PRODUCT KANERVA MACHINES: FACTORIZED BAYESIAN MEMORY

Adam H. Marblestone^{*}, Yan Wu^{*} & Greg Wayne Google DeepMind London, N1C 4AG, UK {amarbles, yanwu, gregwayne}@google.com

Abstract

An ideal cognitively-inspired memory system would compress and organize incoming items. The Kanerva Machine (Wu et al., 2018b;a) is a Bayesian model that naturally implements online memory compression. However, the organization of the Kanerva Machine is limited by its use of a single Gaussian random matrix for storage. Here we introduce the Product Kanerva Machine, which dynamically combines many smaller Kanerva Machines. Its hierarchical structure provides a principled way to abstract invariant features and gives scaling and capacity advantages over single Kanerva Machines. We show that it can exhibit unsupervised clustering, find sparse and combinatorial allocation patterns, and discover spatial tunings that approximately factorize simple images by object.

1 INTRODUCTION

Neural networks may use external memories to flexibly store and access information, bind variables, and learn online without gradient-based parameter updates (Fortunato et al., 2019; Graves et al., 2016; Wayne et al., 2018; Sukhbaatar et al., 2015; Banino et al., 2020; Bartunov et al., 2019; Munkhdalai et al., 2019). Design principles for such memories are not yet fully understood.

The most common external memory is slot-based. It takes the form of a matrix M with columns considered as individual slots. To read from a slot memory, we compute a vector of attention weights w across the slots, and the output is a linear combination $z_{\text{read}} = M$ w. Slot memory lacks key features of human memory. First, it does not automatically compress – the same content can be written to multiple slots. Second, the memory does not naturally organize items according to relational structure, i.e., semantically related items may have unrelated addresses. Third, slot memory is not naturally generative, while human memory supports imagination (Schacter & Madore, 2016). In addition, human memory performs novelty-based updating, semantic grouping and event segmentation (Gershman et al., 2014; Franklin et al., 2019; Howard et al., 2007; Koster et al., 2018). It also seems to extract regularities across memories to form "semantic memory" (Tulving et al., 1972), a process likely related to systems consolidation (Kumaran et al., 2016).

The Kanerva Machine (Wu et al., 2018b;a) uses Bayesian inference for memory update, and is naturally compressive and generative. Instead of a matrix M, it maintains a distribution p(M). Reading with a vector of weights w over columns of M – which specify a query or address for lookup– corresponds to computing a conditional probability of an observation $p(\mathbf{z}|M, \mathbf{w})$, while writing corresponds to computing the posterior given an observation $p(M|\mathbf{z}, \mathbf{w})$. The Kanerva Machine has a few disadvantages, due to its flat memory structure. First, it distributes information across all parameters of the memory distribution without natural grouping. Second, computationally, it scales poorly with the number of columns m of M, $O(m^3)$ for inferring optimal weights w.

We introduce the hierarchically structured *Product Kanerva Machine*, which organizes memories into products of k independent parts in the form $P(M) = \prod_{i=1}^{k} p(M_i)$. This statistical assumption, similar to ICA (Bell & Sejnowski, 1995), is closely related to models of disentanglement (Higgins et al., 2018; Locatello et al., 2018), which assume independent priors within each sample, i.e., $p(z) = \prod_{i=1}^{k} p(z_i)$ — here, product factorization of the memory flexibly encompasses independence within versus across episodes. Multi-component memory is inspired by neuroscience models of the gating of memory by contextual signals (Podlaski et al., 2020; Basu et al., 2016; Pignatelli et al., 2019). Factorizing a Kanerva Machine brings a computational speed advantage, and allows the individual machines within to specialize, leading to meaningful grouping of information.

^{*}Equal contribution



Figure 1: Product Kanerva Machine concept and scaling. A) Architecture. B) Generative model. C) Theoretical scaling of run time with machines k and total columns m.D) Fit to empirical scaling for m = 100 : 600. E) Training curves for MNIST reconstruction at fixed m = 30 and varying k.

2 THE PRODUCT KANERVA MACHINE

The memory system (Fig. 1A) is composed from k small Kanerva Machines each with m_i columns and c rows, where c is the latent code size and $m_i = m/k$ is the number of memory columns per single Kanerva Machine. In our experiments, an encoder/decoder pair is used to map between images x and latent codes z. An assignment network, here a simple multilayer perceptron (see Supp. A.1 for details), is used to compute soft weights $\{r_i\}$ that define the relative strength of reading from or writing to the *i*th Kanerva Machine. The model supports writing, queried reconstruction (i.e., "reading"), and generation. When generating (Fig. 1B), the assignment network is conditioned on a history variable h, while when writing to the memory it is conditioned on the current z, and when reading it is conditioned on the read query z_{query} . Column weights w_i are optimized by least-squares for reconstructing the query z_{query} (see Supp. A.1 for details).

The *i*th Kanerva Machine has the matrix normal distribution $p(\text{vec}(M_i)) \sim \mathcal{N}(\text{vec}(R_i), V_i \otimes I)$ where R_i is a $c \times m_i$ matrix containing the mean of M_i , with m_i its number of columns, V_i is a $m_i \times m_i$ matrix giving the covariance between columns, vectorization means concatenation of columns, and the identity matrix I is $c \times c$. Given addressing weights \mathbf{w}_i for machine *i*, the read-out from memory is the conditional distribution $p(\mathbf{z}|M_i) \sim \mathcal{N}(M_i \mathbf{w}_i, \sigma_i^2 \cdot I)$.

Two possible factorizations are mixtures and products. A product model assumes a factorized likelihood $p(\mathbf{z}) \propto \prod_{i=1}^{k} p(\mathbf{z}|M_i)$, which encourages each component $p(\mathbf{z}|M_i)$ to extract combinatorial, i.e., statistically independent, features across episodes (Williams et al., 2002; Hinton, 1999; 2002; Welling, 2007). A product factorization could therefore comprise a prior encouraging disentanglement across episodes, a milder condition than enforcing factorization across an entire dataset (Locatello et al., 2018; Burgess et al., 2019; Higgins et al., 2018; Watters et al., 2019). A mixture model (or the related switching models (Fox et al., 2009)), on the other hand, tends to find a nearest-neighbour mode dominated by one canonical prototype (Hasselblad, 1966; Shazeer et al., 2017). To address both scenarios, we use a "generalized product" (Cao & Fleet, 2014; Peng et al., 2019) (Eq. 1), containing both products and mixtures as limits (see Supp. A.5). We thus consider a joint distribution between z and all k memory matrices, with each term raised to a power $r_i \geq 0$

$$p(\mathbf{z}, M_1, \dots M_i, \dots) \propto \prod_{i=1}^k p(\mathbf{z}, M_i)^{r_i}$$
 (1)

During writing, $\{r_i\}$ are inferred from the observation \mathbf{z} , and a variable \mathbf{h} which stores information about history. We use $p(\{r_i\}|\mathbf{h})$ during generation and an approximate posterior $q(\{r_i\}|\mathbf{z},\mathbf{h})$ during inference (see Supp. A.2.1 for details). Once $\{r_i\}$ are given, Eq. 1 becomes a product

of Gaussians, which stays in the linear Gaussian family, allowing tractable inference (Roweis & Ghahramani, 1999). Given z and $\{r_i\}$, writing occurs via a Bayesian update to each memory distribution $p(M_i|\mathbf{z}, \{r_i\})$. Updates for the k memories, given z are (see Supp. A.3 for derivation)

$$\Delta = \mathbf{z} - \mu_{\mathbf{z}} \tag{2}$$

$$R_i \leftarrow R_i + \beta_i \Delta \mathbf{w}_i^\top V_i \tag{3}$$

$$V_i \leftarrow V_i - \beta_i \, V_i \, \mathbf{w}_i \mathbf{w}_i^\top \, V_i \tag{4}$$

where

$$\beta_i = \frac{1}{\mathbf{w}_i^\top V_i \, \mathbf{w}_i + \sigma_i^2 / r_i} \text{ and } \mu_{\mathbf{z}} = \frac{\sum_{i=1}^k \frac{r_i}{\sigma_i^2} R_i \, \mathbf{w}_i}{\sum_{j=1}^k \frac{r_j}{\sigma_j^2}} = \sum_{i=1}^k \gamma_i R_i \mathbf{w}_i$$

For reading, μ_z is used as the memory readout. Note how the prediction error term Δ (as in a Kalman Filter) now *couples* the k machines, via μ_z . Algorithms for writing/reading are given in Supp. A.4. The generative model (Fig. 1B) is trained by maximizing a variational lower bound (Kingma & Welling, 2013) \mathcal{L} on $\ln p_{\theta}(\mathbf{x})$ derived in Supp. A.2 (see Supp. A.9 for conditional generations).

3 **Results**

3.1 SCALING

We first asked if a product factorization could give a computational advantage. For a single Kanerva Machine, solving for w scales as $\mathcal{O}(m^3)$ due to the use of a Cholesky decomposition in the least-squares optimization. Parallel operation across k machines gives theoretical scaling of $\mathcal{O}((\frac{m}{k})^3)$. If there are substantial fixed and per-machine overheads, we predict a scaling of the run time of $c + ak + b(m/k)^3$, with optimum at $k_{\text{opt}} = (\frac{3bm^3}{a})^{\frac{1}{4}}$. The empirically determined scaling of the run-time matches¹ this model (Fig. 1C-D). For m > 500, a large speed advantage results even for moving from k = 1 to 2, showing computational benefit for product factorization.

3.2 QUERIED RECONSTRUCTION

We began with a simple queried reconstruction task. A memory of m = 30 total columns was divided into k = 1 to k = 15 machines. A set of 45 MNIST digits were written, and then the memory was read using each item as a query, and the average reconstruction loss $\langle \ln p_{\theta}(\mathbf{x}|\mathbf{z}) \rangle_{q_{\phi}(\mathbf{z})}$ was computed. All factorizations eventually achieved similar reconstruction accuracy (Fig. 1E), showing that product factorization does not incur a loss in representational power.

3.3 PATTERN COMPLETION



Figure 2: RGB binding task demonstrating sparse machine usage and unsupervised classification. A) For each step in an episode, machine weights γ_0 , γ_1 are displayed as a thorn plot, showing sparsity of usage. B) From left to right: original image, query image, reconstructions from each machine (with other machine blanked) given query image, reconstruction from full product model given query image. C) Performance when factorizing m = 60 total columns into k = 2 machines of size 30 vs. a single Kanerva Machine (and stop gradient to assignment weights r_i demonstrating that **r** is optimized). D) Unsupervised partial MNIST digit classification by the model with k = 2: the machine assignment weight γ_1 is highly correlated with the Red digit class.

We next tested the Product Kanerva Machine on the storage of associations/bindings between highdimensional variables, to ask if a product (k = 2) model might show an advantage over a single

¹Parameters $a = 3.318e-08 \pm 1.035e-09$, $b = 2.176e-01 \pm 3.913e-02$, $c = 3.676e-02 \pm 3.787e-03$ fit to m = 400 with $R^2 = 0.996$, which then explain m = 600 with $R^2 = 0.992$ and m = 100 with $R^2 = 0.958$.

Kanerva Machine (k = 1). A set of 45 triplets of MNIST digits were stored in memories with m = 60 total columns, each triplet consisting of an MNIST digit for the Red, Green and Blue channels of an image. Partial queries consisting of the R and B channels, but not the G channel, were presented and the average reconstruction loss was computed across all 3 channels.

For k = 2 machines, the system finds a sparse machine usage pattern (Fig. 2A), with $r_i \approx 0$ or 1, but both machines used equally overall. When a given individual machine is used, it reconstructs the full bound pattern, while the unused machine produces a fixed degenerate pattern (Fig. 2B). A product of k = 2 machines each of 30 columns outperforms a single machine of 60 columns, while a stop gradient on r abolishes this (Fig. 2C). Choice of r depends on digit class (Fig. 2D), e.g., in Fig. 2D predominantly on the R digit but not the B or G digits (symmetry is broken between B and R from run to run leading to horizontal or vertical stripes) – see Supp. A.6 for full RGB class selectivity matrix. Thus, the model optimizes allocation via sparse, dynamic machine choice which becomes selective to digit class in an unsupervised fashion.





Figure 3: "Dancing dSprites" task, demonstrating item factorization across multiple machines. A) Task design. B) Factorization examples. C) Reconstruction loss vs. k and m_i . D-E) Spatial tunings of individual machines: cosine distance between the reconstruction from each machine and template dSprite vs. template position. Diagonal in E spans from lower right to upper left in D.

To test whether our model can factorize memory in a semantically meaningful manner, this section introduces a task based on dSprites (Matthey et al., 2017) that reflects the persistence of objects over episodes (Bambach et al., 2018). *Within* each episode, two randomly sampled objects are consistently presented, but moved independently to different positions across different frames. Shapes, orientations and scales varied *across* episodes. This setup mirrors the strong correlations between features that define an object, persisting over a period of time, coupled with the relative independence of the positions of any two objects as they move, which corresponds to the product structure $p(\text{frame, object}_1, \text{object}_2) \propto p(\text{frame, object}_1) \cdot p(\text{frame, object}_2)$ that our model assumes.

For k = 4 machines, each with 5 columns, we observed a form of object-based factorization (Fig. 3B): *individual* machines typically reconstructed distorted images at the positions of single objects (more examples in Supp. A.7). The k = 4 model outperformed a k = 1 model with the same total number of columns (Fig. 3C). Individual machine reconstructions exhibited localized spatial tunings to the positions of the individual dSprites (Fig. 3D,E and Supp. A.8). In contrast, tuning was invariant to the shape, orientation and size of the dSprites (Supp. A.8). Weights r_i were nearly fixed, suggesting that selectivity was not due to varying r. The model thus spontaneously factored according to localized spatial tunings, such that single machines typically reconstructed single objects.

4 FUTURE DIRECTIONS

Product Kanerva Machines could be extended in several ways. Attention-based selection of input elements could be added, or explicit event segmentation over time, or alternative gating methods and forms of communication between machines, as in (Goyal et al., 2019; Santoro et al., 2018; Hinton et al., 2018; Veness et al., 2017; Kipf et al., 2018). Auxiliary losses could encourage richer unsupervised classification (Makhzani et al., 2015) for class-dependent routing. The generative model can be extended with richer distribution families (Rezende & Mohamed, 2015). Joint inference of **r** and **w**_i using Expectation Maximization (EM) algorithms may be possible (Dempster et al., 1977). Further understanding of when and how factorized memories can encourage extraction of objects or other disentangled features may also be of interest. Ultimately, we hope to use compressive, semantically self-organizing and consolidating memories to solve problems of long-term credit assignment (Ke et al., 2018; Hung et al., 2019), continual learning (van de Ven & Tolias, 2018; Rolnick et al., 2019) and transfer (Higgins et al., 2017).

ACKNOWLEDGMENTS

We thank Andrea Banino, Charles Blundell, Matt Botvinick, Marta Garnelo, Jason Ramapuram, Murray Shanahan and Chen Yan for discussions, Sergey Bartunov for initial review of the manuscript, Loic Matthey for help with dSprites, and Seb Noury for assistance with speed profiling.

REFERENCES

- Sven Bambach, David Crandall, Linda Smith, and Chen Yu. Toddler-inspired visual object learning. In Advances in neural information processing systems, pp. 1201–1210, 2018.
- Andrea Banino, Adri Puigdomnech Badia, Raphael Kster, Martin J. Chadwick, Vinicius Zambaldi, Demis Hassabis, Caswell Barry, Matthew Botvinick, Dharshan Kumaran, and Charles Blundell. Memo: A deep network for flexible combination of episodic memories. In *International Conference on Learning Representations*, 2020.
- Sergey Bartunov, Jack W Rae, Simon Osindero, and Timothy P Lillicrap. Meta-learning deep energy-based memory models. *arXiv preprint arXiv:1910.02720*, 2019.
- Jayeeta Basu, Jeffrey D Zaremba, Stephanie K Cheung, Frederick L Hitti, Boris V Zemelman, Attila Losonczy, and Steven A Siegelbaum. Gating of hippocampal activity, plasticity, and memory by entorhinal cortex long-range inhibition. *Science*, 351(6269):aaa5694, 2016.
- Anthony J Bell and Terrence J Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159, 1995.
- Christopher M Bishop. Pattern recognition and machine learning. springer, 2006.
- Christopher P Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. Monet: Unsupervised scene decomposition and representation. arXiv preprint arXiv:1901.11390, 2019.
- Yanshuai Cao and David J Fleet. Generalized product of experts for automatic and principled fusion of gaussian process predictions. arXiv preprint arXiv:1410.7827, 2014.
- Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1): 1–22, 1977.
- Meire Fortunato, Melissa Tan, Ryan Faulkner, Steven Hansen, Adrià Puigdomènech Badia, Gavin Buttimore, Charles Deck, Joel Z Leibo, and Charles Blundell. Generalization of reinforcement learners with working and episodic memory. In *Advances in Neural Information Processing Systems*, pp. 12448–12457, 2019.
- Emily Fox, Erik B. Sudderth, Michael I. Jordan, and Alan S. Willsky. Nonparametric bayesian learning of switching linear dynamical systems. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou (eds.), Advances in Neural Information Processing Systems 21, pp. 457–464. 2009.
- Nicholas Franklin, Kenneth A Norman, Charan Ranganath, Jeffrey M Zacks, and Samuel J Gershman. Structured event memory: a neuro-symbolic model of event cognition. *BioRxiv*, pp. 541607, 2019.
- Samuel J Gershman, Angela Radulescu, Kenneth A Norman, and Yael Niv. Statistical computations underlying the dynamics of memory updating. *PLoS computational biology*, 10(11):e1003939, 2014.
- Anirudh Goyal, Alex Lamb, Jordan Hoffmann, Shagun Sodhani, Sergey Levine, Yoshua Bengio, and Bernhard Schölkopf. Recurrent independent mechanisms. arXiv preprint arXiv:1909.10893, 2019.
- Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538 (7626):471, 2016.
- Victor Hasselblad. Estimation of parameters for a mixture of normal distributions. *Technometrics*, 8(3):431–444, 1966.

- Irina Higgins, Arka Pal, Andrei Rusu, Loic Matthey, Christopher Burgess, Alexander Pritzel, Matthew Botvinick, Charles Blundell, and Alexander Lerchner. Darla: Improving zero-shot transfer in reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1480–1490. JMLR. org, 2017.
- Irina Higgins, David Amos, David Pfau, Sebastien Racaniere, Loic Matthey, Danilo Rezende, and Alexander Lerchner. Towards a definition of disentangled representations. *arXiv preprint arXiv:1812.02230*, 2018.
- Geoffrey E Hinton. Products of experts. 1999.
- Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002.
- Geoffrey E Hinton, Sara Sabour, and Nicholas Frosst. Matrix capsules with em routing. 2018.
- Marc W Howard, Bing Jing, Kelly M Addis, and Michael J Kahana. Semantic structure and episodic memory. *Handbook of latent semantic analysis*, pp. 121–141, 2007.
- Chia-Chun Hung, Timothy Lillicrap, Josh Abramson, Yan Wu, Mehdi Mirza, Federico Carnevale, Arun Ahuja, and Greg Wayne. Optimizing agent behavior over long time scales by transporting value. *Nature communications*, 10(1):1–12, 2019.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv* preprint arXiv:1611.01144, 2016.
- Weonyoung Joo, Wonsung Lee, Sungrae Park, and Il-Chul Moon. Dirichlet variational autoencoder. arXiv preprint arXiv:1901.02739, 2019.
- Nan Rosemary Ke, Anirudh Goyal, Olexa Bilaniuk, Jonathan Binas, Michael C Mozer, Chris Pal, and Yoshua Bengio. Sparse attentive backtracking: Temporal credit assignment through reminding. In Advances in Neural Information Processing Systems, pp. 7640–7651, 2018.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint* arXiv:1412.6980, 2014.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Thomas Kipf, Yujia Li, Hanjun Dai, Vinicius Zambaldi, Alvaro Sanchez-Gonzalez, Edward Grefenstette, Pushmeet Kohli, and Peter Battaglia. Compile: Compositional imitation learning and execution. *arXiv preprint arXiv:1812.01483*, 2018.
- Raphael Koster, Martin J Chadwick, Yi Chen, David Berron, Andrea Banino, Emrah Düzel, Demis Hassabis, and Dharshan Kumaran. Big-loop recurrence within the hippocampal system supports integration of information across episodes. *Neuron*, 99(6):1342–1354, 2018.
- Dharshan Kumaran, Demis Hassabis, and James L McClelland. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in cognitive sciences*, 20 (7):512–534, 2016.
- Francesco Locatello, Stefan Bauer, Mario Lucic, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. *CoRR*, abs/1811.12359, 2018.
- Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- Loic Matthey, Irina Higgins, Demis Hassabis, and Alexander Lerchner. dsprites: Disentanglement testing sprites dataset. https://github.com/deepmind/dsprites-dataset/, 2017.
- Josh Merel, Leonard Hasenclever, Alexandre Galashov, Arun Ahuja, Vu Pham, Greg Wayne, Yee Whye Teh, and Nicolas Heess. Neural probabilistic motor primitives for humanoid control. *arXiv preprint arXiv:1811.11711*, 2018.

- Tsendsuren Munkhdalai, Alessandro Sordoni, Tong Wang, and Adam Trischler. Metalearned neural memory. In *Advances in Neural Information Processing Systems*, pp. 13310–13321, 2019.
- Xue Bin Peng, Michael Chang, Grace Zhang, Pieter Abbeel, and Sergey Levine. Mcp: Learning composable hierarchical control with multiplicative compositional policies. *arXiv preprint arXiv:1905.09808*, 2019.
- Michele Pignatelli, Tomás J Ryan, Dheeraj S Roy, Chanel Lovett, Lillian M Smith, Shruti Muralidhar, and Susumu Tonegawa. Engram cell excitability state determines the efficacy of memory retrieval. *Neuron*, 101(2):274–284, 2019.
- William F Podlaski, Everton J Agnes, and Tim P Vogels. Context-modular memory networks support high-capacity, flexible, and robust associative memories. *bioRxiv*, 2020. doi: 10.1101/2020. 01.08.898528.
- Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. *arXiv* preprint arXiv:1505.05770, 2015.
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In Advances in Neural Information Processing Systems, pp. 348– 358, 2019.
- Sam Roweis and Zoubin Ghahramani. A unifying review of linear gaussian models. *Neural computation*, 11(2):305–345, 1999.
- Adam Santoro, Ryan Faulkner, David Raposo, Jack Rae, Mike Chrzanowski, Theophane Weber, Daan Wierstra, Oriol Vinyals, Razvan Pascanu, and Timothy Lillicrap. Relational recurrent neural networks. In *Advances in Neural Information Processing Systems*, pp. 7299–7310, 2018.
- Daniel L Schacter and Kevin P Madore. Remembering the past and imagining the future: Identifying and enhancing the contribution of episodic memory. *Memory Studies*, 9(3):245–255, 2016.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. End-to-end memory networks. In Advances in neural information processing systems, pp. 2440–2448, 2015.
- Endel Tulving et al. Episodic and semantic memory. Organization of memory, 1:381-403, 1972.
- Gido M van de Ven and Andreas S Tolias. Generative replay with feedback connections as a general strategy for continual learning. *arXiv preprint arXiv:1809.10635*, 2018.
- Joel Veness, Tor Lattimore, Avishkar Bhoopchand, Agnieszka Grabska-Barwinska, Christopher Mattern, and Peter Toth. Online learning with gated linear networks. *arXiv preprint arXiv:1712.01897*, 2017.
- Nicholas Watters, Loic Matthey, Christopher P Burgess, and Alexander Lerchner. Spatial broadcast decoder: A simple architecture for learning disentangled representations in vaes. *arXiv preprint arXiv:1901.07017*, 2019.
- Greg Wayne, Chia-Chun Hung, David Amos, Mehdi Mirza, Arun Ahuja, Agnieszka Grabska-Barwinska, Jack Rae, Piotr Mirowski, Joel Z Leibo, Adam Santoro, et al. Unsupervised predictive memory in a goal-directed agent. arXiv preprint arXiv:1803.10760, 2018.
- Max Welling. Product of experts. Scholarpedia, 2(10):3879, 2007.
- Christopher Williams, Felix V Agakov, and Stephen N Felderhof. Products of gaussians. In Advances in neural information processing systems, pp. 1017–1024, 2002.
- Yan Wu, Greg Wayne, Alex Graves, and Timothy Lillicrap. The kanerva machine: A generative distributed memory. *arXiv preprint arXiv:1804.01756*, 2018a.
- Yan Wu, Gregory Wayne, Karol Gregor, and Timothy Lillicrap. Learning attractor dynamics for generative memory. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), Advances in Neural Information Processing Systems 31, pp. 9379–9388. 2018b.

A SUPPLEMENTAL MATERIALS

A.1 EXPERIMENTAL DETAILS AND HYPER-PARAMETERS

A.1.1 HYPER-PARAMETERS

The model was trained using the Adam optimizer (Kingma & Ba, 2014) with learning rate between $5e^{-5}$ and $1e^{-3}$ and batch size 24.

For the RGB binding task, a high learning rate of $1e^{-3}$ was used and encouraged fast convergence to the sparse machine weights solution. For learning rate of $1e^{-4}$ or below, the k = 2 model initially underperformed the k = 1 model with the same total columns but then rapidly switched to the sparse solution and superior performance after roughly 200000 train batches.

Latent code sizes c were typically 50 but were 100 for the RGB binding task. The size of the history variable **h** was 10.

Convolutional encoders/decoders with ReLU activations were used

- Encoder: output channels [16, 32, 64, 128], kernel shapes (6, 6), strides 1
- Decoder: output channels [32, 16, 1 for grey-scale or 3 for RGB images], output shapes [(7, 7), (14, 14), (28, 28)], kernel shapes (4, 4), strides 2

except in the case of the dancing dSprites task where a small ResNet (2 layers of ResNet blocks with leaky ReLU activations each containing 2 convolutional layers with Kernel size 3, with an encoder output size 128 projected to c and using pool size 3 and stride 2) was used in order to improve reconstruction quality for dSprites.

A.1.2 TREATMENT OF ADDRESSING WEIGHTS **w**

For solving for the least-squares optimal read weights w, we used the matrix solver $matrix_solve_ls$ in TensorFlow, in Fast mode with L2 regularizer $\lambda = 0.1$ to 0.5, typically 0.35.

A.1.3 TREATMENT OF MACHINE ASSIGNMENT WEIGHTS r

Logits for choosing the k machine weights $\{r_i\}$, in $q_{\phi}(r_{1:k}^t | \mathbf{h}_{t-1}, \mathbf{z}_t)$ or $P_{\theta}(r_{1:k}^t | \mathbf{h}_{t-1})$, parametrizing a diagonal Gaussian $\mathcal{N}(\mu_r, \sigma_r^2)$ in the $\ln(r)$ space (see Supp. A.2), were created as follows. During reading and writing, we used $\mu_{\mathbf{r}} = \text{Linear}_1([\text{MLP}_1(\mathbf{z}_t), \mathbf{h}_{t-1}])$, $\sigma_{\mathbf{r}} = \text{Linear}_2([\text{MLP}_1(\mathbf{z}_t), \mathbf{h}_{t-1}])$. During generation, we used $\mu_{\mathbf{r}} = \text{Linear}_3(\mathbf{h})$, $\sigma_{\mathbf{r}} = \text{Linear}_4(\mathbf{h})$. MLP₁ had layer widths [40, 20, k]. Samples from the resulting Gaussian were passed through a SoftPlus function to generate effective machine observation noises (Wu et al., 2018a) $\sigma_i/\sqrt{r_i}$, and then squared, inverted and normalized to generate the overall machine weight $\gamma_i = \frac{r_i}{\sigma_i^2} / \sum_{j=1}^k \frac{r_j}{\sigma_j^2}$.

See Supp. A.4 for full machine choice algorithm and Supp. A.2 for definitions of the distributions in the generative and inference models.

A.1.4 SPEED TESTS

Speed tests were performed on a V100 GPU machine with 8 CPU cores, with memory operations assigned to CPU to encourage parallelization and encoder/decoder operations assigned to the GPU.

A.1.5 TUNING ANALYSIS

To analyze the tunings of dSprite reconstruction to dSprite properties (Fig. 3D-E), we used a template matching procedure. A template image with individual dSprites at their original individual positions in the stored images was matched to each machine's reconstruction \hat{x} via a cosine distance on the image pixel vector $\frac{\hat{x}_{\text{machine}i} \cdot \text{dSprite}_j}{\|\hat{x}_{\text{machine}i}\|\|\text{dSprite}_j\|}$ where *i* indexes over the *k* machines and *j* indexes over the two dSprites in each image.

A.2 GENERATIVE MODEL DEFINITION

The generative model (Fig. 1B, see Supp. A.9 for example conditional generations) is trained by maximizing a variational lower bound (Kingma & Welling, 2013) \mathcal{L} on $\ln p_{\theta}(\mathbf{x})$. For fixed machine weights r_i , we would use an ELBO

$$\mathcal{L} = \left\langle \ln p_{\theta}\left(\mathbf{x}|\mathbf{z}\right)\right\rangle_{q_{\phi}(\mathbf{z})} - \sum_{i=1}^{k} \left[\mathbb{D}_{\mathrm{KL}}[q_{\phi}\left(\mathbf{w}_{\mathbf{i}}\right)||p_{\theta}\left(\mathbf{w}_{\mathbf{i}}\right)] + \mathbb{D}_{\mathrm{KL}}[q_{\phi}\left(M_{i}\right)||p_{\theta}\left(M_{i}\right)]\right]$$
(5)

where $p_{\theta}(\mathbf{w_i}) = \mathcal{N}(0, 1)$. Here, we further consider the exponential weightings r_i in the generalized product to depend on a latent variable **h** that summarizes the history, via $p(r_i|\mathbf{h})$. This gives a joint distribution

$$p(\mathbf{z}, \{M_i\}_{i=1}^k, \{r_i\}_{i=1}^k, \mathbf{h}) = p(\mathbf{z}, \{M_i\}_{i=1}^k | \{r_i\}_{i=1}^k) \prod_{i=1}^k p(r_i | \mathbf{h}) p(\mathbf{h})$$
(6)

where $p(\mathbf{z}, \{M_i\}_{i=1}^k | \{r_i\}_{i=1}^k) \propto \prod_{i=1}^k p(\mathbf{z}, M_i)^{r_i}$ and results in additional KL divergence terms in the ELBO

$$-\sum_{i=1}^{k} \mathbb{D}_{\mathrm{KL}}[q(r_i)||p(r_i|\mathbf{h})] - \mathbb{D}_{\mathrm{KL}}[q(\mathbf{h})||p(\mathbf{h})]$$
(7)

The full joint distribution is

$$P_{\theta}(\mathbf{z}_{1:T}, M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, r_{1:k}^{1:T}, \mathbf{h}) = P_{\theta}(M_{1:k}) \prod_{t=1}^{T} \left(P_{\theta}(\mathbf{z}_{t} | M_{1:k}^{t}, \mathbf{w}_{1:k}^{t}, r_{1:k}^{t}) P_{\theta}(r_{1:k}^{t} | h) P_{\theta}(\mathbf{w}_{1:k}^{t}) \right) P_{\theta}(\mathbf{h})$$
(8)

$$=\prod_{t=1}^{T}\prod_{i=1}^{k}P_{\theta}(M_{i}^{t})\prod_{t=1}^{T}\left(P_{\theta}(\mathbf{z}_{t}|M_{1:k}^{t},\mathbf{w}_{1:k}^{t},r_{1:k}^{t})P_{\theta}(r_{1:k}^{t}|\mathbf{h}_{t-1})\prod_{i=1}^{k}P_{\theta}(\mathbf{w}_{i}^{t})\right)\prod_{t=1}^{T}P_{\theta}(\mathbf{h}_{t-1})$$
(9)

Marginalizing out $\mathbf{z}_{1:T}$, we have

$$P_{\theta}(M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, r_{1:k}^{1:T}, \mathbf{h}) = \prod_{t=1}^{T} \prod_{i=1}^{k} P_{\theta}(M_{i}^{t}) \prod_{t=1}^{T} \prod_{i=1}^{k} P_{\theta}(\mathbf{w}_{i}^{t}) \prod_{t=1}^{T} P_{\theta}(r_{1:k}^{t} | \mathbf{h}_{t-1}) \prod_{t=1}^{T} P_{\theta}(\mathbf{h}_{t-1})$$
(10)

 $P_{\theta}(r_{1:k}^t|\mathbf{h}_{t-1})$ is a top down generative model for the grouping across machines of content from the individual component machines of the product model. In order to be able to generate sequentially, it must only depend on the history up to but not including the outputs from the present timestep, i.e., \mathbf{h}_{t-1} . Rather than parameterizing \mathbf{r} as a distribution over categorical distributions, we instead parameterize $\ln(\mathbf{r})$ as a Gaussian (with trainable mean and diagonal variances), and then use a deterministic trainable network to produce \mathbf{r} .

 $P_{\theta}(\mathbf{w}_{i}^{t})$ is a standard Gaussian prior.

 $P_{\theta}(\mathbf{h}_{t-1})$ is a standard Gaussian prior.

(As an alternative prior on h, we can use a time-varying AR(1) process as the prior, $P_{\theta}(\mathbf{h}_{t-1}|\mathbf{h}_{t-2})$: this will allow the history variable to perform a random walk within an episode while slowly decaying to a standard Gaussian over time, as was used in Merel et al. (2018).)

 $P_{\theta}(M_i^t)$ is the trainable matrix Gaussian prior of each Kanerva Machine in the product model.

 $P_{\theta}(\mathbf{z_t}|M_{1:k}^t, \mathbf{w}_{1:k}^t, r_{1:k}^t)$ is the generation procedure for one step of the Product Kanerva Machine model as described elsewhere in this document. It will output the mean $\mathbf{w_i}^{\top}M_i$ from each machine and then combine them using the machine weights \mathbf{r} .

Note: we do not use any *additional* prior $P_{\theta}(\mathbf{z})$, such as a standard Gaussian, and likewise when we have an encoder from the image to the latent $e(\mathbf{z}|\mathbf{x})$ we do not subject it to a VAE-style standard Gaussian prior, instead just using a simple autoencoder with conv-net encoder $e(\mathbf{z}|\mathbf{x})$ and deconv-net decoder $d(\mathbf{x}|\mathbf{z})$ outputting the p parameter of a Bernoulli distribution for each image pixel.

A.2.1 INFERENCE MODEL

We use the following factorization of the approximate posterior to infer the hidden variables given a sequence of observations

$$q_{\phi}(M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, r_{1:k}^{1:T}, \mathbf{h} | \mathbf{z}_{1:T}) = q_{\phi}(M_{1:k}^{1:T} | \mathbf{w}_{1:k}^{1:T}, r_{1:k}^{1:T}, \mathbf{h}, \mathbf{z}_{1:T}) q_{\phi}(r_{1:k}^{1:T} | \mathbf{h}, \mathbf{z}_{1:T}) q_{\phi}(\mathbf{w}_{1:k}^{1:T} | \mathbf{z}_{1:T}) q_{\phi}(\mathbf{h} | \mathbf{z}_{$$

(11)

$$=\prod_{t=1}^{T} q_{\phi}(M_{1:k}^{t} | \mathbf{w}_{1:k}^{t}, r_{1:k}^{t}, \mathbf{h}, \mathbf{z}_{t}, M_{1:k}^{t-1}) \prod_{t=1}^{T} q_{\phi}(\mathbf{w}_{1:k}^{t} | \mathbf{z}_{t}, r_{1:k}^{t}, M_{1:k}^{t-1}) \prod_{t=1}^{T} q_{\phi}(r_{1:k}^{t} | \mathbf{h}_{t-1}, \mathbf{z}_{t}) \prod_{t=1}^{T} q_{\phi}(\mathbf{h}_{t-1} | \mathbf{z}_{1:t-1}) \prod_{t=1}^{T} q_{\phi}(\mathbf{x}_{1:k}^{t} | \mathbf{z}_{t}, r_{1:k}^{t}, M_{1:k}^{t-1}) \prod_{t=1}^{T} q_{\phi}(\mathbf{z}_{1:k}^{t} | \mathbf{z}_{t-1}, \mathbf{z}_{t}) \prod_{t=1}^{T} q_{\phi}(\mathbf{z}_{1:k}^{t} | \mathbf{z}_{t-1}, \mathbf{z}_{t-1}^{t}) \prod_{t=1}^{T} q_{\phi}(\mathbf{z}_{1:k}^{t} | \mathbf{z}_{t-1}^{t}) \prod_{t=1}^{T}$$

 $q_{\phi}(M_{1:k}^t | \mathbf{w}_{1:k}^t, r_{1:k}^t, \mathbf{h}, \mathbf{z}_t, M_{1:k}^{t-1})$ is the write step of our Product Kanerva Machine and is described elsewhere in this document

 $q_{\phi}(\mathbf{w}_{1:k}^{t}|\mathbf{z}_{t}, r_{1:k}^{t}, M_{1:k}^{t-1})$ is the "solve for **w** given query" step of our Product Kanerva Machine and is performed by least-squares optimization.

 $q_{\phi}(r_{1:k}^t | \mathbf{h}_{t-1}, \mathbf{z}_t)$ is a bottom-up inference model producing the machine weights variable **r**. Rather than parameterizing **r**, we instead parameterize $\ln(\mathbf{r})$ as a Gaussian (with trainable mean and diagonal variances), and then use a deterministic trainable network to produce **r**.

 $q_{\phi}(\mathbf{h}_{t-1}|\mathbf{z}_{1:t-1})$ is where we will use a superposition memory to store a record of previous \vec{z} and their associated \mathbf{r} variables which will be used to produce the history variable \mathbf{h} . The superposition buffer takes the form $\Omega_t = \frac{1}{t}\Psi([\mathbf{z}_t, r_{1:k}^t]) + \frac{t-1}{t}\Omega_{t-1}$ where Ψ is a trainable embedding function. Then the distribution over the history variable \mathbf{h} can be a diagonal Gaussian $q_{\phi}(\mathbf{h}_{t-1}|z_{1:t-1}) = N(\mu, \sigma)$ where $\mu = \mathrm{MLP}_a(\Omega_{t-1})$ and $\sigma = \mathrm{MLP}_b(\Omega_{t-1})$ (we used small MLPs with layer widths [10, 10] here for MLP_a and MLP_b).

A.2.2 ELBO

The full ELBO is

$$\ln P_{\theta}(\mathbf{z}_{1:T}) \geq \left\langle \frac{P_{\theta}(\mathbf{z}_{1:T}, M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, r_{1:k}^{1:T}, \mathbf{h})}{q_{\phi}(M_{1:k}^{1:T}, w_{1:k}^{1:T}, r_{1:k}^{1:T}, \mathbf{h}|\mathbf{z}_{1:T})} \right\rangle_{q_{\phi}(M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, r_{1:k}^{1:T}, h|\mathbf{z}_{1:T})}$$
(13)

$$= \left\langle \ln P_{\theta}(\mathbf{z}_{1:T} | M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, \mathbf{h}) \right\rangle_{q_{\phi}(M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, \mathbf{h}|z_{1:T})}$$
(14)

$$-\mathbb{D}_{\mathrm{KL}}[q_{\phi}(M_{1:k}^{1:T}, w_{1:k}^{1:T}, \mathbf{r}_{1:k}^{1:T}, \mathbf{h} | \mathbf{z}_{1:T})||P_{\theta}(M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, \mathbf{r}_{1:k}^{1:T}, \mathbf{h})]$$
(15)

$$= \left\langle \ln P_{\theta}(\mathbf{z}_{1:T} | M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, r_{1:k}^{1:T}, \mathbf{h}) \right\rangle_{q_{\phi}(M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, r_{1:k}^{1:T}, \mathbf{h} | \mathbf{z}_{1:T})}$$
(16)

$$-\mathbb{D}_{\mathrm{KL}}\left[\prod_{t=1}^{T} q_{\phi}(M_{1:k}^{t} | \mathbf{w}_{1:k}^{t}, r_{1:k}^{t}, \mathbf{h}, \mathbf{z}_{t}, M_{1:k}^{t-1}) || \prod_{t=1}^{T} \prod_{i=1}^{k} P_{\theta}(M_{i}^{t})\right]$$
(17)

$$-\mathbb{D}_{\mathrm{KL}}\left[\prod_{t=1}^{T} q_{\phi}(\mathbf{w}_{1:k}^{t} | \mathbf{z}_{t}, r_{1:k}^{t}, M_{1:k}^{t-1}) || \prod_{t=1}^{T} \prod_{i=1}^{k} P_{\theta}(\mathbf{w}_{i}^{t})\right]$$
(18)

$$-\mathbb{D}_{\mathrm{KL}}\left[\prod_{t=1}^{T} q_{\phi}(r_{1:k}^{t} | \mathbf{h}_{t-1}, \mathbf{z}_{t}) || \prod_{t=1}^{T} P_{\theta}(r_{1:k}^{t} | \mathbf{h}_{t-1})\right]$$
(19)

$$-\mathbb{D}_{\mathrm{KL}}\left[\prod_{t=1}^{I} q_{\phi}(\mathbf{h}_{t-1}|\mathbf{z}_{1:t-1})||\prod_{t=1}^{I} P_{\theta}(\mathbf{h}_{t-1})\right]$$
(20)

$$= \left\langle \ln P_{\theta}(\mathbf{z}_{1:T} | M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, \mathbf{h}) \right\rangle_{q_{\phi}(M_{1:k}^{1:T}, \mathbf{w}_{1:k}^{1:T}, \mathbf{h}|\mathbf{z}_{1:T})}$$
(21)

$$-\sum_{t=1}^{r}\sum_{i=1}^{\kappa} \mathbb{D}_{\mathrm{KL}}[q_{\phi}(M_{i}^{t}|\mathbf{w}_{i}^{t}, r_{i}^{t}, \mathbf{h}, \mathbf{z}_{t}, M_{i}^{t-1})||P_{\theta}(M_{i}^{t})]$$
(22)

$$-\sum_{t=1}^{T}\sum_{i=1}^{k}\mathbb{D}_{\mathrm{KL}}[q_{\phi}(\mathbf{w}_{i}^{t}|\mathbf{z}_{t}, r_{i}^{t}, M_{i}^{t-1})||P_{\theta}(\mathbf{w}_{i}^{t})]$$
(23)

$$-\sum_{t=1}^{I} \mathbb{D}_{\mathrm{KL}}[q_{\phi}(r_{1:k}^{t}|\mathbf{h}_{t-1}, \mathbf{z}_{t})||P_{\theta}(r_{1:k}^{t}|\mathbf{h}_{t-1})]$$
(24)

$$-\sum_{t=1}^{T} \mathbb{D}_{\mathrm{KL}}[q_{\phi}(\mathbf{h}_{t-1}|\mathbf{z}_{1:t-1})||P_{\theta}(\mathbf{h}_{t-1})]$$

$$(25)$$

Regarding the term $\mathbb{D}_{\mathrm{KL}}[q_{\phi}(r_{1:k}^{t}|\mathbf{h}_{t-1}, \mathbf{z}_{t})||P_{\theta}(r_{1:k}^{t}|\mathbf{h}_{t-1})]$, this should ideally be a KL between two Dirichlet distributions (Joo et al., 2019), i.e., between distributions over categorical distributions. Rather than parameterizing \mathbf{r} , we instead parameterize $\ln(\mathbf{r})$ as a Gaussian (with trainable mean and diagonal variances), and then use a deterministic trainable network to produce \mathbf{r} itself. We are then left with Gaussian KLs which are easy to evaluate and Gaussian variables which are easy to re-parametrize in training.

Note: For Fig. 1 and Fig. 2, we relaxed distributional constraints on \mathbf{w} in the loss function in order to lower variance, by removing the KL loss $\mathbb{D}_{\mathrm{KL}}[q_{\phi}(\mathbf{w}_{i}^{t}|\mathbf{z}_{t}, r_{i}^{t}, M_{i}^{t-1})||P_{\theta}(\mathbf{w}_{i}^{t})]$ on \mathbf{w} in the writing step (but not the reading step), and by using the mean \mathbf{w} rather than sampling it. The full model was used in Fig. 3. The mixture model of Supp. A.5 was trained without the KL penalty $\mathbb{D}_{\mathrm{KL}}[q_{\phi}(M_{i}^{t}|\mathbf{w}_{i}^{t}, r_{i}^{t}, \mathbf{h}, \mathbf{z}_{t}, M_{i}^{t-1})||P_{\theta}(M_{i}^{t})]$ on M to reduce variance, and also did not include KL terms for \mathbf{h} or \mathbf{r} .

A.2.3 SAMPLING FROM THE GENERATIVE MODEL

To generate full episodes autoregressively:

- We first sample priors $P_{\theta}(M_i)$, $P_{\theta}(\mathbf{w}_i)$ and $P_{\theta}(\mathbf{h})$ and then sample $P_{\theta}(r_{1:k}^t | \mathbf{h}_{t-1})$ and $P_{\theta}(\mathbf{z_t} | M_{1:k}^t, \mathbf{w}_{1:k}^t, r_{1:k}^t)$ to produce $r_{1:k}^1$ and then \mathbf{z}_1 .
- \mathbf{z}_1 is then decoded to an image $\hat{\mathbf{x}}_1 = d(\mathbf{z}_1)$, each pixel of the image rounded to 0/1 and then the image re-encoded as $e(\hat{\mathbf{x}}_1)$. We then query the memory with the re-encoded image $e(\hat{\mathbf{x}}_1)$ to obtain an updated \mathbf{z}_1 . This step is repeated several times, 12 times here, to allow the memory to "settle" into one of its stored attractor states (Wu et al., 2018b).
- We then write z_1 into the product memory M using the analytical memory update $q_{\phi}(M_{1:k}^t | \mathbf{w}_{1:k}^t, r_{1:k}^t, \mathbf{h}, \mathbf{z}_t, M_{1:k}^{t-1})$ of Eqs. 2-4 and $[\mathbf{z}_1, r_{1:k}^1]$ into the history \mathbf{h} via Ω by using $\Omega_t \leftarrow \frac{1}{t} \Psi([\mathbf{z}_t, r_{1:k}^t]) + \frac{t-1}{t} \Omega_{t-1}$.
- We then sample $q_{\phi}(\mathbf{h}_{t-1}|\mathbf{z}_{1:t-1})$ to produce \mathbf{h}_1 and sample $P_{\theta}(r^t|\mathbf{h}_{t-1})$ to produce r^2 .
- We then read the memory, using as read weights a draw from the priors on \mathbf{w}_i , $P_{\theta}(\mathbf{w}_i)$, and as machine weights our r^2 , which allows us to produce \mathbf{z}_2 .
- ...and so on, until finished generating.

Note that if a partial episode has been written to begin with, we will simply have M and Ω and hence **h** pre-initialized before starting this process rather than using their priors.

A.3 DERIVATION OF PRODUCT KANERVA WRITE AND READ OPERATIONS

A.3.1 REVIEW OF KANERVA MACHINE

To derive the Product Kanerva Machine, we first reformulate a single Kanerva Machine in terms of a precision matrix rather than covariance matrix representation.

For a single Kanerva Machine, recall that the posterior update of the memory distribution $P(M_i|\mathbf{z})$ is given by the Kalman filter-like form

$$R_i \leftarrow R_i + (\mathbf{z} - R_i \,\mathbf{w}_i) \frac{1}{\mathbf{w}_i^\top V_i \,\mathbf{w}_i + \sigma_i^2} \mathbf{w}_i^\top V_i$$
(26)

$$V_i \leftarrow V_i - V_i \, w_i \, \frac{1}{\mathbf{w}_i^\top \, V_i \, w_i + \sigma_i^2} \, \mathbf{w}_i^\top \, V_i \tag{27}$$

In addition, recall that we can analytically compute the mean and covariance of the joint distribution of z and M_i , as well as of the marginal distribution of z (integrating out M_i):

$$p(\mathbf{z}, M_i) = p(\mathbf{z}|M_i) p(M_i) \sim \mathcal{N}(\mu_i, \Sigma_i)$$
(28)

$$p_i(\mathbf{z}) = \int p(\mathbf{z}|M_i) \, p(M_i) \, \mathrm{d}M_i \sim \mathcal{N}\left(R_i \, \mathbf{w}_i, \underbrace{(\mathbf{w}_i^\top V_i \, \mathbf{w}_i + \sigma_i^2)}_{\Sigma_{z_i}} \cdot I\right)$$
(29)

where

$$\mu_i = \begin{bmatrix} R_i \, \mathbf{w_i} \\ \operatorname{vec} \left(R_i \right) \end{bmatrix} \tag{30}$$

$$\Sigma_{i} = \begin{bmatrix} \Sigma_{z_{i}} \Sigma_{c_{i}}^{\top} \\ \Sigma_{c_{i}}, V_{i} \end{bmatrix} \otimes I$$
(31)

The joint covariance is a Kronecker product of a block matrix, where the upper left block is 1×1 (a scalar), the upper right is $1 \times m_i$, the lower left is $m_i \times 1$ and the lower right is $m_i \times m_i$, and I is the $c \times c$ identity matrix.

To convert to the precision matrix representation, we can use the block matrix inversion rule to obtain the precision matrix for a single Kanerva Machine, similar to eqns. (10) and (11) in Williams et al. (2002):

$$\Lambda_i = \Sigma_i^{-1} = \begin{bmatrix} \sigma_i^{-2} & \Lambda_{c_i}^\top \\ \Lambda_{c_i} & \Lambda_{M_i} \end{bmatrix} \otimes I$$
(32)

where

$$\Lambda_{c_i} = -\sigma_i^{-2} \mathbf{w_i} \tag{33}$$

$$\Lambda_{M_i} = V_i^{-1} + \mathbf{w}_i \, \sigma_i^{-2} \, w_i^\top \tag{34}$$

$$= \left(V_i - V_i \mathbf{w}_i \left(\sigma_i^2 + \mathbf{w}_i^\top V_i \mathbf{w}_i\right)^{-1} \mathbf{w}_i^\top V_i\right)^{-1}$$
(35)

(36)

with the last step due to the Woodbury identity (Bishop, 2006). $\Lambda_{M_i}^{-1}$ is the updated posterior covariance matrix of the memory after an observation of z.

A.3.2 PRODUCTS OF KANERVA MACHINES

So far, we have dealt only with reformulating the notation for a single Kanerva Machine. What about the product of many Kanerva Machines? We can now consider the full joint distribution between observed z and all of the memory matrices, which we assume to factor according to the product of the individual joint Gaussian distributions between z and each memory:

$$p(\mathbf{z}, M_1, \dots, M_i, \dots) \propto \prod_{i=1}^k p(\mathbf{z}, M_i)$$
(37)

$$\sim \mathcal{N}\left(\mu, \Lambda^{-1}\right)$$
 (38)

From the mean and precision form of $p(\mathbf{z}, M_i)$, and using the fact that the precision matrix of a product of Gaussians is the sum of the individual precision matrices and that the mean is a precision weighted average of the individual means

$$\mu_{\text{product}} = \frac{\frac{\mu_1}{\sigma_1^2} + \frac{\mu_2}{\sigma_2^2}}{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}}$$
(39)

$$\Lambda_{\text{product}} = \Lambda_1 + \Lambda_2 \tag{40}$$

we have the joint precision matrix

$$\Lambda = \begin{bmatrix} \Lambda_z & \Lambda_{c_1}^\top & \Lambda_{c_2}^\top & \dots \\ \Lambda_{c_1} & \Lambda_{M_1} & \\ \Lambda_{c_2} & & \Lambda_{M_2} \\ \vdots & & & \end{bmatrix}$$
(41)

By completing the square, we can compute the parameters of the conditional $p(\mathbf{z}|M_{1:m})$:

$$\Lambda_z = \sigma_z^{-2} = \sum_{i=1}^k \sigma_i^{-2} \tag{42}$$

and the joint mean

$$\mu = \begin{bmatrix} \mu_{\mathbf{z}} \\ \operatorname{vec} (R_1) \\ \operatorname{vec} (R_2) \\ \vdots \end{bmatrix}$$
(43)

$$\mu_{\mathbf{z}} = \sum_{i=1}^{k} \gamma_i \, R_i \, \mathbf{w}_{\mathbf{i}} \tag{44}$$

where the coefficient γ_i is the normalised accuracy

$$\gamma_{i} = \frac{\sigma_{i}^{-2}}{\sum_{j=1}^{k} \sigma_{j}^{-2}}$$
(45)

and k is the number of machines.

Note that in the block matrix of equation 41, only the upper left corner couples between the blocks for the different machines/memories. Thus, the posterior update of the covariance, which does not depend on this term, is unmodified compared to the case of individual uncoupled Kanerva Machines.

The memory update rule for $p(M_i|\mathbf{z})$ is modified as:

$$\Delta = \mathbf{z} - \mu_{\mathbf{z}} \tag{46}$$

$$R_i \leftarrow R_i + \beta_i \,\Delta \,\mathbf{w}_i^\top \,V_i \tag{47}$$

$$V_i \leftarrow V_i - \beta_i \, V_i \, \mathbf{w}_i \mathbf{w}_i^\top \, V_i \tag{48}$$

where

$$\beta_i = \frac{1}{\mathbf{w}_i^\top V_i \, \mathbf{w}_i + \sigma_i^2} \tag{49}$$

Note that the only thing that makes this different than independent machine updates is the change in the prediction error term Δ which, which now couples the machines via μ_z from Eq. 44.

Readout takes the form of a simple precision weighted average μ_z of the outputs of each individual machine, again from Eq. 44.

A.3.3 GENERALIZED PRODUCTS OF KANERVA MACHINES

Following Cao & Fleet (2014) we further consider a "generalized product model" in which each term in the product of joint distributions may be weighted to a variable amount by raising it to a positive power r_i , such that

$$p(\mathbf{z}, M_i | r_i) \propto p(\mathbf{z}, M_i)^{r_i} \tag{51}$$

Since a Gaussian raised to a power is equivalent multiplication of the precision matrix by that power, we may simply replace $\Lambda_i \rightarrow \Lambda_i r_i$ in the above derivation of the product model, for each individual Kanerva Machine, and then proceed with the derivation as normal. The readout equations 44 and 45 for μ_z are

$$\gamma_i \to \frac{r_i/\sigma_i^2}{\sum_{j=1}^k r_j/\sigma_j^2} \tag{52}$$

$$\mu_{\mathbf{z}} = \sum_{i=1}^{k} \gamma_i R_i \,\mathbf{w}_{\mathbf{i}} \tag{53}$$

Meanwhile, in the update equations, we replace $V_i \rightarrow V_i/r_i$ and $\sigma_i^2 \rightarrow \sigma_i^2/r_i$, leading to:

$$\Delta = \mathbf{z} - \mu_{\mathbf{z}} \tag{54}$$

$$R_i \leftarrow R_i + \beta_i \,\Delta \,\mathbf{w}_i^{\top} \,V_i \tag{55}$$

$$V_i \leftarrow V_i - \beta_i \, V_i \, \mathbf{w}_i \mathbf{w}_i^\top \, V_i \tag{56}$$

where

$$\beta_i = \frac{1}{\mathbf{w}_i^\top V_i \, \mathbf{w}_i + \sigma_i^2 / r_i} \tag{57}$$

$$\mu_z = \frac{\sum_{i=1}^k \frac{r_i}{\sigma_i^2} R_i \mathbf{w_i}}{\sum_{j=1}^k \frac{r_j}{\sigma_i^2}}$$
(58)

This gives our update equations 2-4 in the main text.

Note that $1/\eta_i^2 := r_i/\sigma_i^2$ may also be treated as a single parameter here, and η_i generated as the output of a neural network. η_i then serves as an effective observation noise σ_i for the posterior update of machine *i*.

Remark. We can understand the coupling between machines during update by expanding the prediction term in equation 3:

$$\Delta = \left(\mathbf{z} - \sum_{j \neq i} \gamma_j R_j \mathbf{w}_j\right) - \gamma_i R_i \mathbf{w}_i$$
(59)

where terms in the bracket represent the residual from all other $j \neq i$ machines' predictions. Therefore, machine *i* is updated to reduce this residual, which may then change the residual for other machines. Because of this inter-dependency, the updates of machines are coupled and may take several iterations to converge.

In practice, we use a single iteration, making the model fully parallelizable over the k machines.

A.4 ALGORITHMS FOR WRITING AND READING

Here we present pseudocode for writing and reading in the Product Kanerva Machine.

For clarity, generation and optimization of the ELBO is treated separately in Supp. A.2.3.

Algorithm 1 Generalized Product Writing

Input: Input episode $\{x_t\}_{t=1}^T$, and Kanerva Machines $\{M_i\}_{i=1}^k$ with means R_i and column covariances V_i , total columns m, code size c, k machines of size $m_i = m/k$ columns each, and T is the episode length. **Initialization:** Each Kanerva Machine has a trainable prior mean matrix initialized as $R_i^0 \sim \mathcal{N}(0, I)$ and a diagonal prior column covariance with trainable scale, generated as a truncated unit normal $V_i^0 = \psi I$, where ψ is a trainable variable whose logarithm is initialized to $\ln(1.0)$.

The internal slot weighting $\mathbf{w_i}$ of each machine has a prior $p^0(\mathbf{w_i}) = \mathcal{N}(0, I)$. w_i is sampled from a normal distribution with mean produced by a least squares solution, and a diagonal covariance with standard deviation χ , where χ is a trainable variable whose logarithm is initialized to $\ln(0.3)$. Ω_0 is initialized to $\vec{0}$. for t = 1, ..., T do

 $\mathbf{z_t} \leftarrow e(\mathbf{x_t})$ where \mathbf{e} is the ConvNet encoder with output dim c

$$\begin{split} \mathbf{h}_{t-1} &\sim q_{\phi}(\mathbf{h}_{t-1}|z_{1:t-1}) = N(\mu,\sigma) \text{ where } \mu = \mathrm{MLP}_{a}(\Omega_{t-1}) \text{ and } \sigma = \mathrm{MLP}_{b}(\Omega_{t-1}). \\ \mathrm{Define } \mathcal{N}\left(\mu_{r},\sigma_{r}^{2}\right) \text{ with } \mu_{\mathbf{r}} = \mathrm{Linear}_{1}([\mathrm{MLP}_{1}(\mathbf{z}_{t}), \mathbf{h}_{t-1}]), \sigma_{r} = \mathrm{Linear}_{2}([\mathrm{MLP}_{1}(\mathbf{z}_{r}), \mathbf{h}_{t-1}]) \\ \eta &\sim \mathcal{N}\left(\mu_{r}, \sigma_{r}^{2}\right) \text{ with } \eta_{i} \text{ representing } \sigma_{i}/\sqrt{r_{i}} \\ \gamma: \gamma_{i} = \frac{1}{\eta_{i}^{2}}/\sum_{j=1}^{k} \frac{1}{\eta_{i}^{2}} \text{ representing } \gamma_{i} \rightarrow \frac{r_{i}/\sigma_{i}^{2}}{\sum_{j=1}^{k} r_{j}/\sigma_{j}^{2}} \\ \text{for } \mathbf{j}=\mathbf{1}, \dots, k \text{ do: } \\ \mathbf{w}_{\mathbf{j}} \leftarrow \mathrm{LeastSquaresOptimize}(M_{j}, \mathbf{z}_{\mathbf{t}}) \\ \text{end for } \\ \Delta \leftarrow \mathbf{z}_{\mathbf{t}} - \sum_{l=1}^{k} \gamma_{l} R_{l} \mathbf{w}_{\mathbf{l}} \\ \text{for } \mathbf{i}=\mathbf{1}, \dots, k \text{ do } \\ V_{i} \leftarrow V_{i} - \frac{1}{\mathbf{w}_{i}^{\top} V_{i} \mathbf{w}_{i} + \eta_{i}^{2}} V_{i} \mathbf{w}_{i} \mathbf{w}_{i}^{\top} V_{i} \\ R_{i} \leftarrow R_{i} + \Delta \frac{1}{\mathbf{w}_{i}^{\top} V_{i} \mathbf{w}_{i} + \eta_{i}^{2}} \mathbf{w}_{i}^{\top} V_{i} \\ \text{end for } \\ \Omega_{t} \leftarrow \frac{1}{t} \mathrm{Linear}_{0}([\mathbf{z}_{t}, \gamma]) + \frac{t-1}{t} \Omega_{t-1} \\ \mathrm{d} \text{ for } \end{split}$$

Algorithm 2 Generalized Product Reading

en

Input: Input episode $\{\mathbf{x}_t\}_{i=1}^T$, and filled Kanerva Machines $\{M_i\}_{i=1}^k$ and history variable \mathbf{h}_0 . **for** t = 1,...,T **do** $\mathbf{z}_t \leftarrow \mathbf{e}(\mathbf{x}_t)$ where e is the ConvNet encoder with output dim cDefine $\mathcal{N}(\mu_r, \sigma_r^2)$ with $\mu_{\mathbf{r}} = \text{Linear}_1([\text{MLP}_1(\mathbf{z}_t), \mathbf{h}_{t-1}]), \sigma_r = \text{Linear}_2([\text{MLP}_1(\mathbf{z}_t), \mathbf{h}_{t-1}])$ $\eta \sim \mathcal{N}(\mu_r, \sigma_r^2)$ with η_i representing $\sigma_i / \sqrt{r_i}$ $\gamma: \gamma_i = \frac{1}{\eta_i^2} / \sum_{j=1}^k \frac{1}{\eta_i^2}$ representing $\gamma_i \rightarrow \frac{r_i / \sigma_i^2}{\sum_{j=1}^k r_j / \sigma_j^2}$ **for** j=1,...,k **do** $\mathbf{w}_j \leftarrow \text{LeastSquaresOptimize}(M_j, \mathbf{z}_t)$ **end for** $\mu_{\mathbf{z}} \leftarrow \sum_{j=1}^k \gamma_j R_j \mathbf{w}_j$ **end for**

A.5 MIXTURE MODEL

A mixture Kanerva Machine model has k mixture coefficients γ_i , forming a categorical distribution. The categorical distribution is sampled to yield a one-hot vector $\hat{\gamma}_i$. We then have a read output

$$\mathbf{z}_{\text{read}} \leftarrow \sum_{i=1}^{k} \hat{\gamma}_i R_i \mathbf{w}_i \tag{60}$$

and the writing update for machine i is

$$V_i \leftarrow V_i - \hat{\gamma}_i \frac{1}{\mathbf{w}_i^\top V_i \, \mathbf{w}_i + \sigma_i^2} \, V_i \, \mathbf{w}_i \, \mathbf{w}_i^\top \, V_i \tag{61}$$

$$R_i \leftarrow R_i + \hat{\gamma}_i (\mathbf{z_t} - R_i \, \mathbf{w_i}) \frac{1}{\mathbf{w_i}^\top V_i \, \mathbf{w_i} + \sigma_i^2} \mathbf{w_i}^\top V_i$$
(62)

The generalized product model becomes a mixture model when \mathbf{r} is one-hot. To see this, note that in this case, in the Product Kanerva Machine, the prediction error Δ for the single machine *i* for which $r_i = 1$ inside the product becomes $z - R_i \mathbf{w}_i$ while the readout simply becomes $R_i \mathbf{w}_i$, as in a single Kanerva Machine, while in writing we reduce to the formula for β for a single Kanerva Machine. If *i* is such that $r_i = 0$ we have no readout from that machine and β in writing becomes 0 since the denominator becomes ∞ . Thus, choosing \mathbf{r} as one-hot thus corresponds to selecting a single machine while ignoring the others, while a mixture model corresponds to a stochastic choice of such a one-hot r.

We trained such a mixture mixture model using categorical reparametrization via Gumbel-SoftMax sampling (Jang et al., 2016; Maddison et al., 2016) of the machine choice variable. We verified that the Gumbel-SoftMax procedure was resulting in gradient flow using stop-gradient controls.

The mixture model shows MNIST digit class selective machine usage (Fig. 4A-C), but its performance degraded (Fig. 4D) as a fixed total number of slots was divided among an increasing number of machines k, in contrast to the robust performance of the product model in Fig. 1 of the main text.

Note that in the RGB binding task (Fig. 2), the network spontaneously found weights approaching $\{0, 1\}$, but it was able to explore a continuous space of soft weights in order to do so, unlike in a mixture model where the weights are one-hot once sampled.

A.6 FULL RGB BINDING TASK SELECTIVITY MATRIX

Fig. 5 shows the full machine usage matrix for k = 2 machines on the RGB binding task as a function of the R, G and B MNIST digit classes.

A.7 ADDITIONAL REPRESENTATIVE RECONSTRUCTIONS FROM DANCING DSPRITES TASK

Fig. 6 shows reconstructions from four different training runs on the dancing dSprites task with k = 4 machines, episode length T = 15 and m = 20 total columns.

A.8 DANCING DSPRITE SELECTIVITIES AND INVARIANCES

Tunings of individual machine reconstructions to dSprite properties were spatially localized and diverse, and seemed to approximately uniformly tile space across the machines, with machines 0 and 2 responsible for edges (Fig. 7A), but were invariant to shape, orientation and size (Fig. 7B). The slope of the curve with respect to size is an artifact of the template matching procedure and the fact that single machine reconstructions are typically smaller than the template dSprites they are matched to.



Figure 4: Mixture model result on queried MNIST reconstruction. A) Machine usage (norm of vector read from machine *i* during reconstruction, times the read weight for machine *i*) as a function of digit class for k = 2 machines. B) Machine usage as a function of MNIST digit class (0-9) for k = 3 machines. C) Queried single machine reconstructions, and full mixture reconstruction (far right) for a mixture model with 3 machines. Machines 0 and 2 are displaying a degenerate class-agnostic pattern corresponding to the average of all MNIST digits, while machine 1 is responsible for reconstructing this digit, consistent with its dominance for digit class 2 in panel B. D) Training curve for a mixture model at fixed total slots m = 30 and increasing k, showing performance degradation with k.

A.9 CONDITIONAL GENERATIONS

Example conditional generations from the Product Kanerva Machine with k = 4 and $m_i = 5$ after loading a short episode of four dancing dSprite images (Fig. 8). Twelve iterations of "attractor settling" were used (Wu et al., 2018b). In several of the generations the memory has simply retrieved a stored item, whereas in a few generations the model hallucinates noisy spatially localized patterns.



Figure 5: Full selectivity matrix for MNIST RBG binding task with a k = 2 product model. In this example, machine choice is sensitive to the digit class of the Red channel but relatively insensitive to that of the Blue or Green channels. Y: red digit class. X: blue digit class. Image: green digit class. Grayscale within each image: the machine assignment weight γ_1 .

0	0	1	2	3	Ρ	0	0	1	2	3	Ρ	Q	0	1	2	3	Ρ
••					••												
							•			•	••		Ĺ			,	Ţ
					,			·				••	•				•.
۰.			2		2	•,			•.		• ,						Ĺ
•					•					•							
			`		•	;					:						•
`.					•	2			7		:					,	
										,	•				•		
										4		·			•		•
					:	4		·			4						:
•			•		0							•		,		,	•
						i.	-							-			
					' •								ſ		ł		•
:												-					-
-•					•						•						
						•		-				·.					·,
			-														
•••																	
• •								•									
•					•	7			^		•						É.
••	-																
											ľ						

Figure 6: Additional dancing dSprite reconstructions showing object separation and spatial tuning. Four individual runs are shown (columns), with many examples from each run (rows). For each run, reconstructions were taken after each of 30 successive 5000 train steps (rows) starting at 100000 train batches. O: Original. 0-3: Reconstructions from each single machine when queried with the full image. P: Product reconstruction.



A) Additional representations of spatial tuning

Figure 7: Additional selectivities and invariances in the dancing dSprites task. A) Spatial tunings of each of the k = 4 machines to each of the 2 dSprite positions. Note how tunings across the four machines approximately tile space. B) Tunings to other properties of the individual dSprites: shape (oval, square or heart), orientation (0 to 2π) and scale. These properties are invariant across machines. Note that the slope in the curves with respect to scale (right) is due to the fact that single machine reconstructions are typically smaller than the template dSprites they are matched to.



Figure 8: Conditional generation example with k = 4. The top row of images ("presented episode") was loaded into memory, and then 10 generative samples were taken (bottom two rows) without further updates to the memory. 12 iterations of attractor settling (Wu et al., 2018b) were used to generate each image.