Leveraging Multi-task Structure for Cognitive Flexibility

Xiaoyu K Zhang (xiazhan.zhang@ugent.be)
Experimental Psychology Department, Ghent Univerity, Henri Dunantlaan 2, 9000 Gent, Belgium

Mehdi Senoussi (mehdi.senoussi@univ-tlse2.fr)
UMR 5263 CLLE, Université de Toulouse, 5 allées Antonio Machado, 31058 Toulouse, France

Tom Verguts (tom.verguts@ugent.be)
Experimental Psychology Department, Ghent Univerity, Henri Dunantlaan 2, 9000 Gent, Belgium

Abstract

Cognitive flexibility requires both retaining past knowledge (stability) and generalizing to new tasks. While attention mechanisms supporting this tradeoff have been studied, the complementary role of environmental structure—richness and specifically connectivity—remains underexplored. We systematically examine how these factors affect performance in MLPs and in attention-based models. Our results show that richer, more connected environments enhance both generalization and stability, especially for attention models, highlighting the importance of architecture–environment interactions in multitask learning.

Keywords: Multi-task structure; Attention; Cognitive flexibility; Stability; Generalization; Connectivity; Neural networks; Catastrophic forgetting

Introduction

Biological and artificial agents operate in dynamic environments where they must learn and switch between multiple tasks. This raises fundamental challenges in how knowledge is stored and shared across tasks. A key opportunity in multitask learning is the ability to generalize by reusing shared components. Both humans and neural networks benefit from forming shared task representations that support flexible transfer across tasks (Correa, Ho, Callaway, Daw, & Griffiths, 2023; Driscoll, Shenoy, & Sussillo, 2024; Johnston & Fusi, 2023; Yang, Joglekar, Song, Newsome, & Wang, 2019). At the same time, sharing representations poses a risk: Learning new tasks can interfere with prior knowledge, a phenomenon known as catastrophic forgetting (De Lange, van de Ven, & Tuytelaars, 2023; French, 1999; Grossberg, 1980; McCloskey & Cohen, 1989; McClelland, McNaughton, & O'Reilly, 1995; Kim & Han, 2023). This reflects a broader tradeoff between generalization and stability, where improving one often harms the other (Musslick & Cohen, 2021).

While prior research has mainly focused on architectural solutions to address this tradeoff, we highlight the often-overlooked but complementary role of environmental structure (Dorrell et al., 2025; Saxe, McClelland, & Ganguli, 2019; Lee, Mannelli, Clopath, Goldt, & Saxe, 2022). We design a multitask environment defined by combinations of sensory and motor cues. We characterize this environment in terms of richness and connectivity and compare standard MLPs with novel attention-based models that dynamically select relevant information to mitigate forgetting (Hummos, 2023; Sommers, Thorat, Anthes, & Kietzmann, 2025; Verbeke & Verguts, 2022) on their ability to retain and generalize task information as a function of richness and connectivity.

Methods

Richness and Connectivity in the task structure

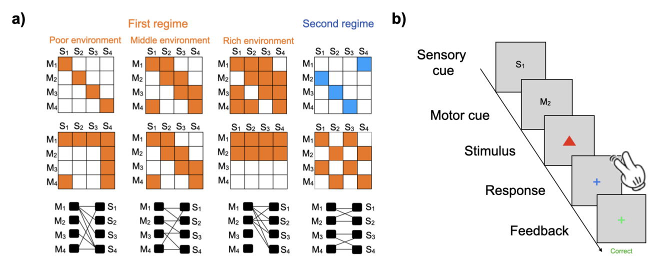

Figure 1: Experimental design for Multi task. (a) Row 1: Multi-4. The first regime has 4/8/12 tasks; the second regime has 4 new ones. Rows 2-3: Different levels of connectivity in the first regime (Row 2), each with their corresponding graphs (Row 3). (b) Trial: sensory cue, motor cue, stimulus, response, feedback.

A Multi-n task structure is an environment with n sensory (e.g. color) and n motor (e.g., left hand) cues. Figure 1a illustrates the Multi-4 task structure, combining 4 sensory and 4 motor cues into 4^2 = 16 tasks, each defined by one cue pair. A regime is a collection of jointly trained tasks; we distinguish the first regime (tasks trained first) from the second regime (tasks trained second). Each regime is divided into trials. Each trial (Figure 1b) includes a sensory cue, motor cue, stimulus, response, and feedback. Stimuli vary over 16 combinations; feedback guides learning of cue-stimulus-response mappings. To explore environmental factors, we used Multi-4 (Figure 1a top) with three richness levels in the first regime poor (4 tasks), middle (8), and rich (12). The second regime always includes 4 new tasks not used during first regime training. Beyond richness, we examine connectivity. 8 out of 16 tasks are selected for the first regime (Figure 1a middle). While many such selections are possible, we group regimes as equivalent if they differ only by cue relabeling or transposition. This yields 32 unique regime configurations (Faradžev, 1978). Each regime forms a bipartite graph where cues are vertices and tasks are edges (Figure 1a bottom).A regime is connected if all cues are linked by paths; otherwise, it's disconnected (see Figure 1a middle-bottom). Of the 32 unique regimes, 17 are connected, 15 disconnected. Connectivity is quantified by average shortest path length (ASPL),which measures the average minimum steps between all cue pairs. Examples in Figure 1a middle show ASPLs of 2, 2.07 (connected) and ∞ (disconnected).

Models

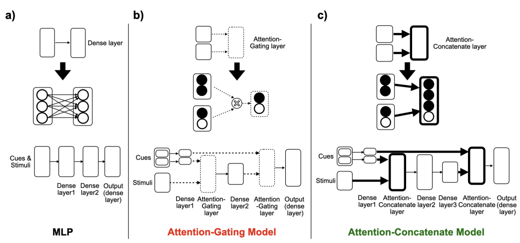

Figure 2: Model architectures. (a) MLP: joint cue-stimulus input. (b) Attention-Gating: cues filter stimulus. (c) Attention-Concatenate: cue and stimulus merged. Lines: dotted = gating, solid = concatenation.

As a baseline, we used MLPs (Figure 2a) that process one-hot encoded sensory and motor cues along with stimuli (input 24, output 8), with 3 and 4 dense layers, trained over 50×5000 trials using Adam, cross-entropy, Xavier initialization, and sigmoid activation functions. Attention models (Figure 2b–c) extend MLPs with cue-guided attention: gating (cues modulate stimulus features via gates) or concatenation (cues merge with stimuli across layers). Both were tested with/without a bottleneck. Training matched MLPs.

Results

Models train on first regime with feedback, then generalize to second regime without feedback. After training on second regime, stability is tested on first regime without feedback. As all models reached 100% accuracy in both training phases, we focus on generalization and stability.

For Multi-4 (Figure 3a rows 1-3), to analyze the impact of environment richness and connectivity on the model performance, we compare MLP_2 (selected for its superior performance among MLP models) and Concat_2 (representing attention models, as all attention models exhibit similar performance). First, we evaluate their generalization and stability across poor, middle, and rich environments(Figure 3a, rows 1–3). Concat_2 outperforms MLP_2 in generalization and stability, especially with limited training. Both improve with richness, but only Concat_2 reaches 100% accuracy in rich settings.

To test connectivity effects (Figure 3a, rows 4–7), MLP_2 shows minor gains in connected regimes. Concat_2 performs well when connected (83–98% generalization, 98–100% stability) but drops in disconnected ones, showing its reliance on structure.

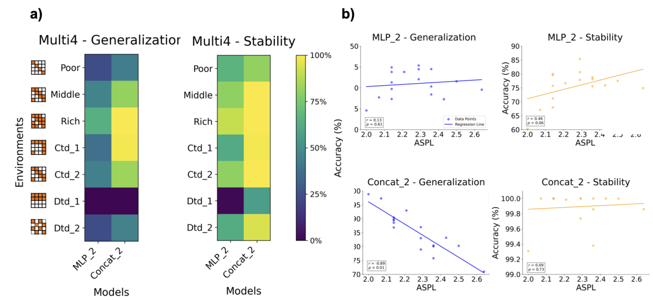

Figure 3: Multi-4 results. (a) Concat_2 outperforms MLP_2 in both generalization and stability, especially in rich or connected regimes. MLP_2 struggles in low-richness or disconnected settings. (b) MLP_2 shows weak correlation with connectivity; Concat_2 generalization improves with higher connectivity. Stability remains at ceiling. (Ctd = Connected; Dtd = Disconnected)

Figure 3b shows that Concat_2 generalization correlates strongly with connectivity (ASPL, $r = 0.89$), while MLP_2 shows no clear pattern. Stability remains near ceiling for Concat_2, with weak correlations.

Discussion

We studied how cognitive architecture interacts with task structure in multitask learning, comparing MLPs and two attention-based models across varying levels of richness and connectivity. Our results show that richer and more connected environments improve performance for all models, but especially for attention models, which consistently generalize better and retain prior knowledge. Attention models uniquely benefit from task connectivity, likely due to better information sharing across tasks. Unlike previous work that combats forgetting through replay or regularization (De Lange et al., 2023; Verbeke & Verguts, 2019), our approach achieves stability through the combination of architectural design and structured training environments. Another aspect of task structure is stimulus ordering (curriculum learning). Its effectiveness is minimal in datasets like CIFAR, but substantial in structured tasks like arithmetic (Wu, Dyer, & Neyshabur, 2021; Matiisen, Oliver, Cohen, & Schulman, 2017). Curriculum learning may be more effective in connected environments, where tasks are built on each other, mirroring human learning (Dekker, Otto, & Summerfield, 2022). Related ideas in reinforcement learning, such as modularity and bottleneck states, also highlight how task structure enables flexible learning (Franklin & Frank, 2018; Tomov, Schulz, & Gershman, 2021; Şimşek & Barto, 2008; Stachenfeld, Botvinick, & Gershman, 2017). Future work will scale up these simulations and test whether humans show similar sensitivity to environmental structure, potentially linking task design and cognitive flexibility more directly.

Acknowledgments

This research is funded by the China Scholarship Council (CSC). We'd like to thank Steven Van Overberghe and Tom Lauwaerts for bringing our attention to the work around Isomorph-Free Exhaustive Generation.

References

Correa, C. G., Ho, M. K., Callaway, F., Daw, N. D., & Griffiths, T. L. (2023, June). Humans decompose tasks by trading off utility and computational cost. *PLOS Computational Biology*, *19*(6), e1011087.

Şimşek, O., & Barto, A. (2008). Skill characterization based on betweenness. In D. Koller, D. Schuurmans, Y. Bengio, & L. Bottou (Eds.), *Advances in neural information processing systems* (Vol. 21). Curran Associates, Inc.

Dekker, R. B., Otto, F., & Summerfield, C. (2022). Curriculum learning for human compositional generalization. *Proceeding of the national academy of science*, *119*, 1–12.

De Lange, M., van de Ven, G. M., & Tuytelaars, T. (2023). Continual evaluation for lifelong learning: Identifying the stability gap. In *ICLR* (pp. 1–21).

Dorrell, W., Hsu, K., Hollingsworth, L., Lee, J. H., Wu, J., Finn, C., ... Whittington, J. C. (2025). *Range, not independence, drives modularity in biologically inspired representations.*

Driscoll, L. N., Shenoy, K., & Sussillo, D. (2024, July). Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. *Nature Neuroscience*, *27*(7), 1349–1363.

Faradžev, I. (1978). Constructive enumeration of combinatorial objects. In *Problèmes combinatoires et théorie des graphes* (pp. 131–135).

Franklin, N. T., & Frank, M. J. (2018). Compositional clustering in task structure learning. *PLoS Computational Biology*, 1–25.

French, R. M. (1999). Catastrophic forgetting in connectionists networks. *Trends in Cognitive Sciences*, *3*(4), 128–135.

Grossberg, S. (1980, January). How does a brain build a cognitive code? *Psychological review*, *87*(1), 1–51.

Hummos, A. (2023). Thalamus: a brain-inspired algorithm for biologically-plausible continual learning and disentangled representations. In *The eleventh international conference on learning representations.*

Johnston, W. J., & Fusi, S. (2023). Abstract representations emerge naturally in neural networks trained to perform multiple tasks. *Nature Communications*, *14*(1), 1040.

Kim, D., & Han, B. (2023). *On the stability-plasticity dilemma of class-incremental learning.*

Lee, S., Mannelli, S. S., Clopath, C., Goldt, S., & Saxe, A. (2022, 17–23 Jul). Maslow's hammer in catastrophic forgetting: Node re-use vs. node activation. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, & S. Sabato (Eds.), *Proceedings of the 39th international conference on machine learning* (Vol. 162, pp. 12455–12477). PMLR.

Matiisen, T., Oliver, A., Cohen, T., & Schulman, J. (2017, 07). Teacher-student curriculum learning. *IEEE Transactions on Neural Networks and Learning Systems*, *PP*.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). 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.

McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In G. H. Bower (Ed.), (Vol. 24, p. 109-165). Academic Press. doi: https://doi.org/10.1016/S0079-7421(08)60536-8

Musslick, S., & Cohen, J. D. (2021). Rationalizing constraints on the capacity for cognitive control. *Trends in Cognitive Sciences*, *25*(9), 757–775.

Saxe, A. M., McClelland, J. L., & Ganguli, S. (2019). A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences*, *116*(23), 11537-11546.

Sommers, R. P., Thorat, S., Anthes, D., & Kietzmann, T. C. (2025). Sparks of cognitive flexibility: self-guided context inference for flexible stimulus-response mapping by attentional routing. *arXiv preprint arXiv:2502.15634*.

Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature neuroscience*, *20*(11), 1643–1653.

Tomov, M. S., Schulz, E., & Gershman, S. J. (2021, January). Multi-task reinforcement learning in humans. *Nature Human Behaviour*, *5*(6), 764–773.

Verbeke, P., & Verguts, T. (2019). Learning to synchronize: How biological agents can couple neural task modules for dealing with the stability-plasticity dilemma. *PLoS computational biology*.

Verbeke, P., & Verguts, T. (2022). Using top-down modulation to optimally balance shared versus separated task representations. *Neural Networks*, *146*, 256-271. doi: https://doi.org/10.1016/j.neunet.2021.11.030

Wu, X., Dyer, E., & Neyshabur, B. (2021, February). When Do Curricula Work? In *International Conference on Learning Representations.*

Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T., & Wang, X.-j. (2019). Task representations in neural networks trained to perform many cognitive tasks. *Nature Neuroscience*, *22*(February).