

# Episodic memory supports the acquisition of structured task representations

Qihong Lu<sup>1,2</sup>, Ali Hummos<sup>3</sup>, & Kenneth A. Norman<sup>1</sup>

<sup>1</sup> Princeton Neuroscience Institute and Department of Psychology, Princeton University

<sup>2</sup> Center for Theoretical Neuroscience and The Mortimer B. Zuckerman Mind Brain Behavior Institute, Columbia University

<sup>3</sup> Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

Correspondence: qihong.lu@columbia.edu

## Abstract

Generalization to new tasks requires learning of task representations that accurately reflect the similarity structure of the task space. Here, we argue that episodic memory (EM) plays an essential role in this process by stabilizing task representations, thereby supporting the accumulation of structured knowledge. We demonstrate this using a neural network model that infers task representations that minimize the current task’s objective function; crucially, the model can retrieve previously encoded task representations from EM and use these to initialize the task inference process. With EM, the model succeeds in learning the underlying task structure; without EM, task representations drift and the network fails to learn the structure. We further show that EM errors can support structure learning by promoting the activation of similar task representations in tasks with similar sensory inputs. Overall, this model provides a novel account of how EM supports the acquisition of structured task representations.

**Keywords:** Episodic memory (EM); task representation; structure learning; neural network modeling.

## Reusing task representations via EM

Prior neural network modeling work has shown that having an explicit representation of the current task (in a separate model layer, projecting to the rest of the model) can facilitate task performance by reducing interference from other tasks (Flesch, Nagy, Saxe, & Summerfield, 2023; Masse, Grant, & Freedman, 2018; ?, ?; Russin, Zolfaghar, Park, Boorman, & O’Reilly, 2022). Crucially, this representation should also reflect the structure *across* related tasks, so knowledge about other tasks can be leveraged for generalization to new tasks (Yang, Joglekar, Song, Newsome, & Wang, 2019; Flesch, Saxe, & Summerfield, 2023; Ito et al., 2022; Tafazoli et al., 2024). Consider the problem of driving. Vehicles travel on the right in the US and on the left in the UK. When traveling from the US to the UK for the first time, one needs to construct a UK representation to avoid interference, but this representation also needs to be structurally related to the US representation to benefit from prior knowledge about driving. How should an intelligent agent build task representations online that satisfies both criteria?

Previous studies have suggested that episodic memory (EM) contributes to this process by retrieving previously formed task representations. This claim is supported by a large body of experimental work (for a review, see Egner, 2023). It also resonates with theoretical work on *amortized inference*, which posits that inferring an appropriate task representation online can be computationally costly, and that per-

formance can be improved by re-using the products of these costly inference processes (Dasgupta, Schulz, Goodman, & Gershman, 2018; Dasgupta & Gershman, 2021); in line with this view, it has been shown that EM in particular can support efficient re-use when the same task repeats (Dasgupta et al., 2018; Lu et al., 2023).

Here, we go beyond prior work by demonstrating that EM helps the model learn task representations that veridically reflect the structure *across* tasks. This contribution of EM to structure learning arises for two reasons: First, inferring a task representation with information about the ongoing task only is underdetermined – representations that support good performance for the ongoing task might not be unique, causing the representation for a task to drift over time. This lack of stability can hinder the accumulation of relational structure across tasks – it is difficult to build an intricately structured building on shifting sands. EM can counteract this instability by biasing the network to consistently re-use previously discovered task representations. We also demonstrate that the stabilizing influence of EM is sometimes not enough to support structure learning, but in this case EM *errors* can help to rescue structure learning by promoting activation of similar task representations in tasks with similar sensory inputs.

## Model architecture

We developed a neural network augmented by a task representation layer and an EM buffer (Fig. 1a). The model infers a task representation at each task switch point (task switches are explicitly signaled to the model – a simplification that allows us to focus on the influence of EM on task representation). Several recent models of task inference have assumed orthogonal or non-overlapping task representations to i) avoid interference across tasks and ii) improve computational feasibility (Franklin, Norman, Ranganath, Zacks, & Gershman, 2020; ?, ?); as a result, these models can not directly represent similarity structure across tasks. To get around this limitation, we leveraged the idea of using gradient-based search in activation space to infer task representations online (Giallanza, Campbell, Cohen, & Rogers, 2023; Hummos, 2022): After a task switch, the model finds a task representation by adjusting the task layer activation to optimize the ongoing task objective. Concretely, the model collects a small sample of observations from the ongoing task. Then, to minimize the loss of this sample, the model performs gradient descent in



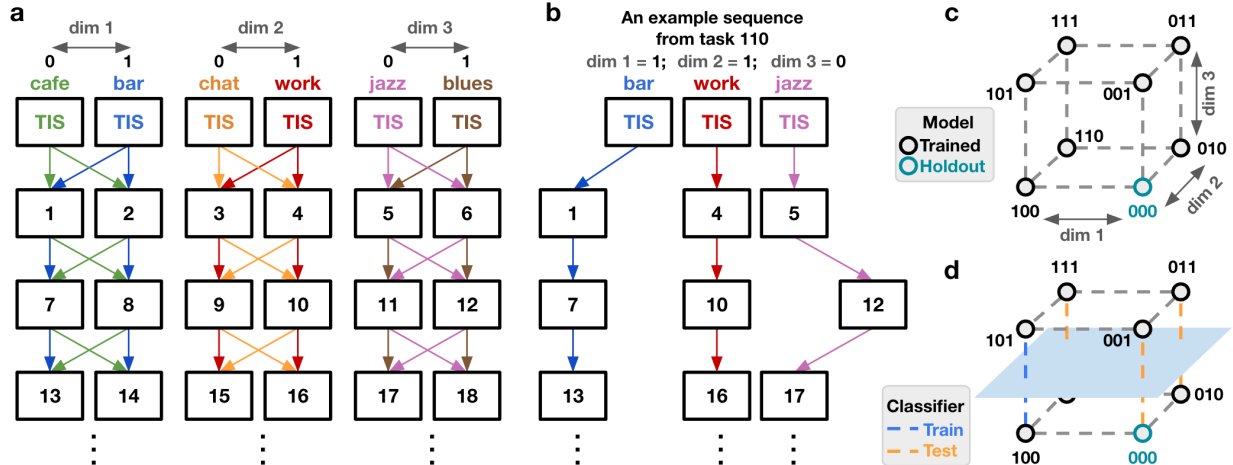


Figure 2: **(a)** The graph used to generate task-dependent sequences in Simulation 1. State transitions are controlled by the task identity (e.g., bar + work + jazz). For example, in a bar, the blue transitions will be used. At the very beginning of the sequence, a three-dimensional task-indicative signal (TIS) is presented. **(b)** In the bar + work + jazz task (i.e., task 110), the state [1, 4, 5] goes to [7, 10, 12], and then [13, 16, 17]. **(c)** The ideal neural representational geometry for the eight tasks should form a cube, as the three binary task dimensions are orthogonal. **(d)** To test if the representation for the third dimension is orthogonal to the other two, we can compute the cross-condition generalization performance (CCGP; Bernardi et al., 2020) along dimension 3: We can start with training a classifier to differentiate the representations for task 101 versus 100. Then, we test the resulting decision plane (marked in light blue) to differentiate the third dimension while varying the other two task dimensions (i.e., 001 versus 000, 111 versus 110, and 011 versus 010). If the test accuracy is 1 for all combinations of train-test splits and across all three dimensions, then CCGP is maximal (i.e., 1).

(i.e., in a cafe + work + jazz task), then [1, 4, 5] leads to [8, 10, 12] – only the first element of the state is impacted.

Given the current state, the model’s objective is to predict the upcoming state. At the beginning of every sequence, a three-dimensional task-indicative signal (TIS) was presented (e.g. the TIS for bar + work + jazz is 110, see Fig. 2b). All models were trained on a blocked curriculum for six epochs. In every epoch, the model cycles in a random order through 7 of the 8 tasks (all tasks except for the holdout task 000; Fig. 2c), receiving a block of 200 sequences per task.

At the end of every sequence, the model first replays the just-experienced sequence to update its task representation. Then, it encodes a memory consisting of the TIS as the memory key and the current task representation. Whenever there is an existing memory indexed by the same memory key, the encoded task representation for that old memory is updated using the new one. Then, the weights are updated to minimize the loss of this sequence. At the beginning of every new sequence, the model first uses the just-observed TIS to retrieve a previously formed task representation to initialize task inference. Then, the model performs task inference using the initial observations of the sequence.

Ideally, the neural representations for the eight tasks should form a cube (Fig. 2c). In other words, the coding direction representing the cafe-bar difference should be “abstract” (as defined by Behrens et al., 2018; Bernardi et al., 2020; Vaidya & Badre, 2022) – it should be invariant to whether the current task is chat or work and jazz or blues. We used cross-condition generalization performance (CCGP; Bernardi et al., 2020) to measure abstraction along the three task dimensions.

For example, consider dimension 3: We can start with training a classifier to differentiate the neural representations for task 000 versus 001. This classifier should generalize regardless of what the first two task dimensions are. Namely, the resulting decision plane (the light blue plane in Fig. 2d) should also reliably classify 100 versus 101, 010 versus 011, and 110 versus 111. CCGP is maximal (i.e., 1) if the classifier generalization performance is 1 for all three task dimensions.

To understand the effect of EM, we trained three kinds of models: Our baseline was a recurrent neural network (RNN) with Gated Recurrent Units (GRU; Cho et al., 2014) but no task layer and no EM; we compared this to a *no-EM* model (a GRU with a task layer but no EM) and a *with-EM* model (a GRU with a task layer and EM). Then we measured performance during training (Fig. 3b) and test (Fig. 3c) as well as CCGP for the learned task representations during test (Fig. 3a). During training, the (baseline) GRU was slow at learning (Fig. 3b) as blocked training led to catastrophic interference across tasks (McClelland et al., 1995). The no-EM model showed a distinctive pattern of performance where its performance was poor at the start of task blocks but it improved sharply within blocks. The improvement within blocks can be explained in terms of the model inferring task representations that help it minimize across-task interference; the drop at block boundaries can be explained in terms of the model failing to re-find old task representations during the inference process, which impedes the accumulation of knowledge. With EM, the performance reached the ceiling almost immediately after a task switch (Fig. 3b), as the model was able to re-load previously used task representations that give it access to knowl-

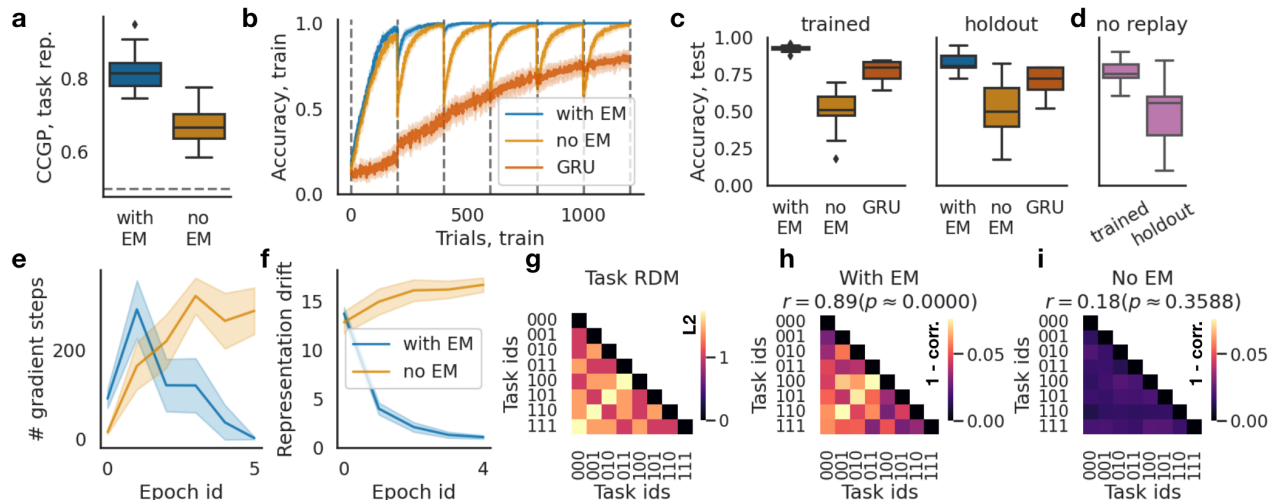


Figure 3: **(a)** CCGP for the task representations is higher for models with EM. **(b)** Prediction accuracy averaged across all seven trained tasks over six training epochs. The dashed lines marked the epoch boundary. GRUs suffered from interference across tasks, slowing learning. No-EM models showed good learning within task blocks but large decreases in performance when the task switched. By contrast, the performance of the with-EM models reached the ceiling more rapidly, as they can retrieve the previously formed task representation. **(c)** During test, accuracy was measured for all eight tasks. The with-EM model was the best at both the trained tasks and the novel holdout task. **(d)** Turning off end-of-sequence replay for the with-EM model reduces performance. This suggests that replaying the just-experienced sequence allowed the with-EM model to encode a more informative task representation. **(e, f)** Models with EM needed fewer gradient steps to find a locally optimal task representation **(e)**, and they had lower representation drift **(f)**, defined as the change of task representation when encountering the same task twice across two consecutive epochs. **(g)** The task RDM computed as the L2 distance between TISs. **(h, i)** The RDMs for task representations during test. The with-EM model RDM **(h)** was highly correlated with the task RDM (Spearman  $r = 0.89, p < 10^{-4}$ ), but the no-EM model RDM was not **(i)**.  $N = 30$  models. Errorbands indicate 2 SEs.

edge acquired in previous blocks.

During the test phase, while the weights were frozen, the model with EM performed the best both for the trained tasks and the holdout task (Fig. 3c). The model with EM also had better CCGP compared to the no-EM model and the GRU (Fig. 3a). These findings of good holdout performance and high CCGP both support the claim that EM promotes structure learning. Additionally, we found that end-of-event replay was important for task performance – when replay was turned off from the with-EM model (Fig. 3d), the performance for both trained tasks and the holdout task was reduced.

We also found that models with EM were more efficient at building the ongoing task representations: They i) needed fewer gradient steps for task inference (Fig. 3e) and ii) had lower representational drift (operationalized as the change of representation when encountering the same task twice) compared to the no-EM model (Fig. 3f).

Finally, to further understand the learned task representations, we compared the representational dissimilarity matrices (RDMs; Kriegeskorte, Mur, & Bandettini, 2008) between the ground truth (Fig. 3g) versus the model representations (Fig. 3h, 3i). Consistent with the CCGP results, the with-EM model RDM was highly correlated with the task RDM (Fig. 3h), and this was not the case for the no-EM model (Fig. 3i). Moreover, this RDM analysis shows that, for the no-EM model, the representations for all tasks became highly similar, which suggests that – by the end of training – the no-EM

model had largely “given up” on using the task layer, instead relying on the RNN to support task-specific responding.

To summarize: Without EM, task representations for the same task can drift perpetually. This suggests that finding a task representation according to the ongoing task objective is underdetermined – there are multiple task representations that satisfy the objective. As a result, the task inference process cannot consistently re-discover previously used representations. This is problematic, as having stable task representations is essential for accumulating knowledge about the temporal structure of the ongoing task as well as the relation across tasks. We found that having EM helped to reduce representation drift; effectively, EM acted as a constraint for the underdetermined task inference problem by requiring the model to re-use previously formed task representations.

## Simulation 2. Learning similarity structure

The model’s goal in Simulation 2 was to learn a set of classification tasks, where the tasks differed only in the angle of their decision boundaries (Fig. 4a). The angles ranged from 0 to 360 degrees at 30-degree intervals. Therefore, the relation across tasks was a ring structure (Fig. 4a). At time  $t$ , the model received a coordinate  $x$ , and it had to predict the associated label  $y$ . As there was no temporal structure in this simulation, we used a feedforward version of our model (the model in Fig. 1a without recurrent connections).

All models were trained for five epochs. In every epoch,




Figure 4: **(a)** The model was given twelve different classification tasks (6 trained, 6 held out), where the ground truth structure across tasks forms a ring. For example, the 8th task is a classification task where the true decision boundary is 240 degrees (counterclockwise from task 0 position). In every epoch, the models experienced tasks with odd indexes in a blocked curriculum (i.e., the trained tasks were 1, 3, 5, ..., 11). After training, all models were tested on all 12 tasks, including the holdout tasks (i.e., 0, 2, 4, ..., 10). **(b, c, d)** The task representations over blocks during training for the model with EM **(b)** formed a ring structure resembling the ground truth. This was not the case for the model without EM **(c)**. This EM-related benefit was largely due to memory errors, as models with perfectly accurate EM **(d)** also did not acquire the ground truth structure. **(e)** The distribution of memory errors during the test phase shows that similar tasks are more confusable. **(f)** The ground truth RDM that encodes the ring structure. **(g)** The average (over models) RDM for models with regular EM closely resembles the ground truth ( $r = 0.75, p < 10^{-4}$ ). **(h, i)** The RDM for no-EM models **(h)** and perfect EM models **(i)** had lower correlations with the ground truth. **(j)** The mean performance across all trained tasks during training. **(k)** The mean performance for all tasks during test. Models with regular EM were much better when generalizing to novel holdout tasks. **(l)** Models with no EM needed more gradient steps to find a locally optimal task representation during training. **(m)** Models with no EM showed much higher representational drift for the same task.  $N = 30$  models. Errorbands indicate 2 SEs.

the model experienced the odd-numbered tasks (1, 3, 5, ..., 11) in a blocked manner. Every block contained 300 trials. The order of the six tasks was randomized. Then, during the test phase, all models were tested on all 12 tasks. The weights were frozen, but the task representation (the purple task layer in 1a) was allowed to be adjusted normally. Note that tasks 0, 2, 4, ..., and 10 were holdout tasks, so they were novel to the model during the test phase.

For the regular EM models, memories were indexed by observations, or  $\{x, y\}$  pairs. At the beginning of every new task, the model collected a small sample of 10 observations,  $\{x_i, y_i\}_{i=0}^{10}$ . For every observation  $\{x_i, y_i\}$ , the model retrieved the task representation with the closest observation, then returned the average of these task representations. This retrieved pattern was used to initialize the task inference process, in which the task representation was optimized to min-

imize the loss over  $\{x_i, y_i\}_{i=0}^{10}$  until convergence. These 10 observations were also used as memory keys to be encoded along with the current task representation. In addition to the regular EM model, we also included a *perfect EM* model that used the true task ID as the memory key and consequently never made memory errors (unlike the regular EM model, which sometimes retrieved memories from the wrong task).

We visualized the task representations discovered for every block by projecting them onto the subspace spanned by the first two principal components (PC; Fig. 4b, 4c, 4d). For the model with EM, the task representations formed a ring structure that resembles the ground truth (Fig. 4b). This was not the case for the model without EM (Fig. 4c). Interestingly, this was also not the case for the perfect EM model (Fig. 4d). The finding that regular EM but not perfect EM learned the ring structure shows that naturally-occurring memory errors

can help to reveal the similarity structure across tasks. For example, if the ongoing task is task 8, the model might retrieve EM of task representation formed for task 7 or 9 (Fig. 4e), as the observations of these tasks are similar. As a result, representations for these three tasks will be used in overlapping sets of scenarios, making them more similar.

To quantitatively compare the structure across learned task representations versus the ground truth during the test phase, we computed the RDMs (Kriegeskorte et al., 2008) for the ground truth relation across tasks (Fig. 4f), models with regular EM (Fig. 4g), models with no EM (Fig. 4h), and models with perfect EM (Fig. 4i). The model RDMs were averaged across models. We then measured the Spearman correlation between each model RDM versus the ground truth task RDM – the results show that the regular EM model was much better at learning the ground truth structure than the no-EM model and the perfect EM model.

Consistent with the previous simulation, learning the ring structure was useful for generalizing to novel tasks consistent with the learned structure. We tracked the performance for all tasks during training (Fig. 4j) and test (Fig. 4k). Importantly, during the test phase, models with regular EM were much better at generalizing to unseen holdout tasks compared to the two alternative models (Fig. 4k).

Also consistent with the previous simulation, models with EM needed fewer gradient steps for task inference (Fig. 4l). Furthermore, representational drift – the change of task representation when revisiting the same task (Fig. 4m) – was much smaller for models with EM. Without EM, representational drift was high (Fig. 4m) even after task performance converged during training (Fig. 4j). This comparison again suggests that task inference is underdetermined, in the sense that there exist many locally optimal task representations for the ongoing task objective; as such, previously discovered task representations cannot be consistently re-discovered. The role of EM is to constrain task inference by biasing the model to re-use previously found task representations. Importantly, though, the results from the perfect EM condition show that stability alone was not sufficient to drive structure learning, as perfect EM minimized drift but did not permit learning of the ring structure – in this case, the model needed the extra factor of memory errors (present in the regular EM condition) to help stitch the task representations together.

## Discussion

We present a neural network model of how EM can support online inference of task representation and shape the relational structure of task representations. Model comparison results suggest EM played a key role in acquiring the compositional structure in Simulation 1 and the similarity structure in Simulation 2; this, in turn, supported generalization to novel tasks that are consistent with the learned task structure.

Our simulation results indicate two ways in which EM can facilitate the acquisition of structured task representations. First, EM reduces representation drift and stabilizes task representations, which is essential for accumulating knowledge

about the relations between tasks. Across the two simulations, we found that, without EM, representations for the same task can drift substantially even after performance converges. This indicates the ongoing task objective is often not constrained enough in the sense that there are many locally optimal task representations. Importantly, not all of these representations reflect how the ongoing task relates to other tasks. Our results show that EM can act as a constraint for task inference by biasing the model to re-use previously discovered task representations.

Simulation 2 also indicates a second, distinct way in which EM can benefit structure learning: When the similarity structure of memory keys is consistent with the similarity structure across tasks, memory errors can help, rather than hinder, the acquisition of the structure. In this situation, memory errors (retrieving a similar task in place of the correct one) will bias the model to use similar representations for similar tasks.

It is notable that perfect (i.e., error-free) EM was sufficient for structure learning in Simulation 1 but not in Simulation 2. We think this is a consequence of the Simulation 1 task being more difficult. That is, in Simulation 1, the model was strongly incentivized to re-use parts of previously learned task representations when learning new tasks because the state transition structure is hard to learn (so “jump-starting” the process with a useful task representation boosts performance a lot). In Simulation 2, the task itself is simple so re-use did not provide as strong of a benefit; in this case, having the extra factor of similarity-based memory errors was helpful for learning similarity structure.

With regard to replay: Our results are consistent with findings showing that humans rapidly replay the just-experienced sequence right before event boundaries (Sols, DuBrow, Davachi, & Fuentemilla, 2017; Silva, Baldassano, & Fuentemilla, 2019) – in Simulation 1, end-of-event replays allowed our model to infer (and then store in EM) a more informative task representation, compared to the task representation inferred based only on data from the start of the task. This predicts that disruption of end-of-event replay would reduce the informativeness of the encoded task representation for later parts of the event but not for the beginning.

In the future, it will be useful to compare our model to other accounts of the role of EM in task representation. Giallanza, Campbell, and Cohen (2023) proposed the episodic generalization and optimization (EGO) framework, which also uses EM to store and retrieve task representations. A key difference is that our model uses EM to initialize task inference, whereas EGO uses backpropagation (through a differentiable EM) to adjust how past memories influence the ongoing task representation. Future work can explore the similarities and differences in the predictions made by EGO versus our model and test them empirically. If the use of EM (shared by the two models) is the key to acquiring structured task representations, then one might expect structured task representations to emerge in EGO in the scenarios described here.

## Acknowledgement

We would like to thank Tyler Giallanza, Declan Campbell, Cody Dong, Jonathan Cohen, and anonymous reviewers for their insightful and constructive feedback.

This work is supported by a Multi-University Research Initiative grant to K.A.N. (ONR/DoD N00014-17-1-2961), an Alan Kanzer Postdoctoral Fellowship to Q.L., and Gatsby Foundation (GAT3708).

## References

- Baldassano, C., Chen, J., Zadbood, A., Pillow, J. W., Hasson, U., & Norman, K. A. (2017, August). Discovering event structure in continuous narrative perception and memory. *Neuron*, 95(3), 709–721.e5.
- Barnett, A. J., Nguyen, M., Spargo, J., Yadav, R., Cohn-Sheehy, B. I., & Ranganath, C. (2023, November). Hippocampal-cortical interactions during event boundaries support retention of complex narrative events. *Neuron*.
- Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018, October). What is a cognitive map? organizing knowledge for flexible behavior. *Neuron*, 100(2), 490–509.
- Ben-Yakov, A., & Henson, R. N. (2018, November). The hippocampal film editor: Sensitivity and specificity to event boundaries in continuous experience. *The Journal of Neuroscience*, 38(47), 10057–10068.
- Bernardi, S., Benna, M. K., Rigotti, M., Munuera, J., Fusi, S., & Salzman, C. D. (2020, November). The geometry of abstraction in the hippocampus and prefrontal cortex. *Cell*, 183(4), 954–967.e21.
- Beukers, A. O., Collin, S. H. P., Kempner, R., Franklin, N. T., Gershman, S. J., & Norman, K. A. (2023, April). Blocked training facilitates learning of multiple schemas. *PsyArXiv*.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014, October). Learning phrase representations using RNN Encoder-Decoder for statistical machine translation. In A. Moschitti, B. Pang, & W. Daelemans (Eds.), *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1724–1734). Doha, Qatar: Association for Computational Linguistics.
- Dasgupta, I., & Gershman, S. J. (2021, January). Memory as a computational resource. *Trends in cognitive sciences*.
- Dasgupta, I., Schulz, E., Goodman, N. D., & Gershman, S. J. (2018, September). Remembrance of inferences past: Amortization in human hypothesis generation. *Cognition*, 178, 67–81.
- Egner, T. (2023, September). Principles of cognitive control over task focus and task switching. *Nature Reviews Psychology*, 1–13.
- Flesch, T., Nagy, D. G., Saxe, A., & Summerfield, C. (2023, January). Modelling continual learning in humans with hebbian context gating and exponentially decaying task signals. *PLoS computational biology*, 19(1), e1010808.
- Flesch, T., Saxe, A., & Summerfield, C. (2023, March). Continual task learning in natural and artificial agents. *Trends in neurosciences*, 46(3), 199–210.
- Franklin, N. T., Norman, K. A., Ranganath, C., Zacks, J. M., & Gershman, S. J. (2020, April). Structured event memory: A neuro-symbolic model of event cognition. *Psychological review*, 127(3), 327–361.
- Giallanza, T., Campbell, D., & Cohen, J. D. (2023, November). Toward the emergence of intelligent control: Episodic generalization and optimization. *PsyArXiv*.
- Giallanza, T., Campbell, D., Cohen, J. D., & Rogers, T. T. (2023, June). An integrated model of semantics and control. *PsyArXiv*.
- Hummos, A. (2022). Thalamus: a brain-inspired algorithm for biologically-plausible continual learning and disentangled representations. *International Conference on Learning Representations*.
- Ito, T., Klingner, T., Schultz, D., Murray, J., Cole, M., & Rigotti, M. (2022). Compositional generalization through abstract representations in human and artificial neural networks. *Advances in neural information processing systems*, 35, 32225–32239.
- Kriegeskorte, N., Mur, M., & Bandettini, P. (2008, November). Representational similarity analysis - connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2, 4.
- Lu, Q., Hasson, U., & Norman, K. A. (2022, February). A neural network model of when to retrieve and encode episodic memories. *eLife*, 11, e74445.
- Lu, Q., Nguyen, T. T., Zhang, Q., Hasson, U., Griffiths, T. L., Zacks, J. M., ... Norman, K. A. (2023, December). Reconciling shared versus Context-Specific information in a neural network model of latent causes. *arXiv*, 2312.08519.
- Masse, N. Y., Grant, G. D., & Freedman, D. J. (2018, October). Alleviating catastrophic forgetting using context-dependent gating and synaptic stabilization. *Proceedings of the National Academy of Sciences of the United States of America*, 115(44), E10467–E10475.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995, July). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3), 419–457.
- Pritzel, A., Uria, B., Srinivasan, S., Badia, A. P., Vinyals, O., Hassabis, D., ... Blundell, C. (2017). Neural episodic control. In D. Precup & Y. W. Teh (Eds.), *Proceedings of the 34th international conference on machine learning* (Vol. 70, pp. 2827–2836). International Convention Centre, Sydney, Australia: PMLR.
- Ritter, S., Wang, J. X., Kurth-Nelson, Z., Jayakumar, S. M., Blundell, C., Pascanu, R., & Botvinick, M. (2018, May). Been there, done that: Meta-Learning with episodic recall. In *Proceedings of the international conference on machine learning (ICML)*.
- Russin, J., Zolfaghar, M., Park, S. A., Boorman, E., &

- O'Reilly, R. C. (2022, July). A neural network model of continual learning with cognitive control. *CogSci ... Annual Conference of the Cognitive Science Society. Cognitive Science Society (U.S.). Conference*, 44, 1064–1071.
- Silva, M., Baldassano, C., & Fuentemilla, L. (2019, October). Rapid memory reactivation at movie event boundaries promotes episodic encoding. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 39(43), 8538–8548.
- Sols, I., DuBrow, S., Davachi, L., & Fuentemilla, L. (2017, November). Event boundaries trigger rapid memory reinstatement of the prior events to promote their representation in Long-Term memory. *Current biology: CB*, 27(22), 3499–3504.e4.
- Tafazoli, S., Bouchacourt, F., Ardalan, A., Markov, N. T., Uchimura, M., Mattar, M. G. G., ... Buschman, T. J. (2024, February). Building compositional tasks with shared neural subspaces. *bioRxiv*, 2024.01.31.578263.
- Vaidya, A. R., & Badre, D. (2022, June). Abstract task representations for inference and control. *Trends in cognitive sciences*, 26(6), 484–498.
- Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T., & Wang, X.-J. (2019, February). Task representations in neural networks trained to perform many cognitive tasks. *Nature neuroscience*, 22(2), 297–306.