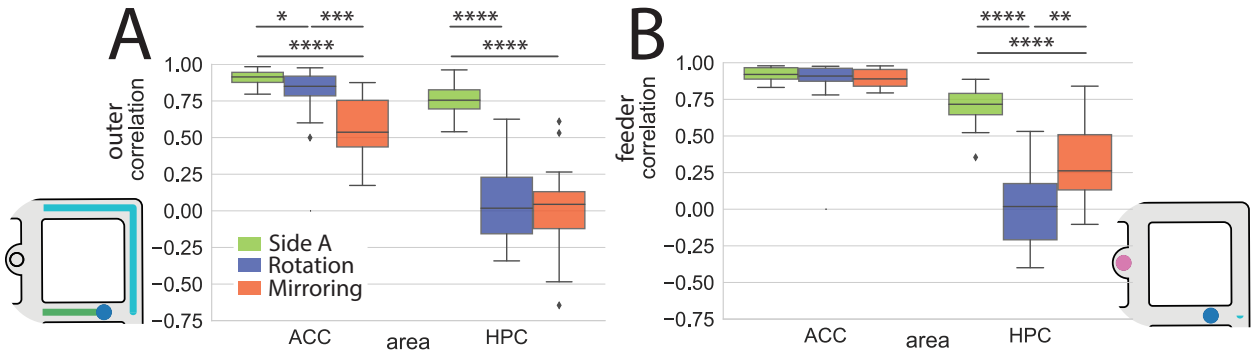

Figure S6: **A**. Correlation of the first 8 principal components between sides for the outer region of the track (pval 1.03×10^{-5} and 5.68×10^{-8} between side A and mirror for ACC and HPC, pval 2.93×10^{-2} and 2.27×10^{-8} between side A and rotation for ACC and HPC , pval 0.39×10^{-4} and 0.65 between rotation and mirroring for ACC and HPC, Welch's t-test) **B**. Correlation of the first 8 principal components between sides for the feeders (pval 0.34 and 6.86×10^{-5} between side A and mirror for ACC and HPC, pval 0.59 and 1.69×10^{-8} between side A and rotation for ACC and HPC , pval 0.69 and 7.02×10^{-3} between rotation and mirroring for ACC and HPC, Welch's t-test).

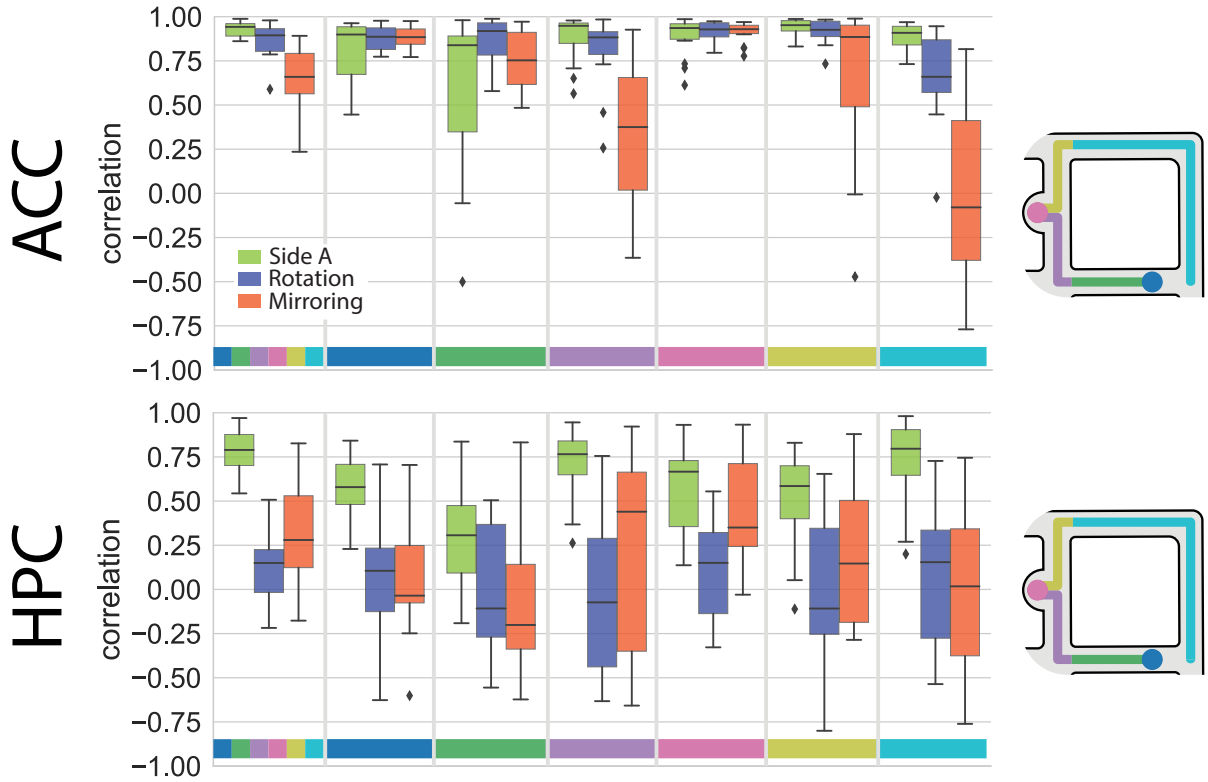

Figure S7: **A**. Similarity of the first 8 principal components for different parts of the track, between first and second halves of side a (side A), or between side A and B parallel loops (mirroring) or contralateral loops (rotation).

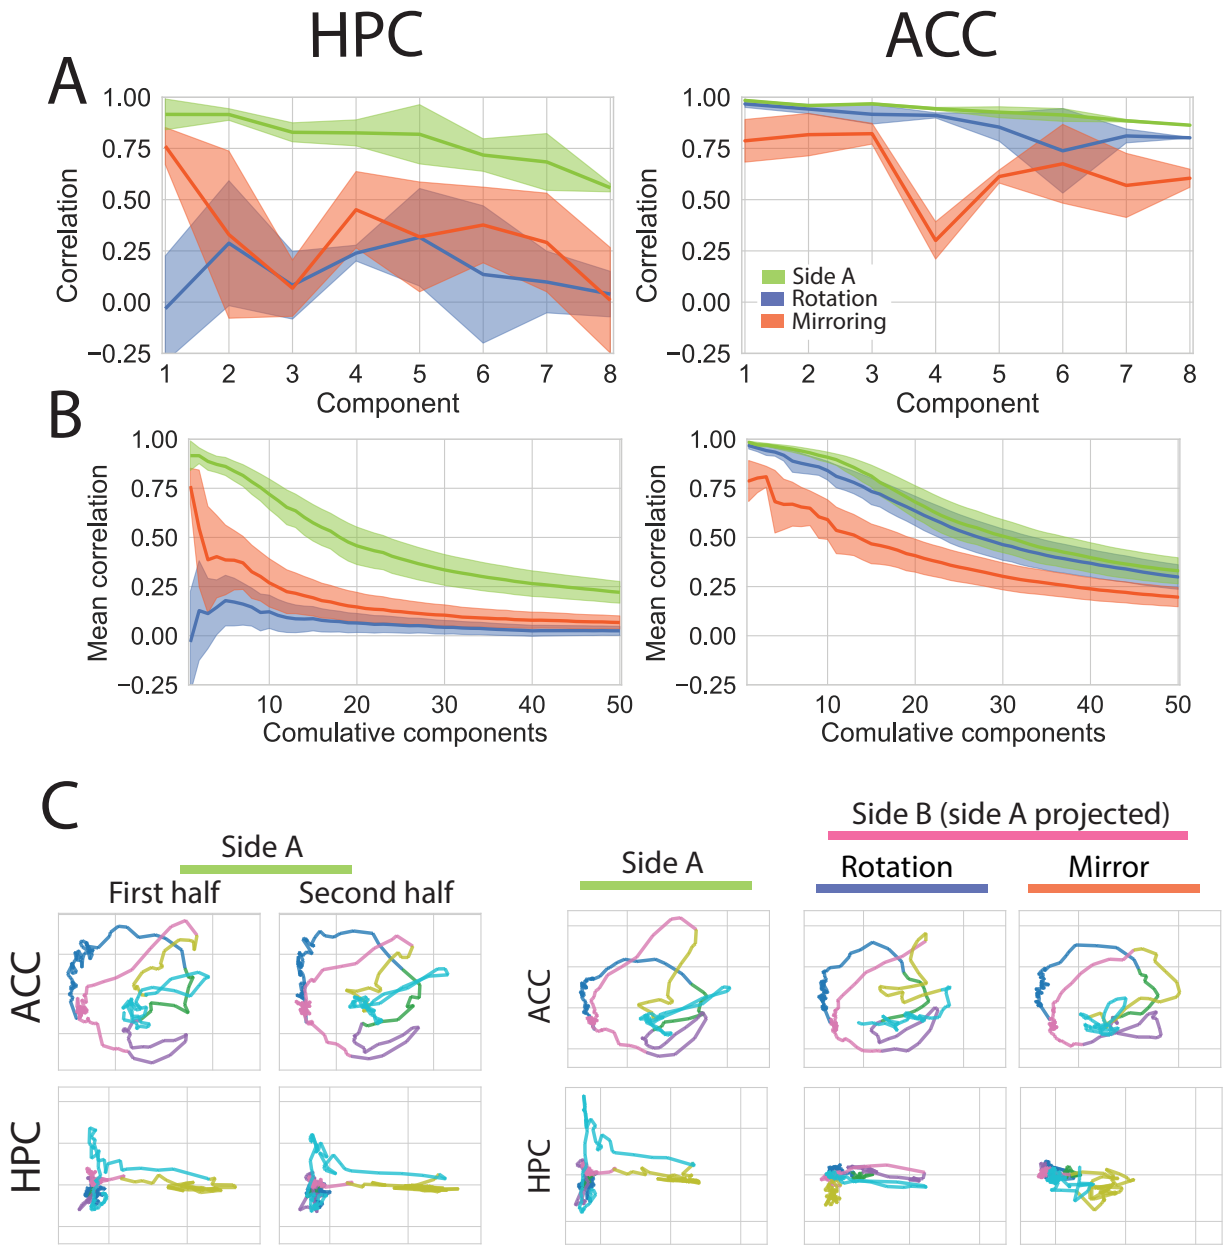

Figure S8: **A**. Pearson correlation between the components of side A and side B (projected using side A loadings), for the first 8 components, and the first and second halves of side A for control. ACC trial dynamics are more similar than HPC. For ACC, rotation is more similar than mirroring. The opposite is true for HPC. Errors show 95% confidence intervals. **B**. The average correlation between the components for side A and side B (using side A projection) and the first and second halves of side A for control, for an increasing number of components. Comparing contralateral loops, corresponding to a rotation mapping between the sides, is closer than comparing parallel loops, corresponding to mirror mapping in the ACC. Both, however, are higher than the correlations in HPC, where mirror mapping beats rotation. All the curves decrease as I average over more components, as latter components encode directions of less variance and are more susceptible to noise in the neural data. **C**. The two first components in the plane.

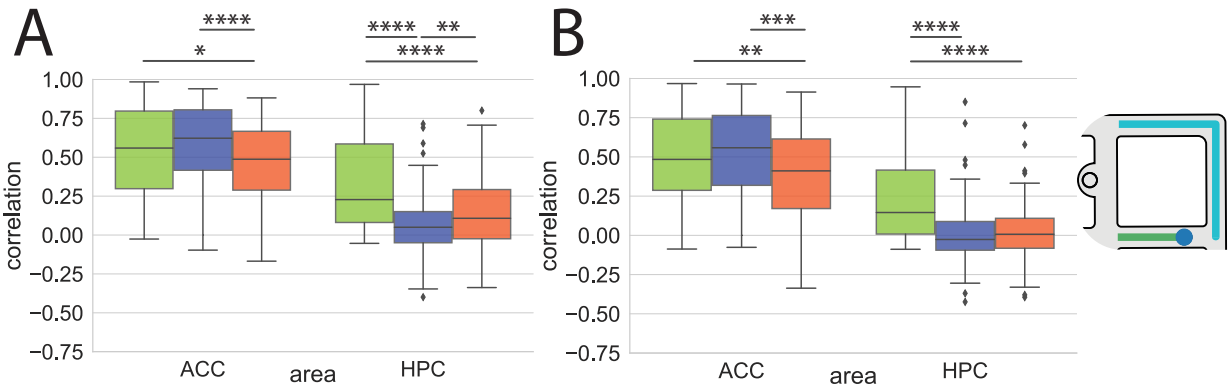

Figure S9: **A**. Correlations between individual neuron trace activities between the two halves of side A, and between sides A and B for rotation (comparing contralateral loops) and mirroring (comparing parallel loops). A higher percentage of neurons preserve trial dynamics, specially in rotation (pval 0.34 and 1.44×10^{-7} between side A and mirror for ACC and HPC, pval 0.03 and 6.75×10^{-15} between side A and rotation for ACC and HPC , pval 7.70×10^{-5} and 2.70×10^{-3} between rotation and mirroring for ACC and HPC, Welch's t-test) **B**. The above pattern holds when considering only the space that is not shared between sides.(pval 8.49×10^{-3} and 1.65×10^{-13} between side A and mirror for ACC and HPC, pval 0.27 and 6.19×10^{-14} between side A and rotation for ACC and HPC , pval 1.88×10^{-4} and 0.66 between rotation and mirroring for ACC and HPC, Welch's t-test).

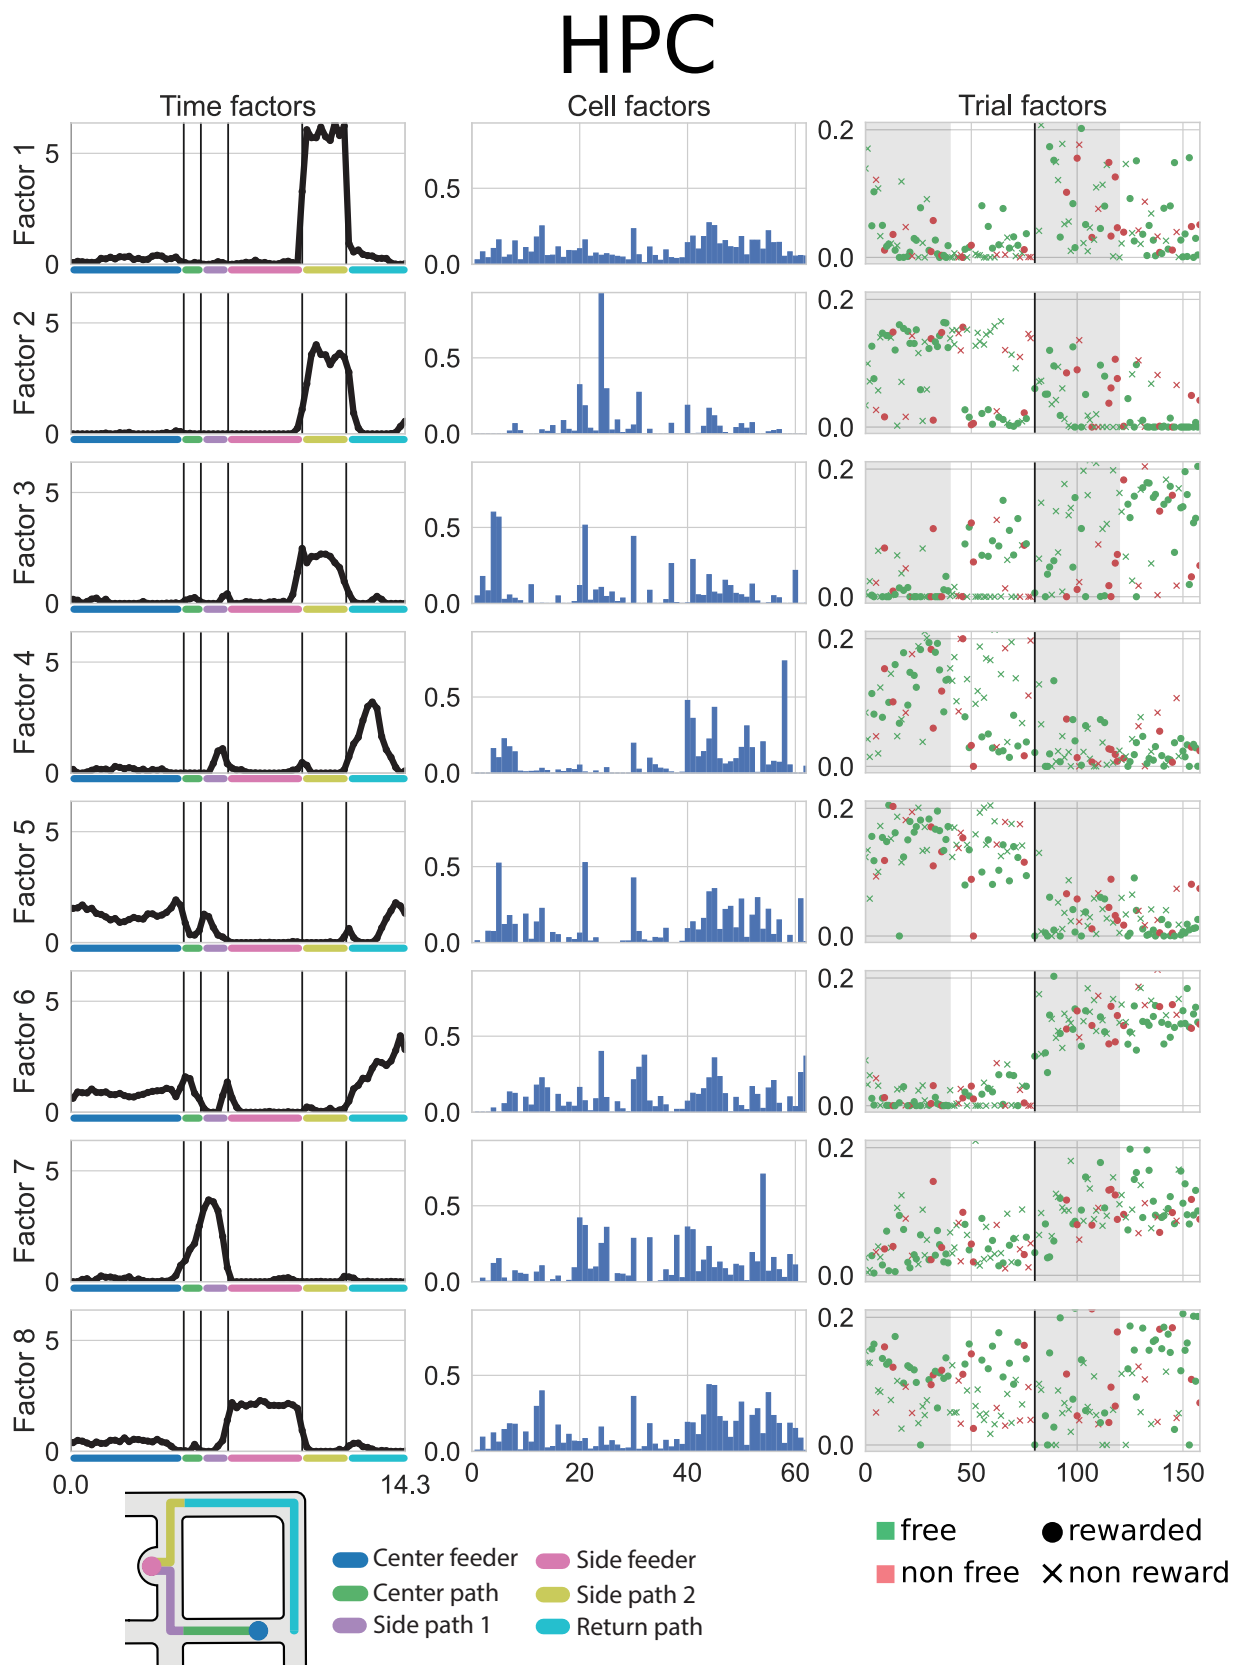

Figure S10: Complete list of factors for the TCA decomposition (see figure 4B for more details) for HPC. Note the discontinuities in the trial factors, indicating a drastic change in the spatial encoding.

.

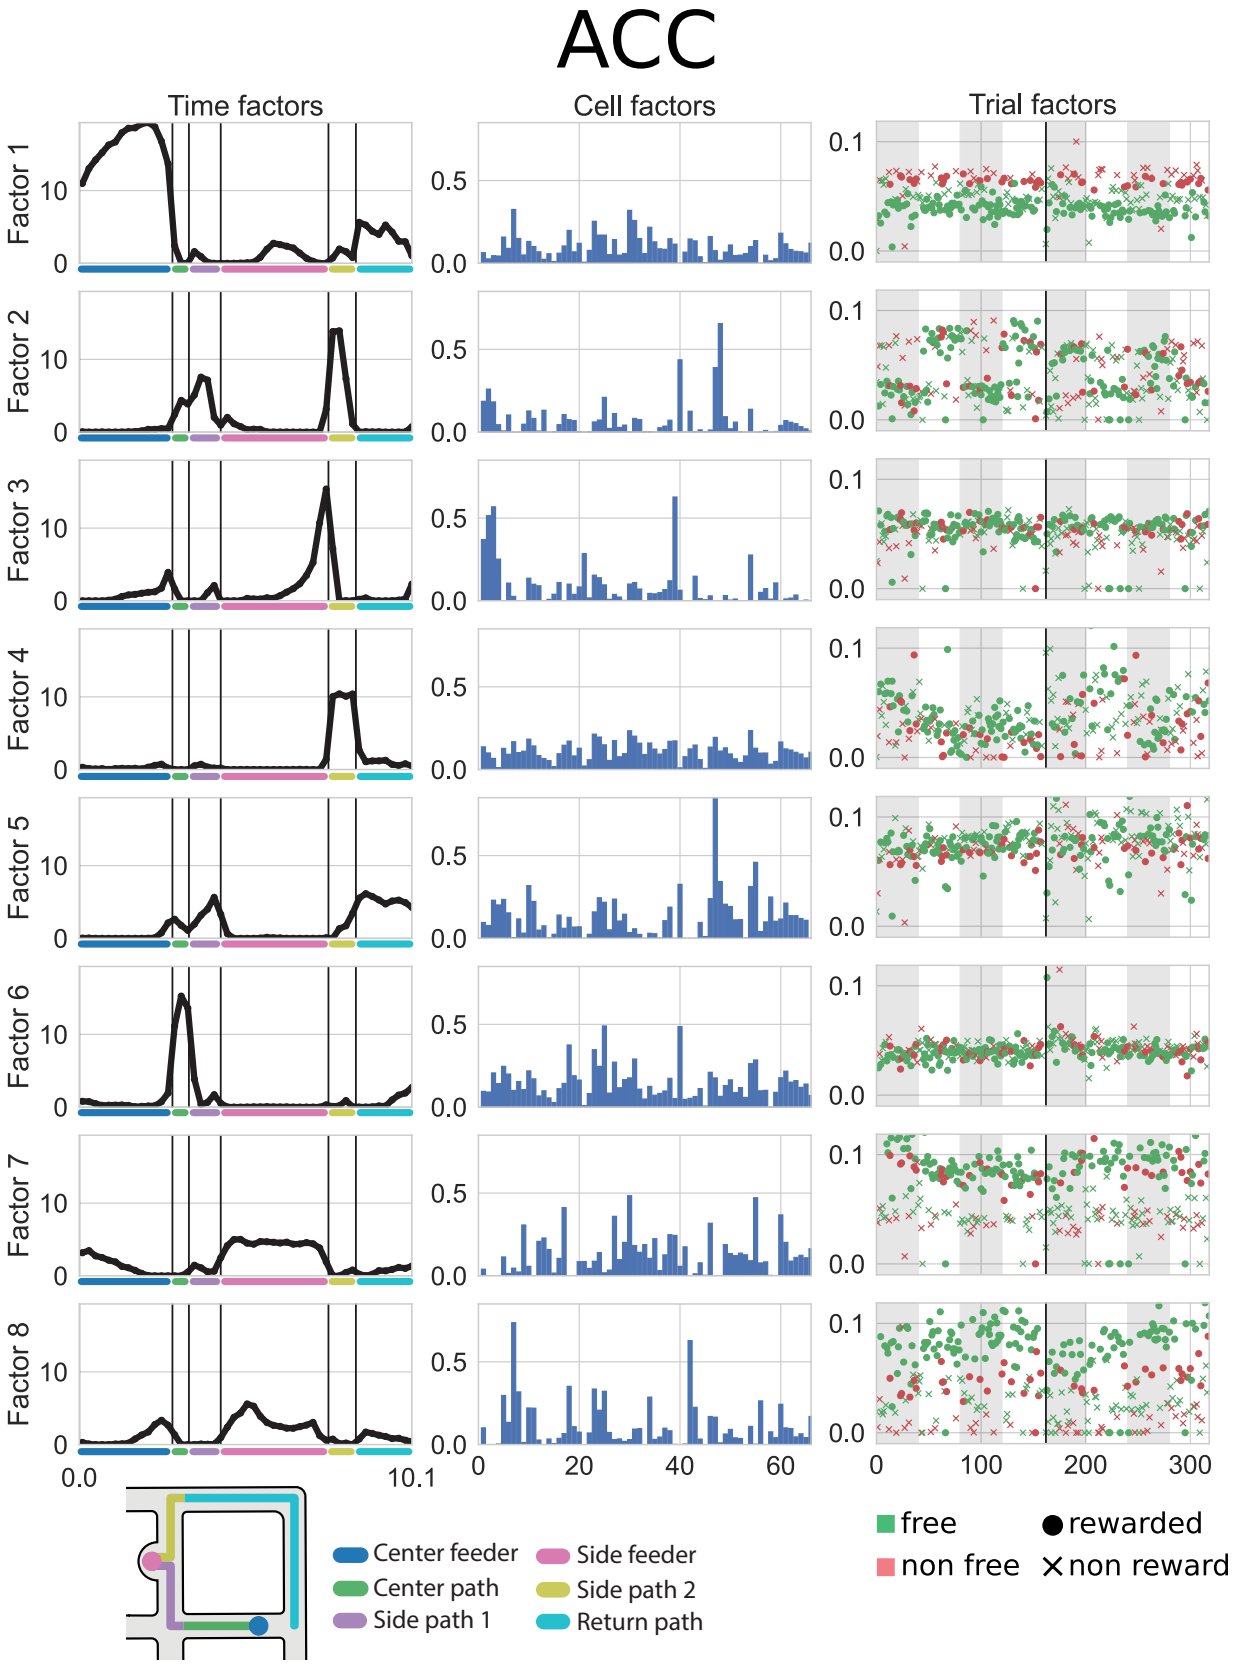

Figure S11: Complete list of factors for the TCA decomposition (see figure 4B for more details) for ACC. The cell factors are also shown, which generally show mixed activities of the cells.

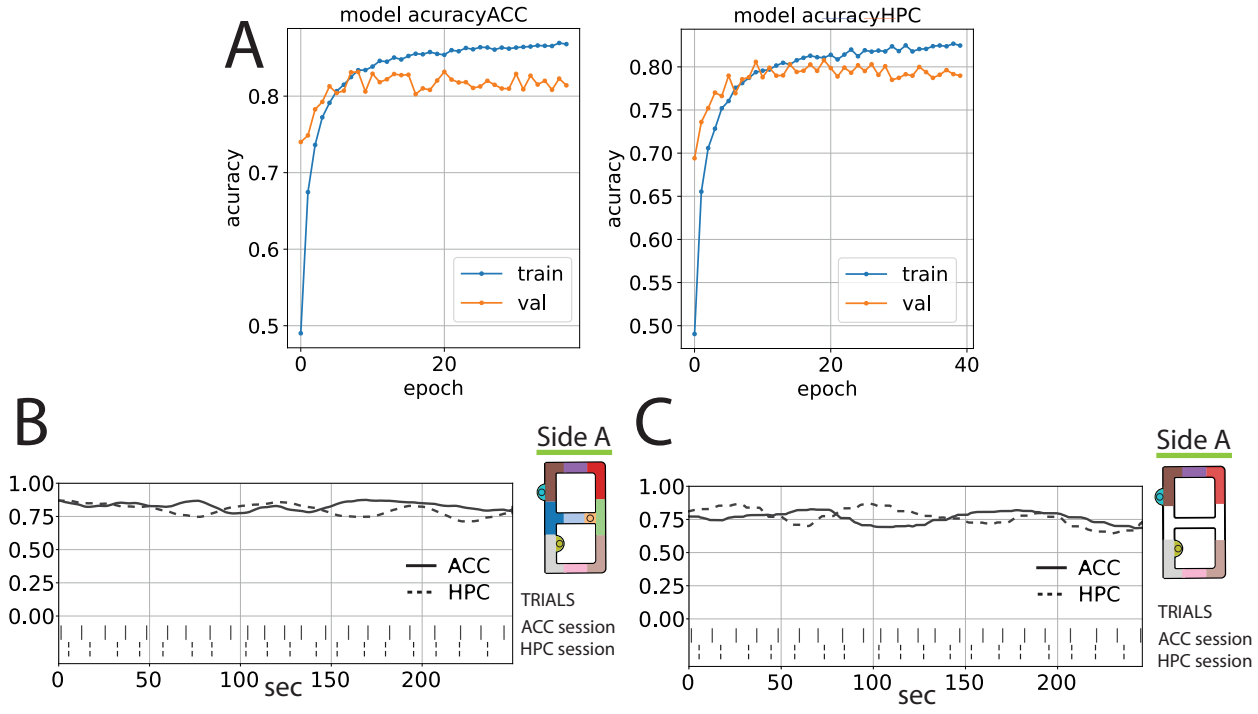

Figure S12: **A**. Train / validation curves show the accuracy of the neural network decoder in the training and validation set, respectively, as it is being trained. **B**. Decoder accuracy on across time for the validation data (side A).**C**. Decoder accuracy on across time for the validation data (side A) for the bins that are note shared between mappings.

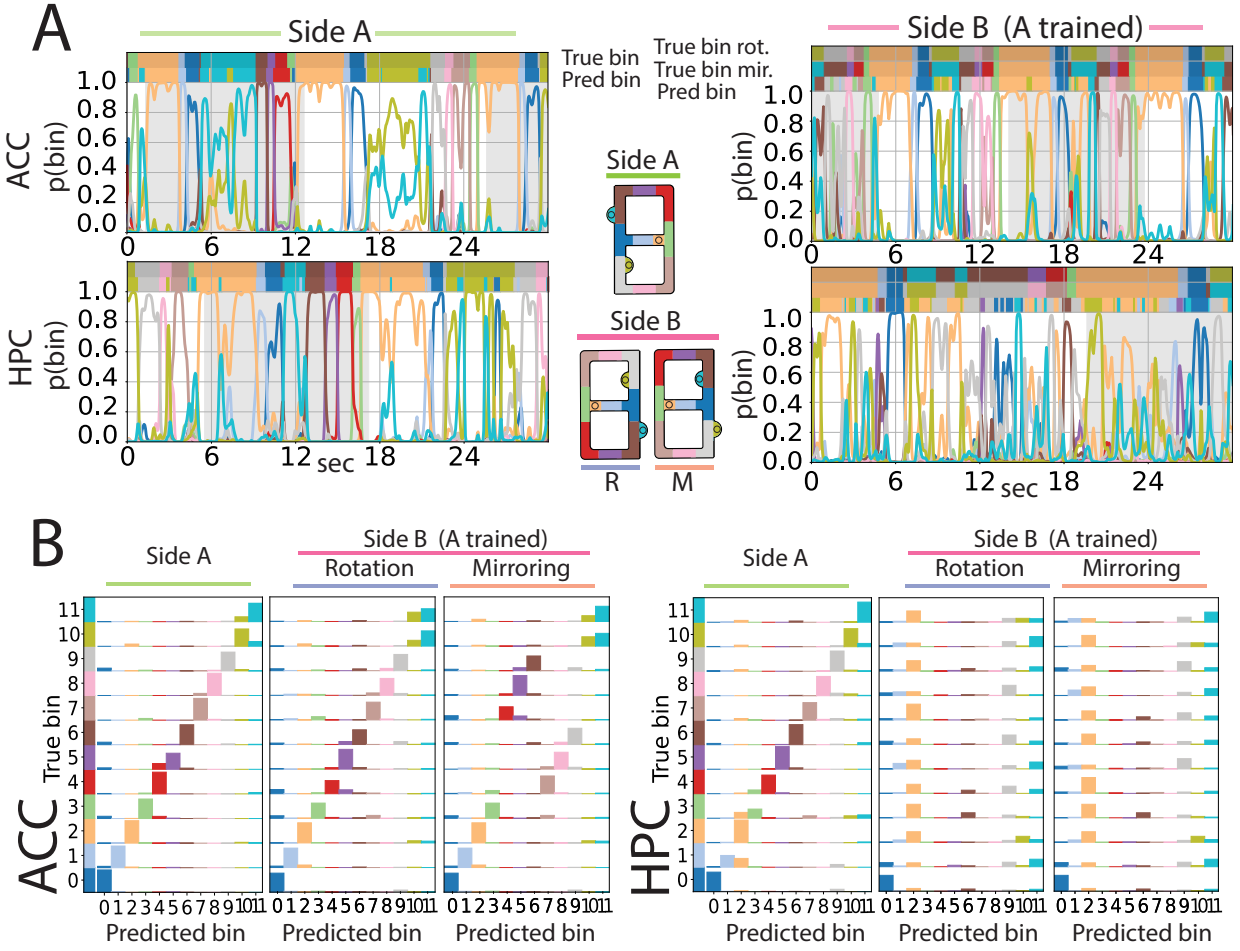

Figure S13: **A**. Sample activities of the network's last layer for a 30sec long time span. The curves indicate the probability that the network assigned to the corresponding bin. The predicted bin is on top, together with the correct bins for rotation and mirror mappings. In side A and ACC side B for rotation, even when the decoder makes a mistake the correct bin is often the next one in prediction strength. **B**. Distribution of predicted bins for each target bin. For the ACC rotation, most of the errors are in predicting the choice feeder bins (bins 10 and 11). For the HPC, the network seems to consistently over represent the central feeder (bin 2).

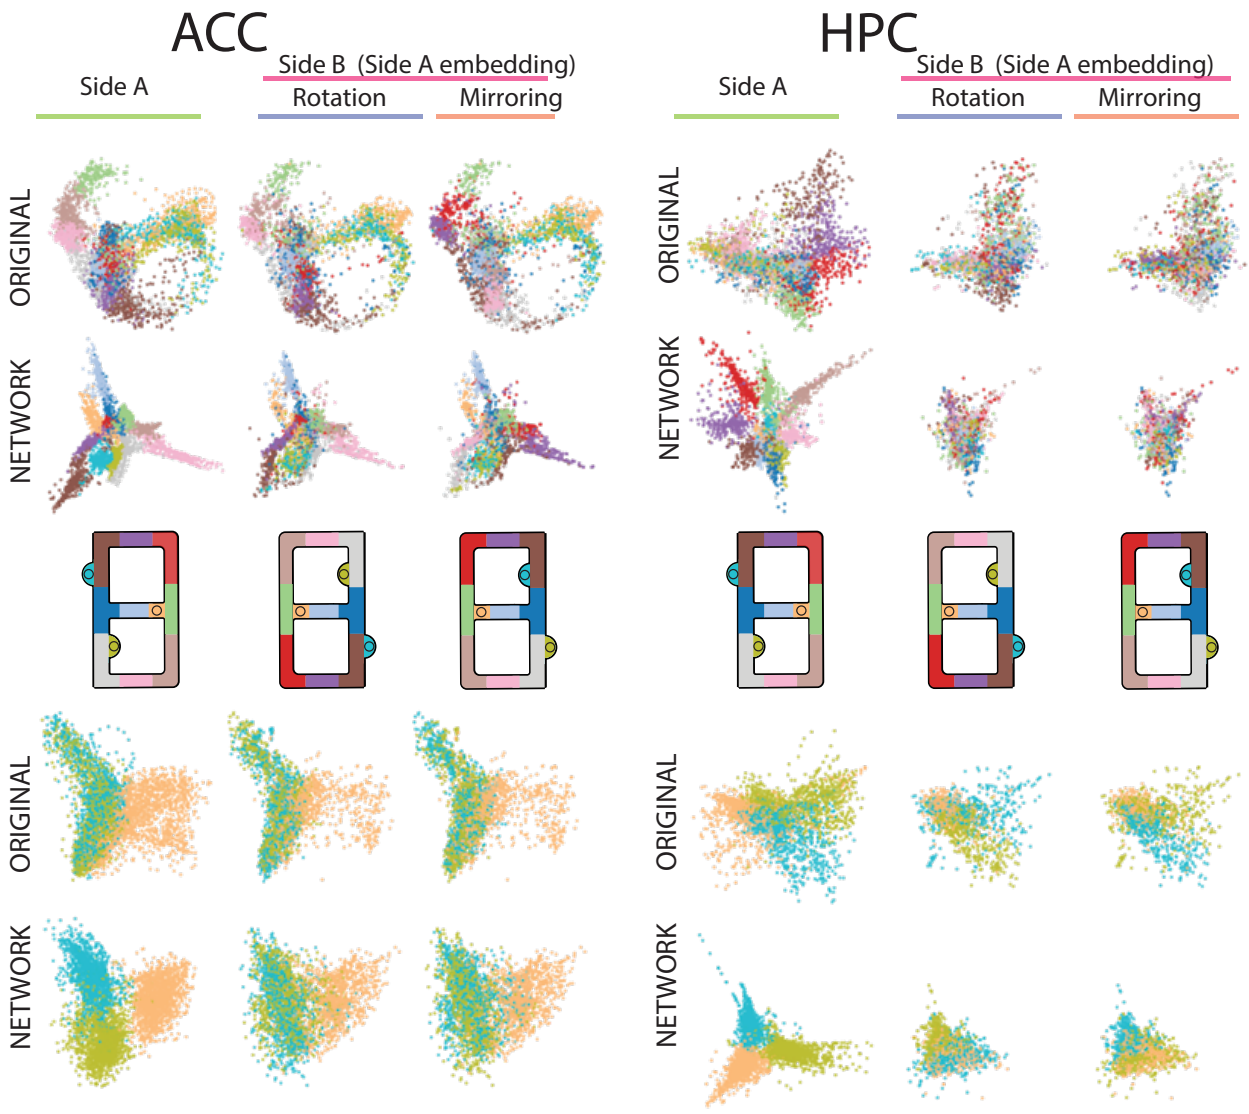

Figure S14: Neural data projected to the plane using nonlinear dimensionality reduction techniques (ISOMAP). Side B is projected using the embedding learned in side A. The corresponding bin labels are shown for every data point. The NETWORK rows are activities processed by the neural network decoder. The ORIGINAL rows are raw activity. On top, all the bins are projected, while on the bottom only the feeder bins are projected. The network successfully disentangles and clusters the bins in side A. For all the bins, only ACC preserves the manifold shape between side A and in side B. For ACC, the network cannot disentangle the choice feeder clusters (at least when projected in the plane), which might explain the decoding results in 3.4.