

Behavioral Neuroscience

Conjunctive Coding in an Evolved Spiking Model of Retrosplenial Cortex

Emily L. Rounds, Andrew S. Alexander, Douglas A. Nitz, and Jeffrey L. Krichmar

Online First Publication, June 4, 2018. <http://dx.doi.org/10.1037/bne0000236>

CITATION

Rounds, E. L., Alexander, A. S., Nitz, D. A., & Krichmar, J. L. (2018, June 4). Conjunctive Coding in an Evolved Spiking Model of Retrosplenial Cortex. *Behavioral Neuroscience*. Advance online publication. <http://dx.doi.org/10.1037/bne0000236>

Conjunctive Coding in an Evolved Spiking Model of Retrosplenial Cortex

Emily L. Rounds
University of California, Irvine

Andrew S. Alexander
University of California, San Diego and Boston University

Douglas A. Nitz
University of California, San Diego

Jeffrey L. Krichmar
University of California, Irvine

Retrosplenial cortex (RSC) is an association cortex supporting spatial navigation and memory. However, critical issues remain concerning the forms by which its ensemble spiking patterns register spatial relationships that are difficult for experimental techniques to fully address. We therefore applied an evolutionary algorithmic optimization technique to create spiking neural network models that matched electrophysiologically observed spiking dynamics in rat RSC neuronal ensembles. Virtual experiments conducted on the evolved networks revealed a mixed selectivity coding capability that was not built into the optimization method, but instead emerged as a consequence of replicating biological firing patterns. The experiments reveal several important outcomes of mixed selectivity that may subserve flexible navigation and spatial representation: (a) robustness to loss of specific inputs, (b) immediate and stable encoding of novel routes and route locations, (c) automatic resolution of input variable conflicts, and (d) dynamic coding that allows rapid adaptation to changing task demands without retraining. These findings suggest that biological retrosplenial cortex can generate unique, first-trial, conjunctive encodings of spatial positions and actions that can be used by downstream brain regions for navigation and path integration. Moreover, these results are consistent with the proposed role for the RSC in the transformation of representations between reference frames and navigation strategy deployment. Finally, the specific modeling framework used for evolving synthetic retrosplenial networks represents an important advance for computational modeling by which synthetic neural networks can encapsulate, describe, and predict the behavior of neural circuits at multiple levels of function.

Keywords: evolutionary algorithms, spiking neural networks, retrosplenial cortex, spatial representation, conjunctive coding

The retrosplenial cortex (RSC) acts as an association cortex for the integration of visual, spatial, and idiothetic information for navigation and memory formation (Alexander & Nitz, 2015; Chrastil, Sherrill, Hasselmo, & Stern, 2015, 2016; Cooper &

Mizumori, 1999; Epstein, 2008; Epstein, Parker, & Feiler, 2007; Hindley, Nelson, Aggleton, & Vann, 2014; Vedder, Miller, Harrison, & Smith, 2016). This integration is expressed in the RSC as a conjunctive code that binds incoming input to multiple internal and external frames of reference simultaneously (Alexander & Nitz, 2015). In principle, such encoding can support scene construction and/or recognition, which the RSC may use to access stored representations of familiar environments (Epstein et al., 2007) and compute path integration (Alexander & Nitz, 2017; Chrastil et al., 2015).

While a great deal is known about the structure and connectivity of RSC (van Groen & Wyss, 1990, 2003; Wyss & van Groen, 1992), the contribution of multisensory integration as performed by RSC to cognition has yet to be fully understood (Pothuizen, Aggleton, & Vann, 2008). Although there is evidence to suggest that RSC may transform inputs from one frame of reference to another (Alexander & Nitz, 2015; Byrne, Becker, & Burgess, 2007), and that it plays a critical role in spatial representation and navigation (Maguire, 2001; Miller, Vedder, Law, & Smith, 2014; Nelson, Hindley, Pearce, Vann, & Aggleton, 2015; Takahashi, Kawamura, Shiota, Kasahata, & Hirayama, 1997; Vedder et al., 2016), the full nature of its representations is unknown. In particular, how the disparate input streams targeting the RSC combine to

Emily L. Rounds, Department of Cognitive Sciences, University of California, Irvine; Andrew S. Alexander, Department of Cognitive Science, University of California, San Diego, and Department of Psychological and Brain Sciences, Boston University; Douglas A. Nitz, Department of Cognitive Science, University of California, San Diego; Jeffrey L. Krichmar, Department of Cognitive Sciences, University of California, Irvine.

This work was supported by the National Science Foundation (Award IIS-1302125). The authors declare no competing financial interests. Experimental design and methodology, Emily L. Rounds, Andrew S. Alexander, Douglas A. Nitz, Jeffrey L. Krichmar; software, Emily L. Rounds; formal analysis, Emily L. Rounds; writing original draft, Emily L. Rounds, and Jeffrey L. Krichmar; writing, review, and editing, Emily L. Rounds, Jeffrey L. Krichmar, Andrew S. Alexander, Douglas A. Nitz; funding acquisition, Jeffrey L. Krichmar and Douglas A. Nitz.

Correspondence concerning this article should be addressed to Emily L. Rounds, Department of Cognitive Sciences, University of California, Irvine, CA, 92697-5100. E-mail: roundse@uci.edu

encode different reference frames that shape animals' navigation strategies is an outstanding question.

Computational neuroscience and theoretical modeling is in a unique position to help fill in the gaps of our knowledge concerning RSC, but explanatory and predictive power in such models requires synthetic neural activity capable of capturing and reproducing the observed dynamics of neurons recorded during task performance. In the case of RSC, such a model would need to faithfully generate the varied and complex firing patterns observed across ensembles of RSC cells according to the same behavioral correlates to which biological activity is anchored.

To these ends, we developed a novel approach for modeling brain regions that replicates ensemble responses based on the behavioral variables and neurophysiological recordings collected during behavioral experiments. In contrast to approaches that emphasize the setting of specific connection weights among neural elements, our approach uses a GPU-accelerated evolutionary process to search metaparameters for learning itself, such as for spike timing dependent plasticity (STDP) and synaptic homeostasis (Rounds et al., 2016). While other groups have used GPU acceleration to evolve neural networks (Carlson, Beyeler, Dutt, & Krichmar, 2014; Dura-Bernal et al., 2016; Fidjeland & Shanahan, 2010; Fountas & Shanahan, 2015; Kasabov, Dhoble, Nuntalid, & Indiveri, 2013; Paine, Jin, Yang, Lin, & Huang, 2013), none have used the method to generate mappings between experimentally recorded behavior and neural activity.

In the present work, models evolved upon RSC ensemble spiking dynamics underwent a series of virtual experiments that would be difficult and time-consuming to conduct in awake, behaving animals. The results indicate that the evolved neuronal activation patterns exhibit a kind of mixed selectivity, in which neuron function and behavior was modulated by the position and orientation of the track in allocentric space. Critically, the flexibility of response observed both at the level of the neuron and the population was not explicitly evolved. Instead, it emerged by virtue of replicating RSC firing patterns in an embedded subset of the population. The results of the simulated experiments demonstrate that, as a consequence of mixed selective responses, model RSC spontaneously generates representations for position, orientation, and action to which it was never exposed during training. This suggests that mammalian RSC is home to a highly adaptable representation of space that can immediately be applied to novel conditions. Furthermore, the model survives the loss of seemingly critical input streams, and disambiguates position when position and orientation variables are in conflict, suggesting the RSC supports a robust, conjunctive encoding of spatial variables.

Method

Experimental Task and Dataset

We used an electrophysiological dataset collected from the retrosplenial cortex (RSC) in six male Long-Evans rats (Alexander & Nitz, 2015). The rats performed a spatial navigation task that involved completing inbound and outbound routes along a W-shaped track that involved either a left-right-left (LRL) or right-left-right (RLR) turn sequence. Four behavioral variables were recorded concurrently with neural activity as the rats performed the task: (a) head direction (HD); (b) allocentric position in

x-y coordinates (Pos); (c) angular velocity (AV); and (d) linear velocity (LV). In each of the 71 recording sessions, rats were required to complete a series of inbound and outbound runs on the track that occupied two separate allocentric locations in the room (α and β). In each recording session, the track was placed in a different location. The rats collectively completed an average of 16 trials on each possible route (α LRL, α RLR, β LRL, β RLR) across all sessions. To analyze neural activity, the W-shaped track was divided into 200 equally sized bins, each approximately 1.2 cm wide. These bins were used to determine the precise points at which a rat reached a turn apex on each run through the track.

CARLsim and Automated Parameter Tuning Framework

Computer simulations were performed by using the CARLsim spiking neural network (SNN) simulator (Beyeler, Carlson, Chou, Dutt, & Krichmar, 2015; Carlson et al., 2014). CARLsim is a C/C++ based simulator that allows the user to construct biologically detailed spiking networks that may be specified at the synapse, neuron, and network level. In the present study, we used CARLsim Version 3.0, which included excitatory spike-timing dependent plasticity (STDP), inhibitory STDP, and a homeostatic mechanism for synaptic scaling to prevent runaway activity (Carlson, Richert, Dutt, & Krichmar, 2013). Both STDP and homeostasis act on the postsynaptic neuron. The CARLsim simulator is GPU-accelerated using off-the-shelf NVIDIA GPUs, which allows detailed large-scale networks to execute rapidly and efficiently and allows multiple SNNs to be run in parallel.

Network Architecture and Evolutionary Algorithm

CARLsim 3.0 includes an automated Parameter Tuning Interface (PTI), which utilizes a plugin to a library for evolutionary computation (Evolutionary Computations in Java, or ECJ; Rounds et al., 2016; White, 2012), in conjunction with a user-defined fitness function, for tuning parameters (Figure 1A). Fitness is calculated on the CARLsim side for each individual in the population. Following the evaluation phase, the fitness values are sent to ECJ for selection. After parent individuals are chosen, their parameters undergo mutation and/or recombination to initialize a new population for each subsequent generation.

Because SNNs feature a large number of open parameters, it can become difficult to tune each parameter manually. Utilizing an evolutionary strategy allows the user to efficiently and automatically search the parameter space. For the present experiments, the PTI used binary tournament selection to select three parent individuals, which underwent mutation to initialize a new generation of 15 children. Recombination was not used in this experiment. The PTI allows the user to specify how many parameters are to be evolved, the range of the parameter space, mutation and recombination parameters, and selection method for choosing individuals for reproduction. We optimized 18 parameters corresponding to E-STDP, I-STDP, and homeostasis on each group, as well as the initial weight ranges for each set of connections (for more details on parameter initialization, see Table 1 in Rounds et al., 2016). Because STDP operates on the postsynaptic neuron, the excitatory neuronal group had both E-STDP (for excitatory input and recurrent projections) and I-STDP (for inhibitory projections), while the




Figure 1. (A) General approach to evolving spiking neural networks (SNNs) using evolutionary computation system written in Java (ECJ). ECJ is used to initialize the parameters being tuned, which are passed to CARLsim and used to set the parameters of each network in the population of SNNs. The SNNs are initialized, run, and evaluated using CARLsim. Following fitness evaluation using some user-defined criteria, the resulting fitness values for each individual in the population are sent to ECJ. ECJ chooses parents from the population and performs selection, replication, and mutation on the parameters of the parent individuals in order to initialize the next generation of SNNs. This continues until an ideal solution is found or until a user-defined number of evolutionary runs are completed. Adapted from Beyeler et al. (2015). (B) Network architecture showing the input groups with the corresponding number of neurons associated with each recorded behavioral variable and how they are connected to the excitatory and inhibitory neuron groups. The network featured four input types, which were input to both the excitatory and inhibitory neurons, recurrent excitation, and feedforward inhibition. (C) Mean fitness values and standard deviation associated with all 10 networks for each evolutionary run. At the beginning of evolution, fitness varied dramatically, but a good solution was achieved by approximately the 20th generation of evolution. Networks were allowed to evolve for 50 generations. See the online article for the color version of this figure.

inhibitory neuronal group was only subject to E-STDP because it only received excitatory projections from the input neurons and received no other sources of input. The excitatory input group received no feedback input and was thus not subject to any kind of STDP. Following evolution, we examined the parameter solutions found by the evolutionary algorithm for each run in order to investigate whether the values between runs tended to fall in similar parts of the parameter space, or if there were multiple possible solutions to the problem. Over 10 networks, with 18 evolved parameters per network, we only found two outliers, one which corresponded to the A^+ parameter for E-STDP on the inhibitory neurons, and another which corresponded to the initial weight value of the $exc \leftrightarrow exc$ connections. Otherwise, the evolved parameters seemed to cluster closely around the mean, suggesting that the solutions found for each of the 10 networks were similar.

Using the PTI, we initialized a population of 15 spiking neural networks. Each network contained 1,017 neurons in total: 480 excitatory neurons, 120 inhibitory neurons (80/20 network), and Poisson spike generators for the input neurons (Figure 1B). The input neuron layer had 390 neurons for allocentric position, eight neurons for head direction, seven neurons for angular velocity, and 12 neurons for linear velocity. The network featured feedforward inhibition and recurrent excitation, resulting in four connection types ($inp \rightarrow exc$, $inp \rightarrow inh$, $exc \leftrightarrow exc$ and $inh \rightarrow exc$). The behavioral inputs were parameterized by creating idealized input neurons using broadly tuned Gaussian and cosine tuning curves. By examining the range of values present for each variable in the dataset, we created an appropriate number of tuning curves to encompass the input space.

To create inputs for linear and angular velocity, the following Gaussian tuning curve was used (Eq. 1):

$$f(n, b) = r_{max} \times e^{-\left(\frac{(input_b - pref_n)^2}{2\sigma}\right)}, \quad b = 1:200 \text{ and } n = 1:N. \quad (1)$$

For linear velocity, the parameters were as follows: $N = 7$, $r_{min} = 0$, $r_{max} = 60$, and $\sigma = 7$. For angular velocity, $N = 12$, $r_{max} = 40$, and $\sigma = 12$.

To create inputs for head direction, the following cosine tuning curve was used (Eq. 2):

$$f(n, b) = \begin{cases} r_{max} \times \cos\left(\frac{\pi x}{360} \times (input_b - \pi)\right) \\ -\left(\frac{\pi x}{360} \times (pref_n - \pi)\right) & \text{if } input_b \geq (pref_n - 45), \\ 0 & \text{else} \end{cases} \quad (2)$$

$b = 1:200$ and $num = 1:N$ where $N = 8$, $x = 3.99$, $r_{max} = 40$, and $\sigma = 8$.

For allocentric position, which was tracked in Cartesian coordinates, a two-dimensional (2D) Gaussian curve was used to create idealized place cells by adapting the equation from (Foster, Morris, & Dayan, 2000; Eq. 3):

$$f(n, b) = r_{max} \times e^{-\left(\frac{\|location_b - pref_n\|^2}{2\sigma^2}\right)}, \quad b = 1:200 \text{ and } n = 1:N, \quad (3)$$

where $N = 390$, $r_{max} = 60$, and $\sigma = 40$. The variables $location$ and $pref$ are x-y-coordinate pairs. $location$ represents the actual position of the agent and $pref$ represents the center of the place cell.

Training and Testing Phase During Evolution

A total of 4,593 trials, each of which consisted of a route (α LRL, α RLR, β LRL, and β RRL), were conducted and recorded

over the course of all 71 recording sessions in the original study conducted by (Alexander & Nitz, 2015). This gave us a sizable pool of data on which to train and test the SNNs. To establish the pools of training and testing trials, we took the number of trials for each route combination in each recording session and divided them in half to create the training and testing sets. The training and testing paradigm consisted of randomly choosing 150 trials for each route combination (600 trials in total; 150×4) in each phase that were presented to the network (Rounds et al., 2016). To ensure that the experimental design was as similar as possible to the original experiment, trials were presented in terms of alternating routes. For example, if the first trial presented was selected from α LRL, then the following trial would be selected from α RRL, which would be repeated 10 times, and then the experiment would switch to track β . Following training, synaptic plasticity was disabled while the network was exposed to the testing set. Following testing, fitness was evaluated and the generation ended, at which point a new generation of individuals was initialized via the PTL. This process continued for 50 generations per each evolutionary run. Ten evolutionary runs were conducted in order to obtain a set of high fitness individual networks for subsequent analysis.

Network fitness was computed by calculating the correlations between synthetic and experimentally observed neuronal firing rates averaged over all trials seen by the individual in the testing phase and choosing best matches (i.e., highest correlations) between synthetic and experimental neurons (Eqs. 4 and 5):

$$f(x) = \begin{cases} \sum_{i=1}^n \rho(\text{realFR}_i, \text{synFR}_{\text{match}}) & \text{if } \text{maxFR} < \text{FR}_{\text{target}} \\ \sum_{i=1}^n \rho(\text{realFR}_i, \text{synFR}_{\text{match}}) - \text{FR}_{\text{error}} & \text{else} \end{cases} \quad (4)$$

where

$$\text{FR}_{\text{error}} = \text{FR}_{\text{max}} - \text{FR}_{\text{target}} \quad (5)$$

and $\text{FR}_{\text{target}} = 250$ Hz was the maximum mean firing rate threshold.

Once a synthetic neuron was matched to an experimental neuron, that neuron was excluded from the pool of neurons that could be matched. After choosing a best match for each experimental neuron, the correlations were summed to give the fitness value. A maximum firing rate component was added to the fitness function to ensure a reasonable firing regime—if the average maximum firing rate (over all trials) of a neuron exceeded a threshold of 250 Hz, the threshold was subtracted from the observed average maximum and the difference was subtracted from the sum of the correlations between matched neurons to give the final fitness. After all generations were complete, the highest fitness individual from each run was saved for further analysis and experiments. Spike trains were converted to mean firing rates as in (Alexander & Nitz, 2015) by applying a Gaussian filter to each firing rate profile for each trial (each firing rate profile had a size of 200 bins; the smoothing filter had a standard deviation of six bins).

Examining the Impact of Lesions on Network Performance

In order to test the conjunctive nature of RSC activity, we simulated lesions of projections from one neural group to another by setting the synaptic weights of those synaptic connections to zero. These simulated lesions were conducted on the SNN with the highest fitness from an evolutionary run. The lesioned SNN underwent a new testing and fitness evaluation phase. The recorded behavioral inputs were presented to each lesioned network again in the same order to ensure consistency between the baseline and lesioned conditions. The individual's performance under each lesion condition was then compared with its intact performance in terms of its fitness, its ability to reconstruct a route, and the distributions of functional neuron types.

Rearranging Turn Sequences on Tracks

To ensure that the network had learned to respond generally to features of a route and not to the specific shape of the route it was trained on, we tested the network with different turn sequences. This was achieved by selecting a recording session and altering the behavioral variables such that the movement would correspond to this new track shape. This was done by rearranging the structure of the track (essentially by switching out bin segments to reflect the desired shape). We altered the trials corresponding to three separate recording sessions to obtain six possible alternate turn sequences: LLL (outbound)/RRR (inbound); LLR (outbound)/LRR (inbound); and RLL (outbound)/RRL (inbound). For LLL/RRR, 10 new trials were constructed and presented to each evolved network; for LLR/LRR, 13 new trials were constructed, and for RLL/RRL, nine new trials were constructed. Network performance was evaluated on fitness, ability to reconstruct the route, and functional neuron type distribution, as with the lesioned networks.

Rotating Track Positions Without Shifting Global Allocentric Track Position

To systematically examine network responses to changing environments, we rotated the W track. Trials corresponded to five recording sessions where we imposed a gradual rotational shift of the track. Each track was shifted around its center point, or the middle turn of the track, in 20° increments. This resulted in 18 positions for each track, including the original position α and β . Positions α and β occupy disparate allocentric positions depending on the recording session. In recording Session 1, the tracks overlap at their center point, meaning that at some point during rotation, it is possible for the tracks to occupy the same approximate allocentric position and orientation. They are also situated close to one another at their center points in recording Session 3. In recording Sessions 2 and 4, their center points are relatively far apart, but are close enough to one another that at some point, track segments of α and β will be overlapping, but not such that similar turns occupy similar allocentric positions. In the case of recording Session 5, the tracks are entirely nonoverlapping. These distances allowed us to investigate network responses under a range of conditions (overlapping, semioverlapping, parallel but nonoverlapping, etc.).

To create these new routes, two behavioral variables were manipulated (Pos and HD). To dissociate differences in network

performance on the fitness metric due to changing one behavioral variable (translation case) or two behavioral variables (orientation case), additional track orientation experiments were carried out in which the head direction was not shifted to be consistent with the new allocentric position. This led to another set of orientation experiments (orientation, fixed HD) in which head direction was always consistent with the original position of the track (Position 1, or 0° shift, out of the 18 rotated positions).

Linearly Translating Track Positions Across the Allocentric Space

To systematically examine network responses to changing allocentric position, we translated the position of the W track. Five separate recording sessions were altered to reflect gradual movement of the track position. The recording sessions were chosen such that the original track positions were parallel to one another but occupied different allocentric positions (with varying amounts of distance between them). For each recording session, this resulted in trials that corresponded to six track positions in which the track shifted in consistently sized steps to a final position. The step size was consistent for all recording sessions—all tracks were moved in increments of 32 cm. To translate track locations, the x-y coordinates corresponding to allocentric position (Pos) were modified. Thus, HD, AV, and LV remain consistent with the original position of the track (track position one out of the six possible positions).

To further investigate network responses under conditions in which the two tracks occupy opposite orientations, trials were constructed in which the orientation of track position α was reversed. This was done for all five recording sessions used for the translation experiments. In this manipulation, the allocentric position (Pos) and the head direction (HD) variables were altered. This also essentially flipped the direction of the route so that if the original start position were on the right-hand side of the track, then to maintain the same turn sequence, the start position would shift to the left-hand side of the track under the reoriented track condition. The inverse translation experimental design constitutes investigation of network response when the track is reoriented as well as linearly translated.

Statistical Analysis

Statistical tests were conducted using MATLAB (MathWorks). Correlations between experimentally recorded and synthetic neural data were calculated using the Pearson's R correlation coefficient. Where relevant, outlier detection was performed using the Generalized Extreme Studentized Deviate (G-ESD) test to find parameter values that fell outside of two standard deviations of the mean, using $r = 4$ (maximum expected number of outliers) for 10 total samples. A Wilcoxon's rank sum test was used to compare ensemble averages. The Kolmogorov–Smirnov test was used to determine significance when comparing randomized and real cumulative distribution functions. Comparisons of fitness, functional neuron type distribution, and positional ensemble reconstruction error were carried out using a Kruskal-Wallis test with Tukey's honestly significant difference (HSD) criterion for post hoc analysis at $p < .05$.

For analysis of reconstruction of the agent's position on a route, positional ensemble reconstruction matrices were made by cross-

correlating the average firing rates for all neurons across trials (either all trials, or divided into pools of even- and odd-numbered trials) associated with track positions α and/or β for both routes (LRL and RLR). The maximum values were found for every column within the resulting correlation matrix, which were connected to form the prediction line. The error for each bin was averaged over all 10 evolved networks to obtain averaged reconstruction errors. For further details, see methodology described in Alexander and Nitz (2015).

For statistical comparisons, n represents the number of synthetic RSC neurons. Data are presented as mean and standard deviation. Where best-fit lines are plotted, we fitted polynomial equations to the relevant data using the MATLAB Curve Fitting Toolbox. This toolbox also allowed for analysis of confidence intervals on the parameters of the equation. Equations in which the confidence interval for a parameter crossed zero were excluded from consideration due to the possibility that the parameter was obsolete.

Code Availability

The CARLsim SNN simulator is freely available from the CARL lab of the University of California, Irvine at <http://www.socsci.uci.edu/~jkrichma/CARLsim/index.html>. The ECJ toolkit is freely available from the ECLab of George Mason University at <https://cs.gmu.edu/~eclab/projects/ecj/>. All analysis and modeling code can be made available by the authors upon request.

Data Availability

All data for the experimental and theoretical work can be provided by the authors upon request.

Results

The CARLsim spiking neural network simulator was used for the present simulations (Beyeler et al., 2015). CARLsim has a Parameter Tuning Interface (PTI) between CARLsim and an evolutionary computation system written in Java (ECJ; White, 2012) that provides a means to automate the parameter tuning process (Figure 1A). Evolutionary algorithms (EAs) enable flexible parameter tuning by means of optimizing a generic fitness function (a metric used to compute a single number that determines the network's adherence to the desired criteria). The automated tuning framework was used to match synthetic neuronal firing rates with neurophysiological recordings, which successfully evolved emergent network functionality. As described in an earlier publication (Alexander & Nitz, 2015), the neuronal activity profiles were obtained from 71 recording sessions between which recording wires were moved in order to maximize independence of neuron sampling. The dataset contained 243 RSC cells, including 15 head direction cells, recorded as rats ran back and forth repeatedly along a W-shaped track. The track was placed in two different locations in the experimenter room (α and β). The rats' angular velocity, linear velocity, head direction, and allocentric position were recorded during behavior via overhead tracking cameras (Alexander & Nitz, 2015). These same behavioral metrics were used as inputs to the SNN (Figure 1B) to approximate the inputs to RSC. However, these behavioral correlates are not meant to be taken to represent literally the nature of inputs to the biological RSC;

instead, they approximate the kinds of information processed and/or utilized by connected regions such as the hippocampus and parietal cortex. These analogue inputs are also not exhaustive because they do not incorporate visual input or feedback from higher cortical regions.

Using an indirect encoding approach, the STDP and homeostasis parameters governing learning were evolved with ECJ until a strong correspondence between synthetic and experimental data was obtained. To compute fitness, we measured the Pearson's R correlation between simulated and real neuron activation patterns and incorporated constraints on the upper and lower bounds of the simulated neuronal firing rate (see Method section for more modeling details). To measure the robustness of the evolved networks, lesions of behavioral inputs were simulated by setting connection weights between the relevant input group(s) and the excitatory and inhibitory groups to zero. In another set of experiments, behavioral data corresponding to certain track positions was altered to simulate either rotation or linear translation of the track, which involved altering positional information for both kinds of manipulations, as well as head direction information for the track rotation manipulations (see Method section for further details).

Statistical Note

Significant differences between multiple groups were assessed using a nonparametric version of ANOVA called the Kruskal-Wallis test. This test computes the χ^2 statistic that determines whether samples in the dataset come from the same distribution. The resulting statistic was corrected with Tukey's HSD post hoc comparison. Any significant p values associated with the Kruskal-Wallis tests, reported here, passed the Tukey's HSD correction at $p < .05$.

Evolved SNNs Capture the Behavior of RSC Neural Ensembles

The automated PTI (Figure 1A) successfully evolved networks (Figure 1B) whose synthetic neuronal firing rate profiles replicated the neurophysiological dataset. All evolutionary runs were allowed to continue for 50 generations. We found, over the course of 10 evolutionary runs, that the algorithm took approximately 20 generations to converge upon a solution, with some fitness gains in subsequent generations. The highest observed fitness was 107.79 (Figure 1C), which roughly corresponds to an average Pearson's $R = 0.44$ between the trial-by-trial firing rates for a synthetic neuron and a match to the trial-by-trial firing rates for each of the neurons recorded in behaving animals. The lowest observed fitness was 104.7, or average Pearson's $R = 0.43$ per neuron (strong correlations by experimental standards). At the start of each evolutionary run, the average maximum fitness score was $F = 84.57 \pm 19.78$.

Similar to RSC neurons recorded in behaving animals, the evolved networks effectively encoded all three spatial frames of reference (allocentric, egocentric, and route-centric). Figure 2 shows one example for each functional neuron type where the black trace is from the biological neuron and the red trace/gray trace is from the simulated neuron. A "turn insensitive" neuron (Figure 2A) is not tied to any particular action correlate by its activation pattern, but nevertheless reliably generates a complex

spatially specific firing pattern. In contrast, purely "turn sensitive" neurons (Figure 2B) respond primarily to a preferred turn type (in this case, left turns), regardless of where such actions occur along a route. Between these extremes, some "turn sensitive" neurons that have spike rates specific to route position (Figure 2C) also respond to specific turning actions. In such cases, higher firing rates are found for a specific instance along a route (e.g., in Figure 2C; higher firing rates for the first right turn on the RLR route and the first left of the LRL route).

The evolved networks replicated all observed features of the RSC data. For example, each network had a similar distribution of functional neuron types to those reported in Alexander and Nitz (2015) where RSC firing according to route position, turning action, environmental position, and conjunctions of these variables were observed (Alexander & Nitz, 2015; Figure 3A). Neurons from each network were categorized into functional types, and then the numbers of neurons of each type were averaged for 10 different network instantiations. About 64% of the neurons in each network ($\bar{n} = 360.5 \pm 26.12$ out of 575.3 ± 15.57) were considered turn sensitive, as determined by comparisons of mean firing rates for all left versus right turn instances (Wilcoxon's rank sum test, $p < .05$). The remaining 37% of the neurons in each network ($\bar{n} = 214.8 \pm 33.05$ out of 575.3) were considered turn insensitive (Wilcoxon's rank sum test, $P \geq 0.05$). Additionally, a subset of turn insensitive neurons ($\bar{n} = 42.8 \pm 8.97$ out of 214.8) exhibited "route modulated" activity in that their mean firing rate profiles for same-route comparisons were highly correlated even when the route was moved to different locations in the larger environment (e.g., firing pattern for α LRL correlated with firing pattern for β LRL, Pearson's $R > 0.4$, $p < .05$).

Turn sensitive neurons were also divided into two groups, in which approximately half were purely turn sensitive, in that they responded to all instances of the preferred turn equally ($\bar{n} = 179.8 \pm 48.36$ out of 360.5), or were route modulated in that they preferred either the first or second instance of the preferred turn along the route regardless of the allocentric position of the track ($\bar{n} = 180.7 \pm 35.07$ out of 360.5). This was determined by comparing all mean firing rates associated with the first left or right turn on the route with those associated with the second left or right turn on the route for each neuron. Those with significant differences associated with the first versus the second turn were considered to exhibit route modulated activity (Wilcoxon's rank-sum, $p > .05$), while those without were considered purely turn sensitive (Wilcoxon's rank-sum, $p < .05$). The former encode conjunctions of route position and associated action, effectively mapping the spatial context within which navigational behaviors take place. The majority of the cells in the evolved networks were turn insensitive (i.e., turn insensitive, no mod. and turn insensitive, route mod. in Figure 3A), which was consistent with the dataset (see Figure 3a of Alexander & Nitz, 2015). Finally, we also found similar distributions of head direction cells in the evolved networks. On average, $4.12\% \pm 2.6\%$ ($\bar{n} = 24.7 \pm 15.57$) of the neurons in each evolved network were head direction cells. This percentage is slightly lower than that reported in Alexander and Nitz (2015; 6%) and Cho and Sharp (2001; 10%).

To ensure that turn preference in the turn sensitive neuron population was not the product of chance, we computed a true and randomized turn bias ratio for each turn sensitive neuron. This was accomplished by creating a turn identity matrix in which the

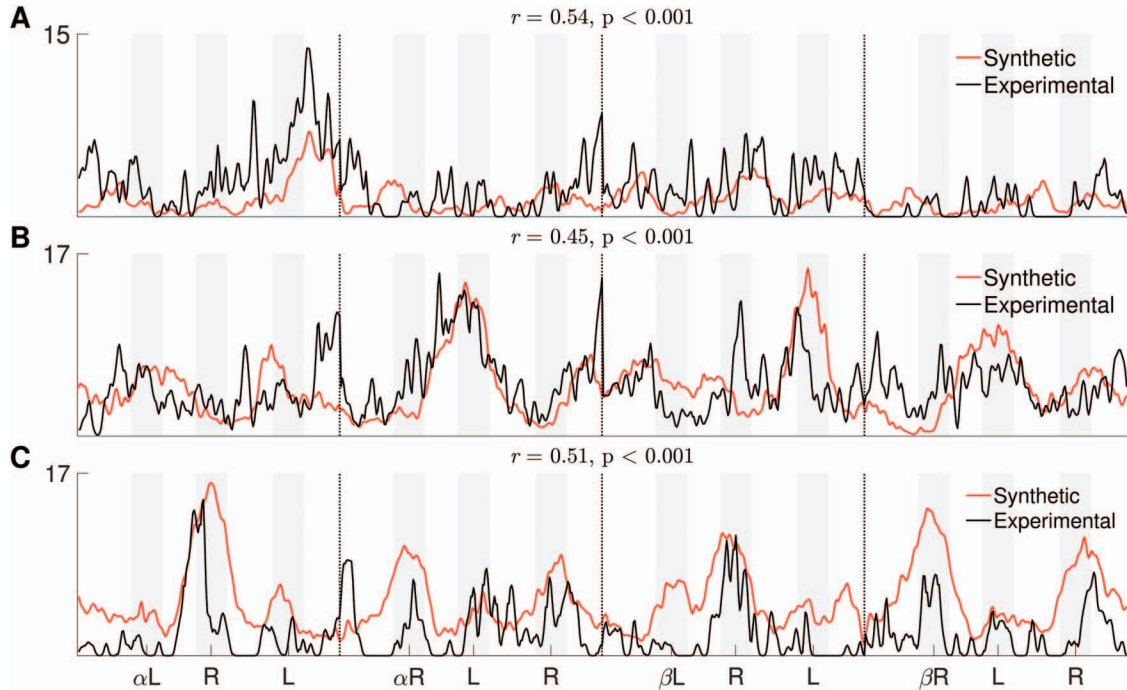


Figure 2. The original experimentally observed activity profile for a biological neuron is shown along with its synthetic matched counterpart. Activity profiles are shown for each route (labeled at the bottom of the plot), which are separated by vertical dotted lines. Turn integration sites are indicated by shaded regions in the background. The Pearson's R correlation coefficient for each match is shown at the top of each plot for each of the three neurons. (A) Example of a turn insensitive neuron (Pearson's $R = 0.54$, $p = 5.12 \times 10^{-62}$). (B) Example of a left turn sensitive neuron without route modulation (Pearson's $R = 0.45$, $p = 7.25 \times 10^{-42}$). (C) Example of a right turn sensitive neuron with route modulation (Pearson's $R = 0.57$, $p = 2.04 \times 10^{-54}$). Though the firing patterns of these neurons suggest that they respond to left turn instances, peak firing rates associated with running behavior occur at the apex of right turn sites for right-left-right (RLR) routes only. These neurons fired significantly more for the first right turn on an RLR turn sequence (Wilcoxon's rank sum test, $p < .05$). See the online article for the color version of this figure.

average firing rate was calculated for all instances of all turn types (12 turns in total, six left and six right), and then the ratio of averaged activity for left versus right turns was taken. The turn identity matrices were subsequently randomized over 25 iterations and recomputed to create a randomized cumulative distribution function (red/lighter gray line in Figure 3B) that was plotted against the actual cumulative distribution function of turn bias ratios (blue/dark gray line in Figure 3B). The true distribution of turn bias ratios for all turn sensitive synthetic cells was significantly greater than the distribution of ratios calculated by randomizing turn identity (25 of 25 iterations, one-tailed Kolmogorov-Smirnov, ($\bar{n} = 360.5$, $D > 0.4$, $p < .0001$).

Similar to Alexander and Nitz (2015), we verified that the evolved networks could reconstruct the agent's position along a route from positional ensemble correlations across the whole neuron population. We created positional ensemble reconstruction matrices following methods described in Alexander and Nitz (2015). This involved creating an average firing rate profile for every neuron based on even trials, odd trials, and all trials. For within-position comparisons (e.g., position α , route LRL even vs. position α , route LRL odd) even and odd average firing rate profiles for all neurons were cross-correlated to generate the re-

construction matrix (Figures 3C–D). Neuronal activations that occur at the same points in route space should be highly correlated, while those that occur at different points should not be. This creates an ideal prediction line that crosses the matrix diagonally from the left top corner to right bottom (dotted white line). A black line was overlaid on the reconstruction matrix that falls along the points of actual highest correlation between bins. When even versus odd trials were compared, the black line fell along the white prediction line (Figure 3C). However, when reconstruction matrices were generated by comparing trials recorded for track position α along the LRL route to trials recorded for track position β along the LRL route, the black reconstruction line did not follow closely along the white prediction line (Figure 3D).

Reconstruction error was calculated by subtracting the black line from the white line and averaging the difference for each of the 200 bins along the track. Reconstruction errors were significantly higher for α/β reconstructions, compared with $\alpha\text{Even}/\alpha\text{Odd}$ and $\beta\text{Even}/\beta\text{Odd}$ reconstructions (Figure 3E), demonstrating that the network can distinguish similar routes that occur in different allocentric positions (Kruskal-Wallis, $\chi^2(5) = 750.96$, $P = 4.68 \times 10^{-160}$, d.f. = 5). This demonstrates that the network was also

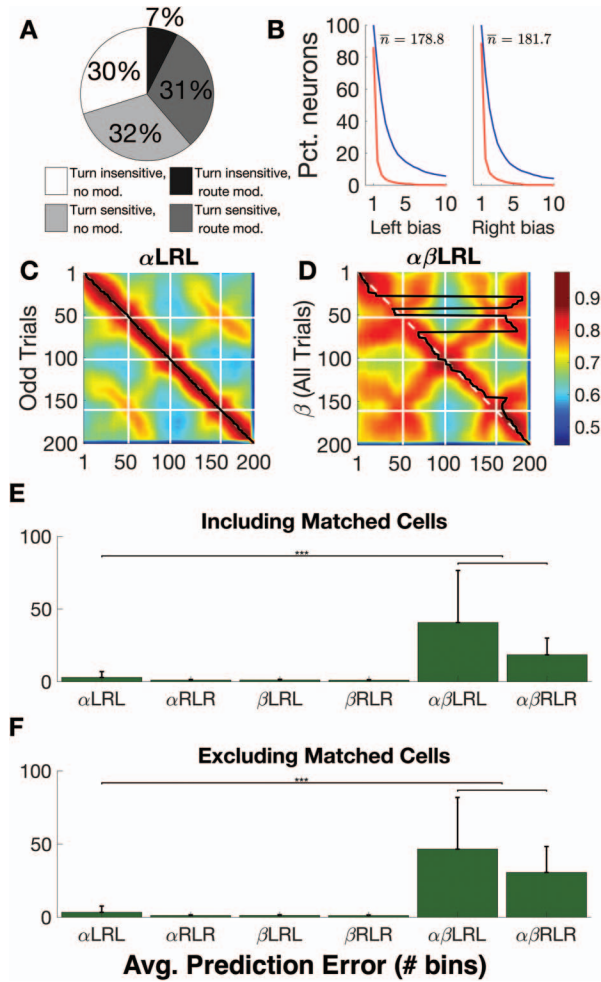


Figure 3. (A) Pie chart showing the average neuronal type distributions for turn sensitive and turn insensitive neurons, similar to those reported in Alexander and Nitz (2015). We add a segment for turn insensitive, route modulated neurons. (B) Random (red/lighter gray line) and nonrandom (blue/dark gray line) cumulative distribution functions (cdfs) are shown for the turn bias ratios associated with all turn sensitive neurons in the population (averaged over all spiking neural networks [SNNs]). (C) A positional ensemble reconstruction matrix that shows the correlations for ensemble neural activity between even and odd trials for the α LRL route, resulting in low reconstruction error. Adapted from Rounds et al. (2016). (D) A positional ensemble reconstruction matrix showing the correlations for ensemble neural activity between all trials associated with the α LRL and β LRL routes. Adapted from Rounds et al. (2016). (E) Bar plot showing average reconstruction error for all positions and routes (arbitrary units). Average error was extremely low (near zero) for reconstructions between even and odd trials for the same position and route (first four bars on plot), while average error was comparatively quite high for reconstructions between tracks α and β for both the LRL and RLR routes (fifth and sixth bars on plot). Significance bars show that reconstruction error between α and β was significantly higher for both routes than reconstruction error between same positions and routes at $p < .0001$ (as revealed by a Kruskal-Wallis test and Tukey's honestly significant difference [HSD] post hoc comparison). Asterisks denote level of significance (***) = $p < .001$). (F) Bar plot showing average reconstruction error for all positions and routes after explicitly matched synthetic neurons were removed from the population. The remaining neurons were also successfully trained to capture the features and behavior of the experimental dataset. See the online article for the color version of this figure.

sensitive to position within the larger environment (a.k.a., the allocentric frame of reference) in addition to a route-centric frame. This finding is consistent with the original result (for further details, see Alexander & Nitz, 2015). Removing the explicitly matched neurons from the analysis did not significantly impact network performance on positional reconstructions (Kruskal-Wallis, $\chi^2(5) = 755.22$, $P = 5.62 \times 10^{-161}$, d.f. = 5; Figure 3F) or on other metrics such as neuronal type distributions or turn bias ratios (not shown).

Control Experiments Demonstrate Necessity of STDP and Training

To assess the contribution of STDP and the evolutionary algorithm to the framework's performance, we ran two control experiments. In the first experiment, 10 networks were initialized and underwent a single generation of testing with no training using STDP. In the second experiment, the same procedure was used, except the population underwent one training session with STDP enabled followed by testing. Because each population was examined after a single generation, neither group was exposed to the evolutionary algorithm process. We measured fitness and additionally compared neuronal type distributions and reconstruction errors with those produced by the fully evolved networks (50 generations of evolution with both training and testing). In both conditions, fitness was poor compared to networks that underwent 50 generations of training and testing (Table 1). In the No STDP condition, the average fitness score was $F = -108.43 \pm 265.44$ because many networks had average maximum firing rates per neuron that exceeded the maximum firing rate threshold (Eqs. 4 and 5). When considering only the sum of the correlations without subtracting the maximum firing rate error, $F = 58.98 \pm 22.47$ (average Pearson's $R \cong 0.24$ per neuron). In the first-generation STDP condition, fitness showed a dramatic increase ($F = 80.5 \pm 12.47$, or Pearson's $R \cong 0.33$ per neuron), but correlations between synthetic and experimental neuron firing patterns were weak compared with those seen in fully evolved networks ($F = 105.93 \pm 0.91$, or Pearson's $R = 0.43$ per neuron).

In terms of our baseline metrics for comparison, we found that neuronal type distributions differed greatly from those expressed by the evolved networks and by the experimental data (Figure 4). In the No STDP condition, we found that $1.18\% \pm 0.98\%$ ($\bar{n} = 7.1 \pm 5.88$) of neurons in each best fit individual could be classified as head direction cells. Of the cells that could not be classified as head direction neurons, $72.48\% \pm 21.63\%$ ($\bar{n} = 429.9 \pm 129.18$ out of 592.9 ± 5.88) were turn insensitive (firing rates associated with left vs. right turns for all neurons, Wilcoxon's rank sum test, $P \geq 0.05$), with the majority ($53.26\% \pm 27.08\%$, $\bar{n} = 316.5 \pm 161.83$ of the total population; Figure 4, top left) being turn insensitive but exhibiting route modulation, although this varied considerably across networks. Neurons were considered route modulated if the correlations between neuronal activation patterns between same routes but different allocentric track locations (e.g., α LRL vs. β LRL) had Pearson's R correlation scores >0.4 . There were $19.22\% \pm 15.43\%$ ($\bar{n} = 113.4 \pm 90.32$) cells that were not modulated by route. The remaining neurons were considered turn sensitive (firing rates associated with left vs. right turns for all neurons, Wilcoxon's rank sum test, $p < .05$), equally divided between route modulated ($11.69\% \pm 15.82\%$, $\bar{n} = 69.2 \pm 47.18$)

Table 1
Average Fitness Scores for No STDP, STDP (First Gen), and STDP (50th Gen)
Experimental Conditions

Condition	No STDP	STDP (first gen)	STDP (50th gen)
Mean fitness	-108.43/265.44 (sum corr. only: 58.98/22.47)	80.5/12.47	105.93/91

Note. STDP = spike timing dependent plasticity.

first left (right) turns versus second left (right) turns in a route, Wilcoxon’s rank sum test, $p < .05$); and not route modulated ($15.82\% \pm 14.72\%$, $\bar{n} = 93.8 \pm 87.57$) first left (right) turns versus second left (right) turns in a route, Wilcoxon’s rank sum test, $P \cong 0.05$).

In the first-generation STDP condition, a similar pattern emerged: $69.74\% \pm 16.12\%$ ($\bar{n} = 392.6 \pm 82.23$) of the neurons were turn insensitive (Figure 4, bottom left; Wilcoxon’s rank sum test, $P \cong 0.05$), with the majority again being turn insensitive with route modulation ($47.81\% \pm 27.07\%$; $\bar{n} = 266.0 \pm 146.86$ of the total population). The rest of the turn insensitive neurons were not modulated by route ($21.93\% \pm 13.62\%$; $\bar{n} = 392.6 \pm 82.23$). The remainder of the neurons evaluated were turn sensitive (Wilcoxon’s rank sum test, $p < .05$). On average, $15.42\% \pm 8.3\%$ ($\bar{n} = 88.5 \pm 50.94$) were route modulated and $14.84\% \pm 8.21$ ($\bar{n} = 85.5 \pm 50.08$) were not route modulated. There were comparable numbers of head direction cells present in first-generation STDP networks as in fully evolved networks, $5.58\% \pm 4.02\%$ ($\bar{n} = 33.4 \pm 24.14$), although percentages varied less in the latter case.

Although reconstruction errors were low comparing odd and even trials with No STDP and first-generation STDP, the discriminability of these networks was worse than experimental data and the 50 generations of STDP networks. Average reconstruction errors of position from ensemble activity patterns were very good in both the No STDP (Kruskal-Wallis, $\chi^2(5) = 747.76$, $p = 2.3 \times$

10^{-159} , d.f. = 5) and first-generation STDP (Kruskal-Wallis, $\chi^2(5) = 744.81$, $p = 1.00 \times 10^{-158}$, d.f. = 5) conditions (Figure 4, top and bottom right), and followed the same pattern observed in fully evolved networks as well as the experimental data.

However, the reconstruction matrices qualitatively differed considerably between one another (Figure 5A–D). In the No STDP condition peak correlations still occurred at same-bin positions for even versus odd trial positional reconstructions even though correlations at all positions were universally high with very low variation by distance (Figure 6). A similar pattern emerged in the first-generation STDP condition, although correlations were slightly more variable and neuronal firing patterns were relatively differentiable by position and action correlates than in the No STDP condition. To demonstrate the influence of evolution, we plotted the average correlation as a function of bin distance from the center ‘perfect prediction line’ for an even versus odd trial positional reconstruction matrix over all best-fit networks of each run for each condition. In the No STDP case, correlations are relatively steady instead of declining with distance (Figure 6, blue/dark gray line), and do not fall below Pearson’s $R \cong 0.9$ before near the end of the trial, suggesting that different spatial locations are not likely to be differentiable, especially in conditions in which there is a great deal of noise. In the first-generation STDP case, correlations drop somewhat more steeply compared to the No STDP case until the bin distance from the center point is about 40,

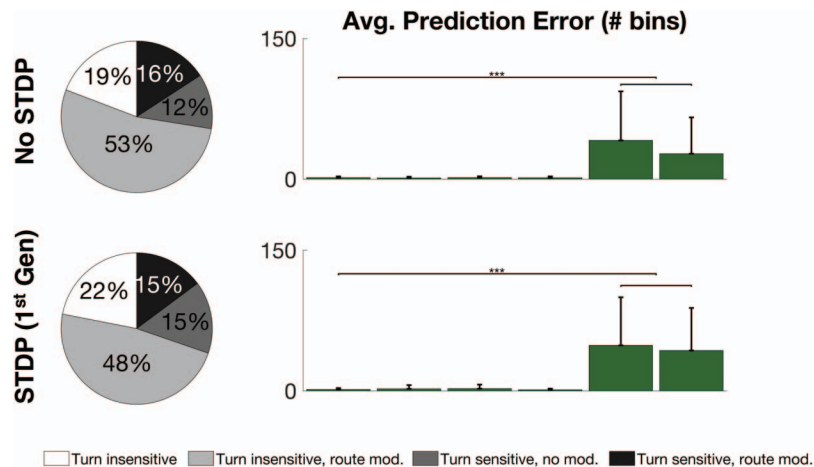


Figure 4. (Top) Neuron type distributions and average reconstruction of position errors for the No spike timing dependent plasticity (STDP) condition. The majority of the neurons in the population are turn insensitive. Average reconstruction errors approximate those seen in the experimental data and evolved networks. (Bottom) Neuron type distributions and average reconstruction of position errors for the STDP (first generation) condition. Most neurons in the population are also turn insensitive. Average reconstruction errors show very good discrimination between α and β track locations and again approximate those seen in the experimental data and evolved network populations. *** $p < .001$. See the online article for the color version of this figure.

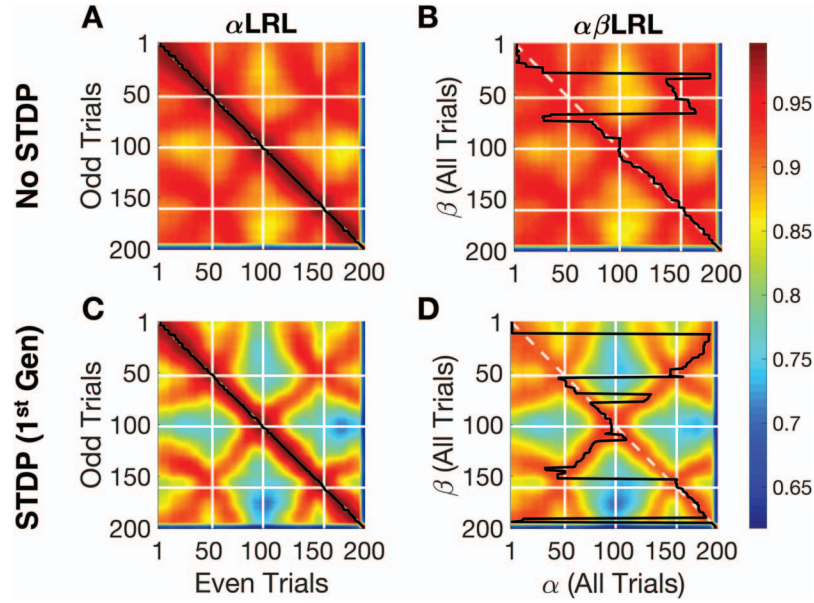


Figure 5. (A) Averaged ensemble correlation matrices over 10 best-fit networks after testing without spike timing dependent plasticity (STDP) or evolution for even versus odd trials corresponding to the α LRL track position and route. (B) Averaged ensemble correlation matrices from the same networks for all α LRL versus all β LRL trials. (C) Averaged ensemble correlation matrices over 10 best-fit networks after training (with STDP) and testing but no evolution. (D) Averaged ensemble correlation matrices from the same networks for all α LRL versus all β LRL trials. See the online article for the color version of this figure.

and then drops very gradually as the distance approaches 80 (Figure 6, red/medium gray line). However, the average correlation as a function of bin distance does not fall below Pearson’s $R \cong 0.8$ before near the end of the trial, suggesting that disparate spatial locations are not much more differentiable than in the No STDP case. In the 50 generations of STDP case (fully evolved networks), correlations drop steeply to Pearson’s $R \cong 0.6$ until the bin distance approaches 80, and closely matches the experimental data (Figure 6, green/light gray and black dotted line, respectively). Taken together, these control experiments show that multiple generations of mutating STDP parameters were necessary to achieve good matches to the experimental RSC data in terms of distributions of neuron activity and discrimination of allocentric space.

Evolved SNNs Are Robust to Loss of Input Streams

To investigate how different input streams might affect RSC responses, we eliminated specific inputs into the evolved networks and tested model fitness. With the exception of allocentric position, the evolved networks were highly robust to lesions of single input streams, suggesting that RSC function in intact animals cannot overcome lesions to the hippocampus whose neuron ensembles accurately encode environmental location. Significant decreases in fitness were observed when two or more input streams were lesioned (Kruskal-Wallis, $\chi^2(8) = 70.66$, $p = 3.63 \times 10^{-12}$, d.f. = 8; Figure 7A). Combined angular velocity and head direction lesions had about the same impact as lesions eliminating environmental position, and, in general, fitness decreased with the number of input lesions.

We also investigated how different intrinsic connections might affect RSC responses. The evolved networks were relatively insensitive to the loss of recurrent excitatory connections. However, lesions to inhibitory connections had a strong effect on fitness ($F = 40.37 \pm 21.46$), primarily because inhibitory connections stabilized the neuronal firing rates and this loss of stability impacted the correlations between synthetic and experimental neural activity. The average firing rate per neuron in unlesioned networks was only 4.41 Hz, with a maximum of 72.07 Hz, while, when inhibitory connections were lesioned, the average firing rate was 181.36 Hz, with a maximum of 481.56 Hz. The result is consistent with the known powerful impact of inhibitory networks on both network stabilization (bounded responsiveness of excitatory elements) and fidelity of representation (Froemke, 2015). It should be noted, however, that most cases in which at least one input was lesioned resulted in much higher firing rates than normal, and the impact on firing rate in the HD lesion case was about on par with that of the inhibitory neuron lesion case. Thus, it must be noted that a major impact on network fitness results from disruptions in constraints on firing rates (Table 2).

To test whether the impact of lesioning individual inputs might be due to the amount of input coming into the network, we quantified the amount of synaptic input to the network over 100 trials (for a total of 20,000 time steps). This was calculated for each input group by multiplying the number of spikes from each input neuron by the synaptic weight of each of its connected output neurons. Synaptic input was then averaged over all 10 networks. The total average amount of synaptic input was 0.103. Angular velocity made up $12.12\% \pm 20.15$ (0.013 ± 0.02) of that input,

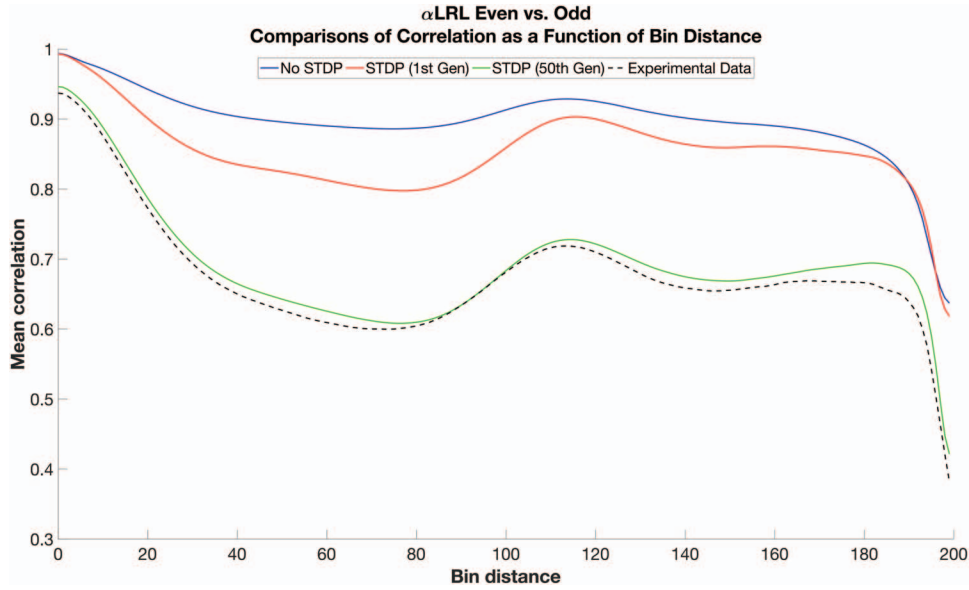


Figure 6. A plot of average correlation value as a function of bin distance from the theoretical perfect prediction line in all ensemble correlation matrices associated with the α LRL (even vs. odd trials) track location and route. Correlations were averaged over all 10 best-fit networks in each experimental condition (No STDP, spike timing dependent plasticity (STDP) first generation, and STDP 50th generation). The average correlation as a function of distance is also plotted for the experimental data (black dotted line), which is best approximated by the evolved networks (green/light gray line). See the online article for the color version of this figure.

linear velocity made up $7.47\% \pm 16.96$ (0.008 ± 0.018), head direction made up $21.51\% \pm 28.64$ (0.022 ± 0.03), and position made up $58.89\% \pm 24.13$ (0.061 ± 0.025). These results suggest that the magnitude of the impact of each input lesion does corre-

spond to the amount of synaptic input coming into the network from that group. However, STDP seems to work by balancing the relative strengths of the input until each is weighted appropriately such that synthetic neuronal firing profiles resemble those ob-

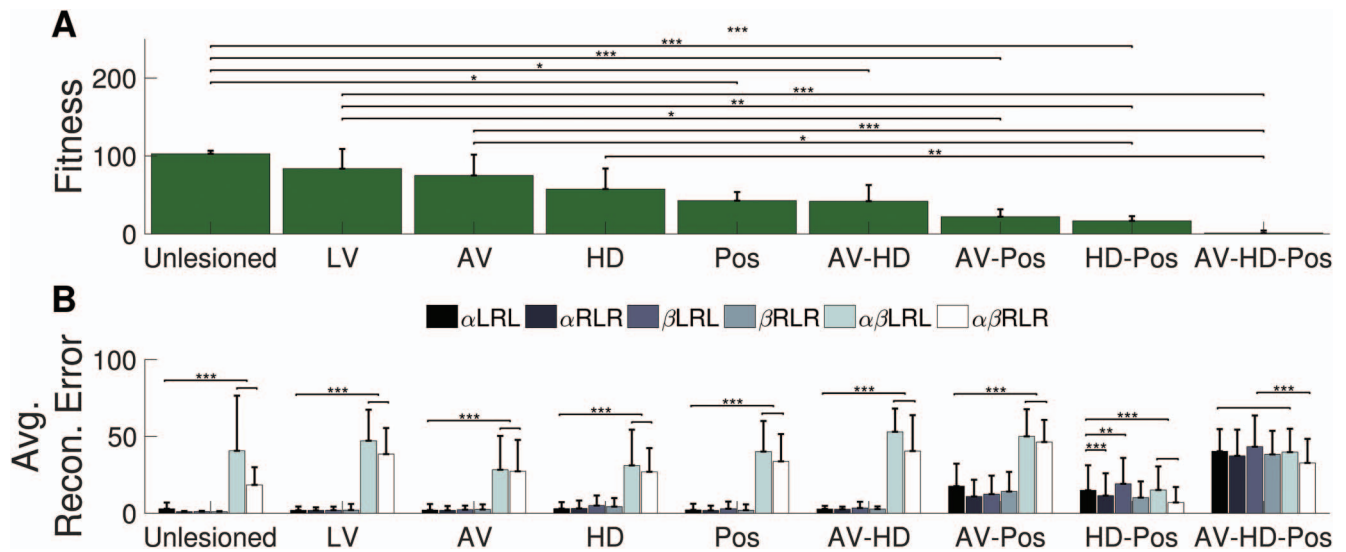


Figure 7. (A) Fitness, as measured by the values of correlations between matched synthetic and experimental neural activity, declines with the number of lesions. Fitness is near zero when all critical inputs are absent from the network (AV-HD-Pos). (B) Bar plots of reconstruction error associated with each lesion case. When head direction (HD) and Pos are lesioned conjunctively, the models lose sensitivity to the allocentric frame of reference. When Angular velocity (AV), HD, and Pos are all lesioned conjunctively, the network loses the ability to reconstruct position within the route. * $p < .05$. ** $p < .01$. *** $p < .001$. See the online article for the color version of this figure.

Table 2
Average Firing Rates (Hz) as a Function of Lesion Type

Lesion type	Mean	Max
Unlesioned	4.4 ± 2.74	72.27 ± 57.76
AV	56.93 ± 122.3	187.62 ± 257.45
HD	128.65 ± 198.05	380.27 ± 431.44
Pos	51.42 ± 151.47	166.96 ± 297.55
AV-HD	136.28 ± 205.82	387.77 ± 440.16
AV-Pos	132.49 ± 210.11	330.64 ± 426.85
HD-Pos	140.61 ± 223.1	371.38 ± 418.16
AV-HD-Pos	134.12 ± 203.76	341.47 ± 437.2
Exc	4.19 ± 2.73	72.62 ± 58.59
Inh	181.36 ± 223.99	481.57 ± 401.2

Note. Shown are the mean firing rates associated with each lesion combination (firing rates were averaged over all neurons and all ten networks) with standard deviations.

served in the experimental data. We compared the average synaptic input from each group with the proportion of input neurons that make up that group (Figure 8) and found that angular velocity is scaled up by 10.44%, linear velocity is scaled up by 4.59%, head direction is scaled up by 19.59%, and position is scaled down by 34.63%. Together with the lesion results, this suggests that the RSC is not solely dependent on any one type of input, and it may be predictive of the relative balance of inputs to this brain area.

Despite the remapping of firing rates, the simulated RSC population could still accurately reconstruct the agent’s position when allocentric position of a given route was consistent. Of the lesion cases examined, only when HD and Pos were lesioned together did the reconstruction errors for α/β fall significantly such that they could not distinguish the difference in the two track positions. α/α and β/β (Kruskal-Wallis, $\chi^2(5) = 89.99$, $p = 6.76 \times 10^{-18}$, d.f. = 5; Figure 7B). By comparison, reconstruction error subsequent to HD and Pos lesions rose more moderately, suggesting that the system cannot effectively disambiguate the environmental locations of routes in the absence of information input streams pro-

viding either head direction and environmental location. This may be related to the finding that some head direction cells in the RSC also conjunctively encode place (Jacob et al., 2017).

On the other hand, when AV, HD, and Pos inputs were lesioned jointly, the network could no longer accurately reconstruct position, yielding universally high reconstruction errors. The reconstruction between α and β for the RLR route was significantly different from all other groups (Kruskal-Wallis, $\chi^2(5) = 33.42$, $p = 3.1 \times 10^{-6}$, d.f. = 5; Figure 7B). Taken together, these simulated lesion experiments predict that conjunctive coding makes the RSC robust to elimination of individual inputs and flexible enough to generate coherent spatial representations when pairs of some input streams are impacted.

Rearranging Turn Sequences Does Not Impact Network Performance, but Does Cause Functional Remapping

To predict how the RSC coding might respond to the imposition of novel routes and paths, we created three new simulated routes that gave rise to six new action sequences: LLL (outbound) and RRR (inbound; Figure 9A), LLR (outbound) and LRR (inbound; Figure 9B), and RLL (outbound) and RRL (inbound; Figure 9C). Remarkably, we found that the networks performed similarly compared with the original LRL/RLR case, in that the progression along the route could be accurately predicted by RSC population activity between even and odd trials associated with the same track position and trajectory (Figures 5D, 5F). The result suggests that the RSC in intact animals is configured to provide first-trial representation of route space, environmental space, and actions when novel navigational contexts are encountered. Significant differences emerged between comparisons of even versus odd trials and comparisons of trials from tracks α and β (Figure 9E–F). Again, the Kruskal-Wallis test indicated the presence of significant differences and Tukey’s HSD was used to determine which groups were different. (LLL/RRR: Kruskal-Wallis, $\chi^2(5) = 764.45$, $p =$

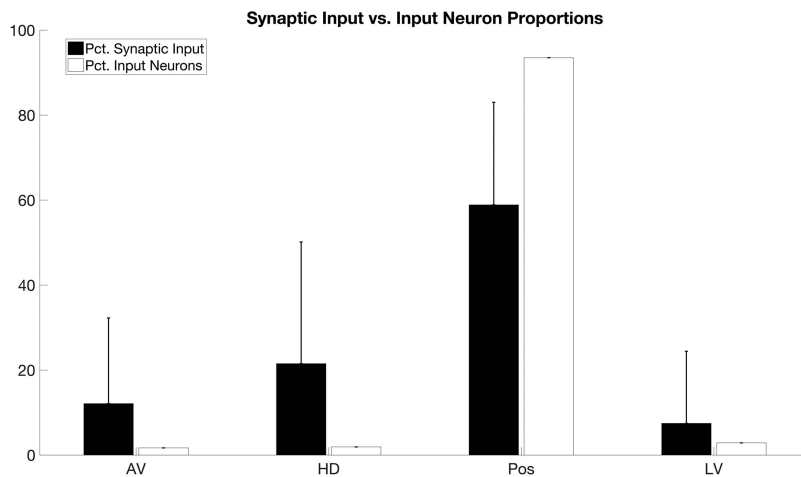
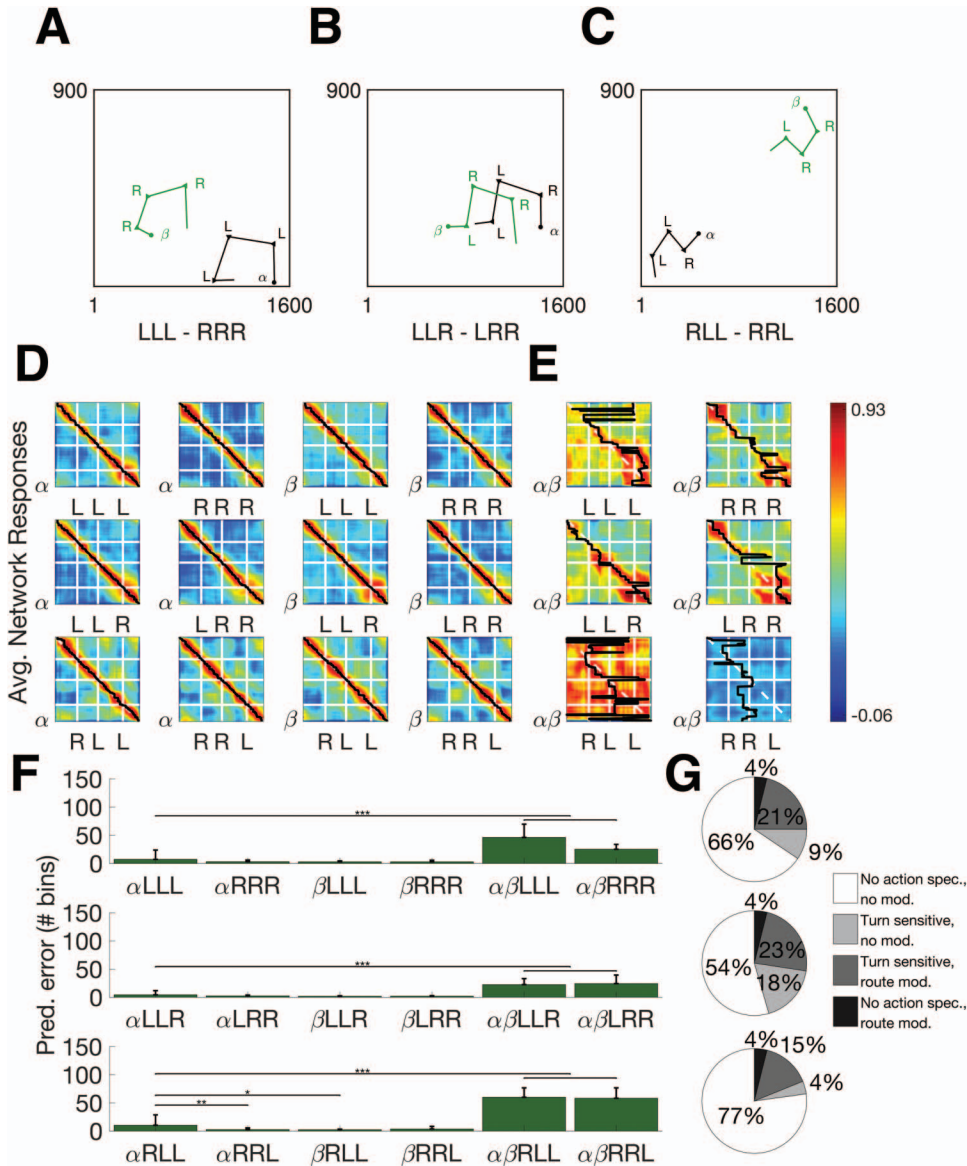


Figure 8. The average synaptic input from each input group to the evolved networks was compared with the corresponding proportion of each group’s input neurons. Angular velocity (AV), head direction (HD), and linear velocity (LV) were scaled up compared with their raw proportions of input neurons, while position (Pos) was correspondingly scaled down.



5.66×10^{-163} , d.f. = 5; LLR/LRR: Kruskal-Wallis, $\chi^2(5) = 773.38$, $p = 6.63 \times 10^{-165}$, d.f. = 5; RLL/RRL: Kruskal-Wallis, $\chi^2(5) = 775.9$, $p = 1.89 \times 10^{-165}$, d.f. = 5). Furthermore, the neuron type distributions were similar to those reported in Alexander and Nitz (2015; Figure 9G).

We again measured network fitness by correlating the synthetic firing rate profiles associated with the new route shapes with the experimentally recorded firing rate profiles associated with the original route shapes in order to see whether each neuron's representation of the new route would still be correlated with its representation for old routes as a way to judge remapping. Despite the rearrangement of turn sequences, the fitness levels of the networks were approximately the same as under the original track conditions when new matches between synthetic and experimental neurons were allowed (Table 3, Unforced match). Because fitness is based on how well synthetic neural activity correlates with experimental neural activity, and because these turn sequences were not investigated experimentally, the biological neurons could not have encoded those specific routes and turn sequences. However, when the synthetic neurons were forced to match the experimentally observed neurons recorded under original route conditions, fitness was very low (Table 3, Forced match). This suggests that although individual neuronal firing rate patterns remap when features of the environment (position, turn sequence, etc.) are altered, the simulated RSC population can flexibly and rapidly remap to novel environments. This suggests that the biological RSC supports flexible encoding of specific features while retaining information about the broader environmental context of the room. Note that measures of fitness are unrelated to either the quality of reconstructions of position within a route since reconstruction depends on ensemble activity across the population.

Changing Allocentric Position and Orientation of the Tracks

We next systematically moved the track, either by linearly translating it across the environmental space or by rotating the track around its center. These manipulations were intended to test the effects of perturbing the allocentric and egocentric frames of reference on RSC responses. We ran five sessions of each condition in which the tracks occupied different starting positions (Orientation: Figure 10A–E; Translation: Figure 10F–J). We also ran additional trials in which the track at position α was rotated 180 degrees so that it was in the opposite orientation and then linearly translated across the space.

Table 3
Fitness Scores for Resequenced Routes Under Different Matching Conditions

Turn sequence	Unforced match	Forced match
LLL/RRR	95.03 ± 3.97	$5.82 \pm 9.36 \times 10^{-16}$
LLR/LRR	95.38 ± 4.42	$10.24 \pm 1.87 \times 10^{-15}$
RLL/RRL	94.46 ± 5.66	$5.79 \pm 9.36 \times 10^{-16}$

Note. When network responses to novel route sequences were correlated with the same biological neuron to which it was originally matched, fitness was low. However, if neurons were assigned a new match, fitness was high.

We found that changing the position or orientation of the track resulted in neuron remapping, as can be observed by the drop in fitness scores at each new position (Figure 10K, 10L). In the orientation case, the fitness reached an asymptote at approximately 10 arbitrary units (a.u.), which corresponds roughly to Pearson's $R = 0.20$, after the track had been rotated 80 degrees (Figure 10K). This drop-in fitness, when the track was rotated around its center point, was more abrupt than when the track was simply translated across allocentric space (Figure 10L), evidencing the powerful effect of head orientation as a variable to disambiguate navigational contexts. During translation, network fitness fell somewhat more gradually to reach an asymptote at approximately 20 a.u. (approximately $R = 0.10$ at about 160 cm from the original position). We plotted lines of best fit associated with each change in fitness across nine positions for translation and rotation (Figure 10K, 10L) to demonstrate the influence exerted over neural activation when two input variables are changed (e.g., sharper decline in fitness when the track is rotated) as opposed to when only one input variable is changed (e.g., gradual decline in fitness when the track is translated). This may suggest a conjunctive coding scheme in the retrosplenial cortex by which novel route representations are generated.

To examine how the networks represent space and how these representations are grounded in both external and internal frames of reference, we calculated the average reconstruction errors, based on the positional ensemble reconstructions, between all possible locations of the track to see how well the network could predict position within route when routes were in orthogonal, parallel, or juxtaposed orientations, or when the routes overlapped in allocentric space for each selected recording session individually (for clarity, relevant examples are shown in Figure 11, but similar results were obtained for all sessions). We calculated the reconstruction errors via positional ensemble correlations for all 18 positions covering the complete 360-degree orientation of the track, which were then averaged and visualized using a heat map. The values were superimposed over a polar plot (Figure 11). Each value in the plot is the averaged reconstruction error for track α versus β for four example sessions of trials in which the track was rotated. Each slice of the polar plot corresponds to the rotation of α 's track position (0 degrees is the original position of α). Inside each slice are 18 radial bars that represent the average reconstruction error between α and β for each orientation of β (the bar at the center is the original position of track β , and the outermost bar is the position of β at 340 degrees).

The networks were most sensitive to rotated routes when allocentric positions closely overlap. For this particular recording session, the tracks' center points occupy the same approximate allocentric position. This means that as each track is rotated, there are points at which they occupy both the same orientation and position in space, and are thus the same route. One would expect that reconstruction error should be low at these points, because this is analogous to comparing even and odd trials from the same track position and route. This would give rise to a spiral of low reconstruction errors on the polar plot in which the positions and orientations systematically match. This expectation was confirmed (Figure 11A). To highlight this, we circle the location of the minimum value of the reconstruction error on the polar plot in black and plot the relative track positions of α and β at that point (Figure 11B, left). We similarly highlight the maximum recon-

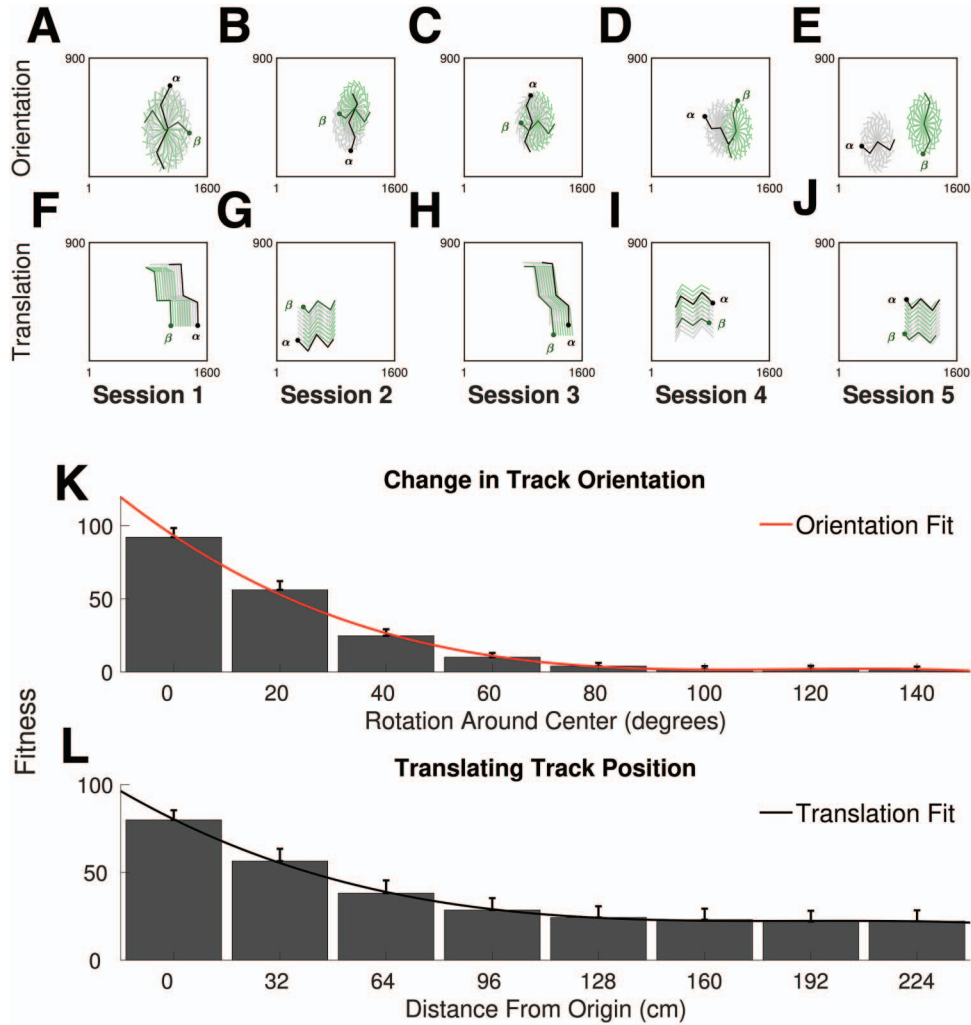


Figure 10. (A–E) We manipulated all trials in five separate recording sessions to simulate track rotation in 20° increments around the center point of the track. (F–J) We manipulated the behavioral position data associated with a separate five recording sessions to simulate movement of each track position linearly across the allocentric space in 32 cm increments. This manipulation did not require altering head direction. These experiments were repeated with one track (α) rotated 180° such that tracks α and β occupied juxtaposed orientations. (K) Track orientation causes neuronal remapping (revealed by changes in network fitness). Fitness drops sharply between the first, second, and third orientations, but reaches an asymptote of approximately 10 a.u. by the fourth position. A best-fit line is plotted along with the bar plot to describe the change in fitness. (L) Track translation also causes neuronal remapping, as revealed by changes in network fitness at each position. Fitness declines steadily when the track is moved linearly across the space as compared to the case in which orientation is changed. See the online article for the color version of this figure.

This document is copyrighted by the American Psychological Association or one of its allied publishers. This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

struction error value by circling that point in red (dark gray) and plotting the relative track positions and positional ensemble reconstruction matrix (Figure 11B, right). This demonstrates that the network is sensitive to the route-centric frame of reference and the allocentric frame of reference because the network recognizes track positions that occur at similar positions and orientations in space, but is highly discriminative of the routes when they are in opposite orientations but share a similar allocentric position.

When the tracks are further apart in allocentric space such that the arms of the track can overlap at some positions, but are not fully overlapping, these spiral lines disappear, in further support of

the hypothesis that the networks do not form a representation of the track on the basis of route alone. Instead, the networks are very good at creating representations that reflect and enhance the differences between ambiguous routes that partially overlap in allocentric space based on conflicting input streams (Figure 11C). Interestingly, when the track positions occupy analogous orientations but do not overlap in allocentric space, reconstruction error is relatively low, suggesting that the differences in representation between the routes do not need to be enhanced once the routes are far enough apart in allocentric space that there is no risk of ambiguity due to points of overlap between them. This is shown by

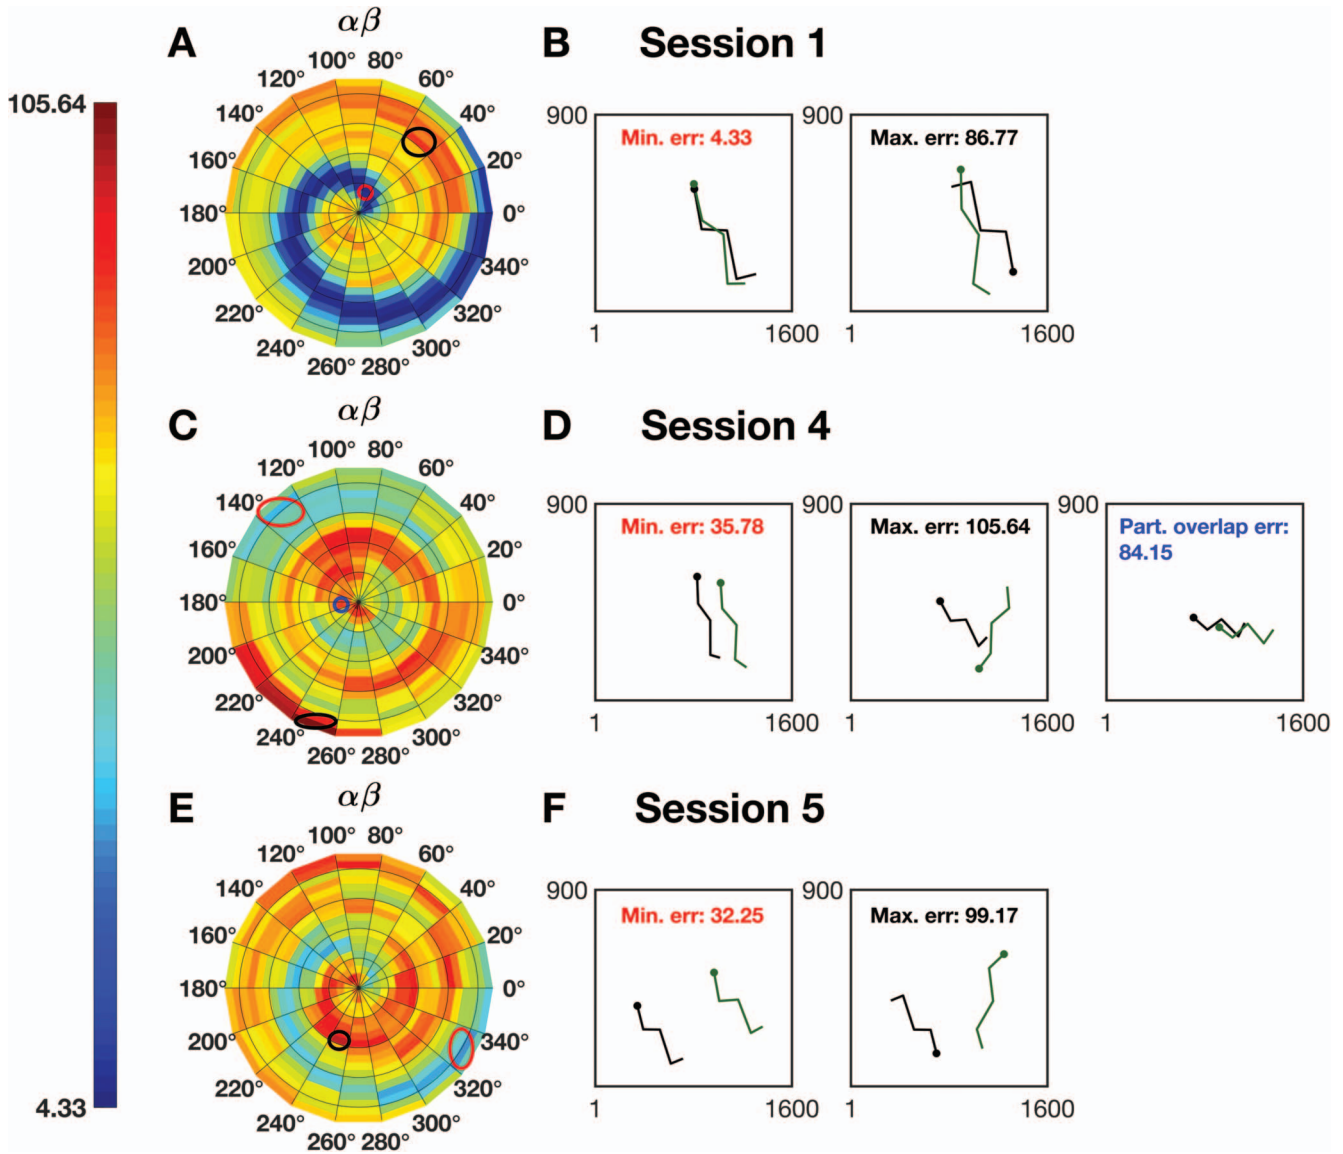


Figure 11. (A) When tracks α and β closely overlap during rotation, the average reconstruction error for correlations between neural activity associated with each track is low, indicated by a line of lowest error (blue/gray values in polar plot). (B) When tracks are in similar positions, average reconstruction error is low (left, track locations with minimum error; red circle/dark gray circle in polar plot). When the tracks are in opposite orientations, average reconstruction error is high (right, see track locations with maximum error; black circle in polar plot). (C) When the center points of the tracks do not overlap in space, the tracks cannot overlap completely but may be parallel and partially overlap at various points. Due to heightened conflict between variables, a spiral line is no longer discernible. (D) When the tracks are parallel and nonoverlapping, error is lowest (left). When the tracks are nonoverlapping but occupy opposite orientations (middle), error is highest. When the tracks are parallel but do not fully overlap (last third of α , containing the second left turn, overlaps with the first third of β , containing the first left turn), there is high behavioral conflict, resulting in high average reconstruction error (small blue circle/small gray circle in polar plot). (E) When the tracks are sufficiently far apart that their positions never overlap as they undergo rotation, the spiral line of lowest error is discernible. (F) When the tracks are parallel but occupy different locations, average error is low (left). When they occupy opposite orientations, average reconstruction error is high (right). See the online article for the color version of this figure.

the point of lowest error circled in red (dark gray) on the polar plot and the corresponding track positions (Figure 11C–D, left). On the other hand, when the track positions occupy opposite orientations with minimal overlap, reconstruction error is relatively high (Fig-

ure 11C–D, middle). Finally, when the tracks are in the same orientation and their allocentric positions overlap closely but the routes are not aligned, reconstruction error is high, suggesting that the behavioral variables at these two positions are in conflict such

that the network representations of each are maximally different (Figure 11C–D, right). This conclusion is consistent with the track translation experiments where head direction and allocentric position were in conflict between routes (Figure 12). In this case, the network’s holistic representations for each route were highly dissimilar, resulting in high reconstruction error.

When positions do not overlap in allocentric space, network sensitivity to similar route orientations is reestablished. That is, when the allocentric positions of the tracks are entirely nonoverlapping, these spirals of low reconstruction errors are again discernible (Figure 11E), suggesting that the networks remain sensitive to routes that occupy the same or opposite orientation, but

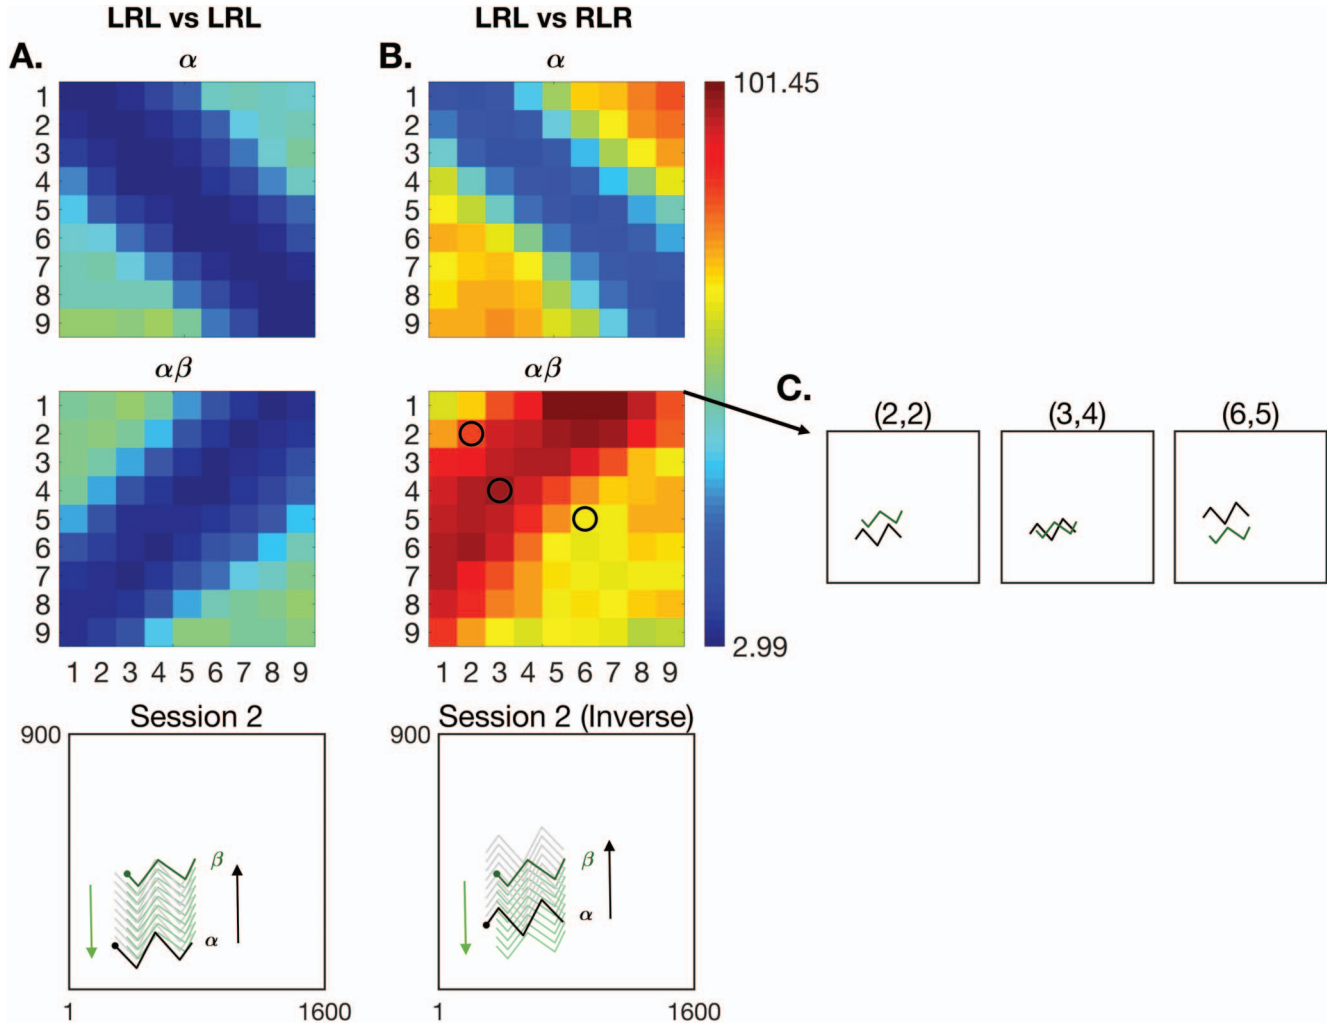


Figure 12. (A) When reconstructing the agent’s position on track α using even and odd trials from the same session, reconstruction error is low. When comparing activity from different positions of track α , error is higher, scaling with the distance between them (top). This pattern is mirrored in reconstructions on track β from α (middle). The original track locations are shown in black and dark green/dark gray, respectively. For reference, all positions of both tracks are shown as they were translated across allocentric space (bottom). (B) When reconstructing position on track α for opposite turn sequences (routes LRL and RLR), reconstruction error is lower when the tracks occupy the same allocentric position versus when they do not (top, blue/dark gray values). Only head direction differed dramatically between the routes, suggesting that the conflict in heading is responsible for the higher error. When reconstruction error was measured for trials where tracks occupied juxtaposed orientations, it scaled with the amount of conflict between the variables associated with each track (middle). For reference, the positions of the two tracks (with inverted α) are shown as they were translated across the environment (bottom). (C) We show three illustrative examples of variable conflict as it relates to reconstruction error. Error is high when the tracks occupy opposite orientations and are near one another, but not overlapping (left). It increases when the tracks overlap in space while occupying opposite orientations (middle). Variable conflict declines as the tracks move away from one another in opposite directions, resulting in reduced error (right). See the online article for the color version of this figure.

entirely different allocentric position, in space. Although reconstruction errors are overall higher due to the distance between the track positions, this demonstrates a coherent and sophisticated representation of the route-centric frame of reference that is contextualized by the allocentric frame of reference (Figure 11E). As before, when the tracks occupy analogous orientations, reconstruction error for the session is at its lowest at a value of 32.25 arbitrary units (Figure 11F, left), while the error is at its highest (a value of 99.17 arbitrary units) when the tracks occupy opposite orientations (Figure 11F, right).

Neurons Undergo Functional Remapping

Finally, we found that neuronal responses changed over the course of trials in which tracks were translated linearly or rotated in allocentric space, and could be reassigned to different functional categories than the one to which it previously belonged at the track's initial position. This was calculated by statistically evaluating the neuron's firing rate profile at each new track position and averaging the number of times it remapped over the course of all manipulated trials. Consistent patterns of remapping emerged across the two manipulations: Turn insensitive neurons not modulated by route were most likely to remap to turn sensitive neurons, but were unlikely to be route modulated following remapping. Turn insensitive neurons whose activity profiles were modulated by route, on the other hand, typically retained their insensitivity to turns, but lost their route modulation. Turn sensitive, not route modulated neurons tended to remap in both regards and became turn insensitive and route modulated. Turn sensitive neurons that were route modulated, in contrast, often retained their turn sensitivity but lost their modulation by route. These patterns were consistent across both linear translation and rotation (Table 4).

It was unusual for turn sensitive neurons to change their preferred type of turn instance (i.e., neurons that preferred left [right] turns continued to respond to left [right] turns if they retained their turn sensitivity). Interestingly, neurons remapped frequently on an individual basis; however, the overall distributions of each functional category across the population remained relatively stable (the maximum shift for each category never surpasses 6.8% for the orientation case and 11.5% for the translation case).

Examples of neurons whose activity profiles demonstrate functional remapping are shown in Figure 13. In one instance, a previously turn insensitive neuron becomes sensitive to left turns

(Figure 13A–B), and in another instance, a previously turn insensitive neuron without route modulation remaps to become route modulated but retains its turn insensitivity (Figure 13C–D). These results suggest, together with analysis of the networks' ability to reconstruct routes under novel conditions in which the track positions are altered, that neurons in the retrosplenial cortex can reliably and flexibly encode multiple task-relevant variables such that their activities may shift to meet changing task demands. This is consistent with the observation that neurons in higher cortical regions may demonstrate category-free coding and/or mixed selectivity to handle more complex cognitive tasks (Eichenbaum, 2017; Raposo, Kaufman, & Churchland, 2014; Rigotti et al., 2013).

Discussion

To gain a better understanding of how the RSC flexibly encodes space, we used a novel method to match neurophysiological data from the rodent RSC to a simulated spiking neural network using behavioral measurements such as allocentric position, head direction, linear and angular velocity as inputs. In a series of virtual experiments where we simulated lesions and changes to the rodent's environment, we showed that the evolved RSC models support: (a) conjunctive encoding of multiple, disparate input streams, consistent with results reported previously in this brain region; (b) resiliency to loss of sensory input streams; and (c) flexible and generalizable encoding at the population level. Lastly, we showed that neuronal firing patterns in the RSC transcend assigned functional categories.

Although randomly initialized SNNs that did not incorporate the evolutionary strategy or training (STDP) could capture some features of the electrophysiological dataset (such as different functional neuron types and significantly higher reconstruction errors for same vs. different allocentric track positions compared with trials compared for the same track position), most of the neurons in these networks had highly correlated patterns of activation across trials and extremely low positional discriminability. This is reflected by the fact that the majority of the neurons in both sets of networks were turn insensitive but route modulated (i.e., were highly correlated by route independent of allocentric track location) and by the plot of correlation as a function of bin distance. These results suggest that training and evolution are necessary for networks to fully capture the descriptive and functional character-

Table 4
Percentages of Neurons That Remap During Track Manipulation Experiments

	Turn insensitive, no mod.	Turn insensitive, route mod.	Turn sensitive, no mod.	Turn sensitive, route mod.
Orientation				
Turn insensitive, no mod.	—	5.66% ± .84	30.3% ± 3.2	7.92% ± 1.02
Turn insensitive, route mod.	45.4% ± 2.25	—	24.01% ± 3.17	7.19% ± 1.2
Turn sensitive, no mod.	3.57% ± .41	39.53% ± 3.2	—	11.65% ± .76
Turn sensitive, route mod.	3.15% ± .49	36.37% ± 4.01	42.81% ± 1.84	—
Translation				
Turn insensitive, no mod.	—	7.44% ± 2.78	20.13% ± 2.17	5.57% ± 1.71
Turn insensitive, route mod.	48.66% ± 4.34	—	12.96% ± 1.9	5.6% ± .95
Turn sensitive, no mod.	3.04% ± 1.75	35.95% ± 3.38	—	10.26% ± 2.11
Turn sensitive, route mod.	3.61% ± 1.61	32.98% ± 3.84	38.11% ± 1.76	—

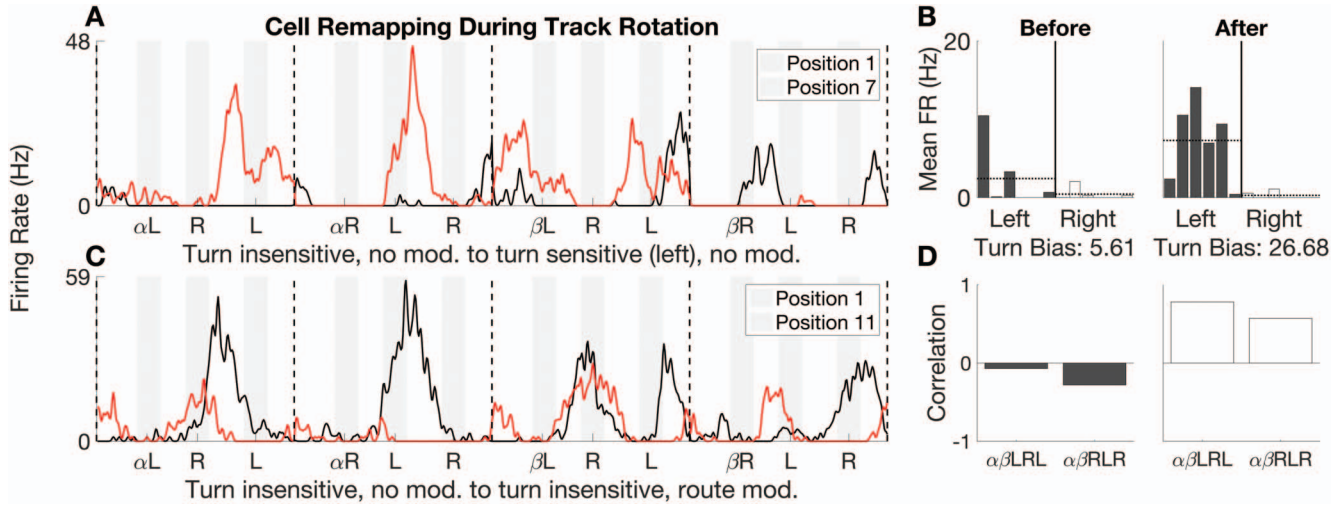


Figure 13. (A) An example neuron whose firing rate pattern was consistent with turn insensitivity at the original track position (black trace), but whose firing pattern changed over the course of track rotation to exhibit turn sensitivity. (B) Bar plots of mean firing rate for all six left (gray bars) and right (white bars) turn instances. The ratio of the neuron’s average activity for left to right turns was 5.61 before remapping, and was 26.68 following remapping. (C) An example neuron whose firing rate pattern is consistent with turn insensitivity both before and after remapping. However, before remapping, the neuron’s activity does not exhibit modulation by route (black trace). Following remapping, the neuron’s activity is route modulated (red/dark gray trace). (D) Bar plots of the Pearson’s *R* correlation coefficient before and after remapping between tracks α and β for both turn sequences (LRL and RLR). Before remapping, the neuron’s activity for the same turn sequences and different track locations was negatively correlated (α LRL vs. β LRL, $r = -0.07$, $p = .35$; α RLR vs. β RLR, $r = -0.28$, $p = 5.59 \times 10^{-5}$). Following remapping, the activity neuron’s activity for those track positions and turn sequences were strongly positively correlated (α LRL vs. β LRL, $r = .78$, $p = 2.1 \times 10^{-42}$; α RLR vs. β RLR, $r = .57$, $p = 9.89 \times 10^{-19}$). See the online article for the color version of this figure.

This document is copyrighted by the American Psychological Association or one of its allied publishers. This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

istics of the electrophysiological data and that can behave in a manner comparable to the biological RSC.

The retrosplenial cortex is well-suited to playing a role in navigation and the generation of routes. It contains an integrated head direction system (Cho & Sharp, 2001; Clark, Bassett, Wang, & Taube, 2010; Taube, 1995) based on input from the anterior thalamic nuclei and targets cells in the entorhinal cortex, which may provide temporal structure to the firing of hippocampal cells (Schlesiger et al., 2015). A computational model of mental imagery suggests a role for the RSC in scene construction underlying imagination and episodic memory which hinges on the transformation of spatial information between different frames of reference (Burgess, Becker, King, & O’Keefe, 2001; Byrne et al., 2007). Because the RSC is situated between the hippocampus and parietal cortex (Robinson, Keene, Iaccarino, Duan, & Bucci, 2011; Todd, Huszár, DeAngeli, & Bucci, 2016; Vann, Aggleton, & Maguire, 2009), it may also support the development and deployment of different kinds of navigation strategies that are tied to different spatial frames of reference (e.g., allocentric, egocentric, and route-centric). A high-level computational model showed that the RSC is capable of choosing and deploying such strategies (Oess, Krichmar, & Röhrbein, 2017).

Several important and specific predictions result from the present experiments. Our simulations of the RSC suggest that new routes through a familiar environment do not need to be learned—if an organism has experience within a broader environmental context, then learning may not be necessary for it to

flexibly adapt to new paths through that environment. This is reminiscent of the schema concept proposed by Tse and colleagues (Preston & Eichenbaum, 2013; Tse et al., 2007, 2011). Specifically, our evolved SNNs adapted flexibly to new routes within allocentric space that it had never seen before during the evolutionary training process, as shown by experiments where we altered the turn sequences associated with specific routes within the simulated room. The networks could reconstruct the agent’s position on the route from neural ensemble activity, but could also disambiguate similar routes that occupied different positions in allocentric space. This may help to explain a phenomenon associated with damage to the RSC (particularly the right RSC) called “topographic disorientation,” a condition in which patients with this kind of damage can recognize familiar landmarks, but can no longer construct or navigate routes between them (Kim, Aminoff, Kastner, & Behrmann, 2015; Takahashi et al., 1997). If the RSC were responsible for planning and/or deploying routes through familiar environments, then damage to the region would predict topographic disorientation. This is also consistent with theoretical results reported in Oess et al. (2017) in which the RSC was used to build a confidence level for the deployment of different kinds of navigation strategies (allocentric, egocentric, or route-centric) based on the kinds of incoming sensory input. Our results may provide a biologically plausible substrate for these navigation strategies.

In contrast, our model is different from Byrne et al. (2007) in that representations of new places or routes in a familiar environ-

ment do not need to be learned. Instead, neuronal activity in the RSC remaps to accommodate new representations of different routes that occupy the same general space. Our model is broadly consistent with the idea that what reference frame is employed by the RSC depends on the sensory inputs to the region, but its activity changes flexibly and immediately without relearning.

We also predict that spatial representations in the RSC are closely tied to specific associations between inputs to the region. In a set of simulated lesion experiments, we found that our model's learned associations between the different input streams were sufficient for the network to perform pattern completion in the absence of any one input. However, if these associations were disrupted in such a way that there was conflict between the behavioral variables associated with a specific representation, then the network's route representations become sufficiently different that they are uncorrelated or anticorrelated in response to that conflict. This is consistent with findings reported by Nelson, Powell, Holmes, Vann, and Aggleton (2015) in which rats with lesions to the RSC were able to solve a standard T maze alteration task if intramaze cues were consistent, but showed impairments when cue placements were in conflict. The findings suggested that the RSC is sensitive to multiple cue types and that damage to the RSC might impair the rats' ability to flexibly switch their attention from between either different kinds of spatial information or between different spatial navigation strategies, but the two possibilities could not be disambiguated. Based on activity seen in our evolved SNNs, we suggest the impairments are likely due to both, because specific kinds of inputs seem to be closely related to specific navigation strategies.

These findings are also consistent with evidence from the RSC that it is an association cortex where multimodal sensory information is integrated. For example, evidence suggests that the RSC integrates visual and idiothetic information for error correction during navigation, and that lesions disrupt cross-modal object recognition (Cooper & Mizumori, 1999; Hindley et al., 2014). Moreover, neurons in the region exhibit conjunctive reward-relocation representation similar to the hippocampus (Vedder et al., 2016). Furthermore, our experiments revealed that the neural activity underpinning these representations are not fixed. Instead, activity is flexible and adaptable, such that neurons can represent changing contexts and task demands in a flexible and parameter-free way, which is similar to what has been observed in the parietal cortex (Raposo et al., 2014; Vedder et al., 2016).

On a separate note, this model represents a breakthrough for theoretical neuroscience. By training generic but biologically plausible SNNs to replicate electrophysiologically recorded neural activity, we have added a level of veracity to the results provided by the simulated experiments. The present experiments further show that such models can be used to make powerful theoretical predictions for biological brain regions, which can be tested and verified by experimentalists to the mutual benefit of computational and experimental researchers. Furthermore, this method can generate networks with potentially powerful capabilities that emerge based on network inputs and the constraints of the evolutionary strategy, such as route encoding in the present case. Although similar methods have been used for computational modeling, they have not replicated data from entire brain regions to test cognition and are not as flexible or generalizable. Such methods have instead focused on the output of specific kinds of behavior (Dura-Bernal et

al., 2016), the replication of specific cell types (Carlson, Nagewaran, Dutt, & Krichmar, 2014), or on the evolution of specific neural models (Fountas & Shanahan, 2015). More abstract models, such as NEAT, have focused on evolving efficient neural network topologies rather than replicating neuronal behavior (Stanley & Miikkulainen, 2002), and other methods of evolving SNNs using GPU-accelerated frameworks have focused on using them for data analysis (Kasabov et al., 2013). Our framework may represent a promising avenue for the generation of new predictions and theories about cognitive function, because simulated experiments can be used to inform researchers of the best avenue of research, which could potentially help to save tremendous time and resources.

We also emphasize that the framework is generalizable and adaptable, in contrast with other computational models that are usually designed with a specific task, behavior, and/or brain region in mind (Byrne et al., 2007; Oess et al., 2017). Our method was not designed for one specific brain region or task; therefore, any model can be evolved in response to any number of user-defined inputs, meaning that it can be applied to electrophysiological data sets from virtually any brain region. The evolved models resulting from the method can also be adapted for use in different tasks than the ones they were designed to solve. This represents a major stepping stone for computational neuroscientists because it is a new, flexible, and dynamic way of creating and testing theoretical models of cognition and behavior. We suggest that these kinds of tools and techniques will be very powerful as we further our efforts to understand how neural activity gives rise to behavior.

For the RSC to efficiently represent multimodal inputs and multiple spatial frames of reference, it needs to incorporate a sparse, reduced neuronal code. We propose that the effectiveness of the evolutionary algorithmic approach used in this experiment may be related to dimensionality reduction as performed by STDP. Recent studies suggest that a statistical learning technique known as nonnegative matrix factorization (NMF) may be able to describe receptive fields and neuronal response properties in the medial superior temporal dorsal (MSTd) region of the brain, which is involved in the perception of visual motion (Beyeler, Dutt, & Krichmar, 2016), and there exists a mathematical proof supporting an equivalence between pair-based STDP with homeostasis and NMF (Carlson et al., 2013). Using a similar methodology to Beyeler et al. (2016), we applied NMF, with a sparsity constraint, to the parameterized behavioral variables that were used as input to the network in order to generate another set of synthetic neuron responses. We found that the response profiles of these neurons were remarkably similar to those that resulted from the evolutionary algorithm (Figure 14). We suggest that STDP and homeostasis may perform dimensionality reduction in the form of NMF on its inputs, allowing the network to extract a compressed representation of the input. This process may explain the rapid convergence of the evolutionary algorithm, and provides a mechanistic explanation for the evolved networks' behavior (Beyeler, Rounds, Carlson, Dutt, & Krichmar, 2017).

In summary, the present results support the idea that the RSC is an association cortex that can rapidly and flexibly integrate multisensory information. In a navigation task, such as that explored here, the RSC contains multiple frames of reference (i.e., allocentric, egocentric, route-centric) necessary to support spatial memory (Alexander & Nitz, 2015; Oess et al., 2017). Moreover, the present results support the idea that the RSC carries a conjunctive code

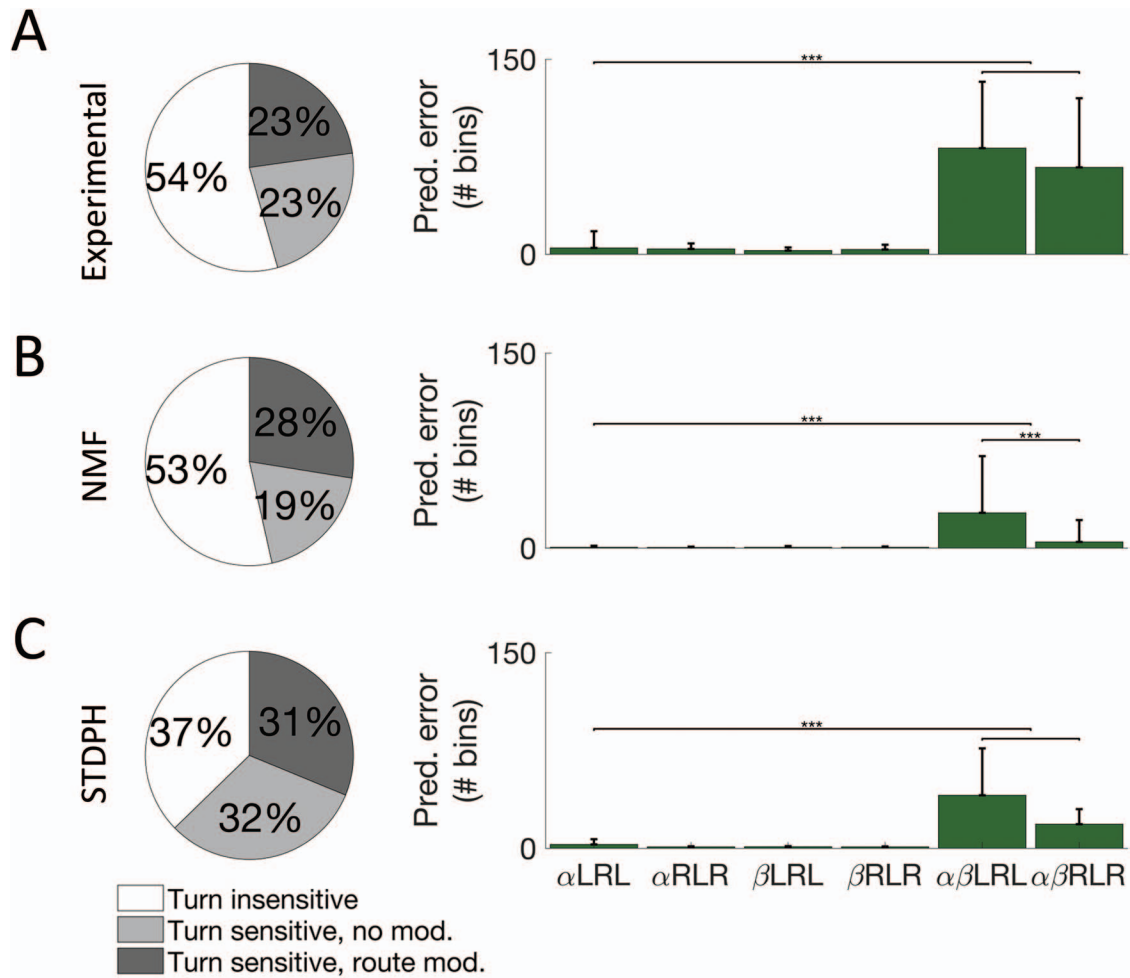


Figure 14. Adapted from Beyeler et al. (2017). (A) Left, neuronal type distributions present in the experimental dataset. Right, average reconstruction errors of position based on ensemble correlation matrices from experimentally recorded neuron responses. (B) Left, neuronal type distributions of synthetic neuron responses generated by the nonnegative matrix factorization (NMF) algorithm. Right, average reconstruction errors of position based on ensemble correlation matrices from NMF-generated synthetic neuron responses. (C) Left, neuronal type distributions of synthetic neuron responses generated by STDPH and the evolutionary algorithm. Right, average reconstruction errors of position based on ensemble correlation matrices from STDPH-generated synthetic neuron responses. *** $p < .001$. See the online article for the color version of this figure.

that makes it resilient to perturbations. Finally, the present results highlight the ability of the RSC to remap due to contextual changes.

References

Alexander, A. S., & Nitz, D. A. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nature Neuroscience*, *18*, 1143–1151. <http://dx.doi.org/10.1038/nn.4058>

Alexander Andrew, S., & Nitz Douglas, A. (2017). Spatially periodic activation patterns of retrosplenial cortex encode route sub-spaces and distance traveled. *Current Biology*, *27*, 1551–1560.

Beyeler, M., Carlson, K. D., Chou, T. S., Dutt, N., & Krichmar, J. L. (2015, July 12–17). CARLsim 3: A user-friendly and highly optimized library for the creation of neurobiologically detailed spiking neural networks. *2015 International Joint Conference on Neural Networks* (pp. 1–8).

Killarney, Ireland: IEEE. <http://dx.doi.org/10.1109/IJCNN.2015.7280424>

Beyeler, M., Dutt, N., & Krichmar, J. L. (2016). 3D Visual Response Properties of MSTd Emerge from an Efficient, Sparse Population Code. *The Journal of Neuroscience*, *36*, 8399–8415. <http://dx.doi.org/10.1523/JNEUROSCI.0396-16.2016>

Beyeler, M., Rounds, E., Carlson, K., Dutt, N., & Krichmar, J. L. (2017). Sparse coding and dimensionality reduction in cortex. *bioRxiv*. Advance online publication. <http://dx.doi.org/10.1101/149880>

Burgess, N., Becker, S., King, J. A., & O’Keefe, J. (2001). Memory for events and their spatial context: Models and experiments. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences*, *356*, 1493–1503. <http://dx.doi.org/10.1098/rstb.2001.0948>

Byrne, P., Becker, S., & Burgess, N. (2007). Remembering the past and imagining the future: A neural model of spatial memory and imagery.

- Psychological Review*, 114, 340–375. <http://dx.doi.org/10.1037/0033-295X.114.2.340>
- Carlson, K. D., Beyeler, M., Dutt, N., & Krichmar, J. L. (2014, February). *GPGPU accelerated simulation and parameter tuning for neuromorphic applications*. Paper presented at the 19th Asia and South Pacific Design Automation Conference (ASP-DAC'14).
- Carlson, K. D., Nageswaran, J. M., Dutt, N., & Krichmar, J. L. (2014). An efficient automated parameter tuning framework for spiking neural networks. *Frontiers in Neuroscience*, 8, 10. <http://dx.doi.org/10.3389/fnins.2014.00010>
- Carlson, K. D., Richert, M., Dutt, N., & Krichmar, J. L. (2013, August 4–9). Biologically plausible models of homeostasis and STDP: Stability and learning in spiking neural networks. *The 2013 International Joint Conference on Neural Networks* (pp. 1–8). Dallas, TX: IEEE. <http://dx.doi.org/10.1109/IJCNN.2013.6706961>
- Cho, J., & Sharp, P. E. (2001). Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behavioral Neuroscience*, 115, 3–25. <http://dx.doi.org/10.1037/0735-7044.115.1.3>
- Chrastil, E. R., Sherrill, K. R., Hasselmo, M. E., & Stern, C. E. (2015). There and back again: Hippocampus and retrosplenial cortex track homing distance during human path integration. *The Journal of Neuroscience*, 35, 15442–15452. <http://dx.doi.org/10.1523/JNEUROSCI.1209-15.2015>
- Chrastil, E. R., Sherrill, K. R., Hasselmo, M. E., & Stern, C. E. (2016). Which way and how far? Tracking of translation and rotation information for human path integration. *Human Brain Mapping*, 37, 3636–3655. <http://dx.doi.org/10.1002/hbm.23265>
- Clark, B. J., Bassett, J. P., Wang, S. S., & Taube, J. S. (2010). Impaired head direction cell representation in the anterodorsal thalamus after lesions of the retrosplenial cortex. *The Journal of Neuroscience*, 30, 5289–5302. <http://dx.doi.org/10.1523/JNEUROSCI.3380-09.2010>
- Cooper, B. G., & Mizumori, S. J. (1999). Retrosplenial cortex inactivation selectively impairs navigation in darkness. *Neuroreport*, 10, 625–630. <http://dx.doi.org/10.1097/00001756-199902250-00033>
- Dura-Bernal, S., Neymotin, S. A., Kerr, C. C., Sivagnanam, S., Majumdar, A., Francis, J. T., & Lytton, W. W. (2016). Evolutionary algorithm optimization of biological learning parameters in a biomimetic neuroprosthesis. *IBM Journal of Research and Development*, 61(2/3), 6:1–14–6. <http://dx.doi.org/10.1147/JRD.2017.2656758>
- Eichenbaum, H. (2017). Barlow versus Hebb: When is it time to abandon the notion of feature detectors and adopt the cell assembly as the unit of cognition? *Neuroscience Letters*. Advance online publication. <http://dx.doi.org/10.1016/j.neulet.2017.04.006>
- Epstein, R. A. (2008). Parahippocampal and retrosplenial contributions to human spatial navigation. *Trends in Cognitive Sciences*, 12, 388–396. <http://dx.doi.org/10.1016/j.tics.2008.07.004>
- Epstein, R. A., Parker, W. E., & Feiler, A. M. (2007). Where am I now? Distinct roles for parahippocampal and retrosplenial cortices in place recognition. *The Journal of Neuroscience*, 27, 6141–6149. <http://dx.doi.org/10.1523/JNEUROSCI.0799-07.2007>
- Fidjeland, A. K., & Shanahan, M. P. (2010, July 18–23). Accelerated simulation of spiking neural networks using GPUs. *The 2010 International Joint Conference on Neural Networks* (pp. 1–8). Barcelona, Spain: IEEE. <http://dx.doi.org/10.1109/IJCNN.2010.5596678>
- Foster, D. J., Morris, R. G. M., & Dayan, P. (2000). A model of hippocampally dependent navigation, using the temporal difference learning rule. *Hippocampus*, 10, 1–16. [http://dx.doi.org/10.1002/\(SICI\)1098-1063\(2000\)10:1<1::AID-HIPO1>3.0.CO;2-1](http://dx.doi.org/10.1002/(SICI)1098-1063(2000)10:1<1::AID-HIPO1>3.0.CO;2-1)
- Fountas, Z., & Shanahan, M. (2015, July 12–17). GPU-based fast parameter optimization for phenomenological spiking neural models. *2015 International Joint Conference on Neural Networks* (pp. 1–8). Killarney, Ireland: IEEE. <http://dx.doi.org/10.1109/IJCNN.2015.7280668>
- Froemke, R. C. (2015). Plasticity of cortical excitatory-inhibitory balance. *Annual Review of Neuroscience*, 38, 195–219. <http://dx.doi.org/10.1146/annurev-neuro-071714-034002>
- Hindley, E. L., Nelson, A. J. D., Aggleton, J. P., & Vann, S. D. (2014). Dysgranular retrosplenial cortex lesions in rats disrupt cross-modal object recognition. *Learning & Memory*, 21, 171–179. <http://dx.doi.org/10.1101/lm.032516.113>
- Jacob, P.-Y., Casali, G., Spieser, L., Page, H., Overington, D., & Jeffery, K. (2017). An independent, landmark-dominated head-direction signal in dysgranular retrosplenial cortex. *Nature Neuroscience*, 20, 173–175. <http://dx.doi.org/10.1038/nn.4465>
- Kasabov, N., Dhoble, K., Nuntalid, N., & Indiveri, G. (2013). Dynamic evolving spiking neural networks for on-line spatio- and spectro-temporal pattern recognition. *Neural Networks*, 41, 188–201. <http://dx.doi.org/10.1016/j.neunet.2012.11.014>
- Kim, J. G., Aminoff, E. M., Kastner, S., & Behrmann, M. (2015). A neural basis for developmental topographic disorientation. *The Journal of Neuroscience*, 35, 12954–12969. <http://dx.doi.org/10.1523/JNEUROSCI.0640-15.2015>
- Maguire, E. A. (2001). The retrosplenial contribution to human navigation: A review of lesion and neuroimaging findings. *Scandinavian Journal of Psychology*, 42, 225–238. <http://dx.doi.org/10.1111/1467-9450.00233>
- Miller, A. M. P., Vedder, L. C., Law, L. M., & Smith, D. M. (2014). Cues, context, and long-term memory: The role of the retrosplenial cortex in spatial cognition. *Frontiers in human neuroscience*, 8, 586.
- Nelson, A. J. D., Hindley, E. L., Pearce, J. M., Vann, S. D., & Aggleton, J. P. (2015). The effect of retrosplenial cortex lesions in rats on incidental and active spatial learning. *Frontiers in Behavioral Neuroscience*, 9, 11. <http://dx.doi.org/10.3389/fnbeh.2015.00011>
- Nelson, A. J. D., Powell, A. L., Holmes, J. D., Vann, S. D., & Aggleton, J. P. (2015). What does spatial alternation tell us about retrosplenial cortex function? *Frontiers in Behavioral Neuroscience*, 9, 126.
- Oess, T., Krichmar, J. L., & Röhrbein, F. (2017). A computational model for spatial navigation based on reference frames in the hippocampus, retrosplenial cortex, and posterior parietal cortex. *Frontiers in Neurobotics*, 11, 4. <http://dx.doi.org/10.3389/fnbot.2017.00004>
- Paine, T., Jin, H., Yang, J., Lin, Z., & Huang, T. (2013). GPU asynchronous stochastic gradient descent to speed up neural network training. *arXiv preprint arXiv*, 1312.6186.
- Pothuizen, H. H. J., Aggleton, J. P., & Vann, S. D. (2008). Do rats with retrosplenial cortex lesions lack direction? *The European Journal of Neuroscience*, 28, 2486–2498. <http://dx.doi.org/10.1111/j.1460-9568.2008.06550.x>
- Preston, A. R., & Eichenbaum, H. (2013). Interplay of hippocampus and prefrontal cortex in memory. *Current Biology*, 23, R764–R773. <http://dx.doi.org/10.1016/j.cub.2013.05.041>
- Raposo, D., Kaufman, M. T., & Churchland, A. K. (2014). A category-free neural population supports evolving demands during decision-making. *Nature Neuroscience*, 17, 1784–1792. <http://dx.doi.org/10.1038/nn.3865>
- Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497, 585–590. <http://dx.doi.org/10.1038/nature12160>
- Robinson, S., Keene, C. S., Iaccarino, H. F., Duan, D., & Bucci, D. J. (2011). Involvement of retrosplenial cortex in forming associations between multiple sensory stimuli. *Behavioral Neuroscience*, 125, 578–587. <http://dx.doi.org/10.1037/a0024262>
- Rounds, E. L., Scott, E. O., Alexander, A. S., De Jong, K. A., Nitz, D. A., & Krichmar, J. L. (2016). An evolutionary framework for replicating neurophysiological data with spiking neural networks. In J. Handl, E. Hart, P. R. Lewis, M. López-Ibáñez, G. Ochoa, & B. Paechter (Eds.), *Parallel problem solving from nature—PPSN XIV: 14th International*

- Conference* (pp. 537–547). Cham, UK: Springer International Publishing.
- Schlesinger, M. I., Cannova, C. C., Boublil, B. L., Hales, J. B., Mankin, E. A., Brandon, M. P., . . . Leutgeb, S. (2015). The medial entorhinal cortex is necessary for temporal organization of hippocampal neuronal activity. *Nature Neuroscience*, *18*, 1123–1132. <http://dx.doi.org/10.1038/nn.4056>
- Stanley, K. O., & Miiikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, *10*, 99–127. <http://dx.doi.org/10.1162/106365602320169811>
- Takahashi, N., Kawamura, M., Shiota, J., Kasahata, N., & Hirayama, K. (1997). Pure topographic disorientation due to right retrosplenial lesion. *Neurology*, *49*, 464–469. <http://dx.doi.org/10.1212/WNL.49.2.464>
- Taube, J. S. (1995). Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *The Journal of Neuroscience*, *15*, 70–86.
- Todd, T. P., Huszár, R., DeAngeli, N. E., & Bucci, D. J. (2016). Higher-order conditioning and the retrosplenial cortex. *Neurobiology of Learning and Memory*, *133*, 257–264. <http://dx.doi.org/10.1016/j.nlm.2016.05.006>
- Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., . . . Morris, R. G. (2007). Schemas and memory consolidation. *Science*, *316*, 76–82. <http://dx.doi.org/10.1126/science.1135935>
- Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., . . . Morris, R. G. (2011). Schema-dependent gene activation and memory encoding in neocortex. *Science*, *333*, 891–895. <http://dx.doi.org/10.1126/science.1205274>
- van Groen, T., & Wyss, J. M. (1990). Connections of the retrosplenial granular a cortex in the rat. *The Journal of Comparative Neurology*, *300*, 593–606. <http://dx.doi.org/10.1002/cne.903000412>
- van Groen, T., & Wyss, J. M. (2003). Connections of the retrosplenial granular b cortex in the rat. *The Journal of Comparative Neurology*, *463*, 249–263. <http://dx.doi.org/10.1002/cne.10757>
- Vann, S. D., Aggleton, J. P., & Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nature Reviews Neuroscience*, *10*, 792–802. <http://dx.doi.org/10.1038/nrn2733>
- Vedder, L. C., Miller, A. M. P., Harrison, M. B., & Smith, D. M. (2016). Retrosplenial cortical neurons encode navigational cues, trajectories and reward locations during goal directed navigation. *Cerebral Cortex*. Advance online publication. <http://dx.doi.org/10.1093/cercor/bhw192>
- White, D. R. (2012). Software review: The ECJ toolkit. *Genetic Programming and Evolvable Machines*, *13*, 65–67. <http://dx.doi.org/10.1007/s10710-011-9148-z>
- Wyss, J. M., & van Groen, T. (1992). Connections between the retrosplenial cortex and the hippocampal formation in the rat: A review. *Hippocampus*, *2*, 1–11. <http://dx.doi.org/10.1002/hipo.450020102>

Received December 2, 2017

Accepted January 11, 2018 ■