The Dynamic and Structured Nature of Learning and Memory

Hanqi Zhou (hanqi.zhou@uni-tuebingen.de)
University of Tübingen, Maria-von-Linden-Str. 6
Tübingen, 72076 Germany

Álvaro Tejero-Cantero* **(alvaro.tejero@uni-tuebingen.de)**
University of Tübingen, Maria-von-Linden-Str. 6
Tübingen, 72076 Germany

Charley M. Wu* **(charley.wu@uni-tuebingen.de)**
University of Tübingen, Maria-von-Linden-Str. 6
Tübingen, 72076 Germany

* These authors contributed equally

Abstract

The dynamic and relational structure of memory has important implications for human learning, education, and artificial intelligence. In this work, we propose the Hierarchical Ornstein–Uhlenbeck Model (HOUM) to capture the dynamics of learning and forgetting, with a GraphHOUM extension leveraging the relational structure of knowledge. This combined approach models the dynamic interplay and structured organization of memory traces in short-term and long-term memory, predicting future recall probability. We demonstrate the effectiveness of our model by outperforming previous models on their own datasets. This work provides important insights into human learning and memory, and lays the foundations for developing future tools using artificial intelligence to recommend learning schedules for self-directed learners.

Keywords: Memory and Learning, Spaced Repetition, Connectionism, Stochastic Process, Knowledge Graphs

Introduction

The interconnected nature of learning and memory has long been a central focus in cognitive psychology and neuroscience, with important implications for education and artificial intelligence. A key observation is that the spacing of study sessions can enhance memory retention (Ebbinghaus, 1885), inspiring various learning techniques utilizing a *spaced repetition* schedule (Settles & Meeder, 2016; Walsh, Gluck, Gunzelmann, Jastrzembski, & Krusmark, 2018).

However, human learning is also characterized by the relational structure of knowledge (Rumelhart, 2017; Piaget, 1970). Accordingly, a number of recent studies suggest that memory traces are formed in a structured manner, where learning one knowledge component can influence others connected to it (Lynn & Bassett, 2020; Karuza, Thompson-Schill, & Bassett, 2016). Thus, a complete picture of human learning not only needs to account for the dynamic maintenance of individual memory traces, but also the structured nature of knowledge.

Here, we propose a novel Hierarchical Ornstein–Uhlenbeck Model (HOUM) to describe the dynamic interplay among

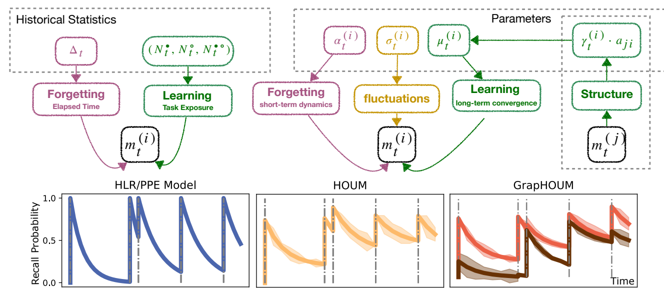

Figure 1: Regression Models HLR/PPE vs. Hierarchical Bayesian Models HOUM/GraphHOUM. Two different knowledge components are presented in the GraphHOUM panel to illustrate relational interactions. $N_t^{\bullet}, N_t^{\circ}, N_t^{\bullet\circ}$ represent the number of recalled, forgotten, and total interactions up to time t. Vertical dashed lines indicate rehearsals.

memory traces in both short-term and long-term memory. We also extend this model to account for the structured organization of knowledge (GraphHOUM). By combining both dynamics and structure, our model aims to provide a more comprehensive and accurate representation of human learning and memory mechanisms, which we intend to deploy as an AI tutor to support self-directed learning in future work.

Models of Memory and Learning

Cognitive models of learning and memory retention have been developed to predict the future recall probability of a learner (Lee & Wagenmakers, 2014) based on past performances. These models are often based on three key findings: i) the *power law of learning* describing increased performance with repeated practice, ii) the *power law of forgetting* describing performance decline as a function of elapsed time, and iii) the *spacing effect* showing distributing practice over time enhances retention compared to “cramming”.

Two prominent examples are the *Half-Life Regression* (HLR; Settles & Meeder, 2016) and *Predictive Performance Equation* (PPE; Walsh et al., 2018) models.

For a given knowledge component, such as the midpoint formula for a segment, HLR predicts the probability of correct recall m_t at time t as a function of the time elapsed since the last exposure Δ_t and the number of previous rehearsals $x_t := (N_t^{\bullet}, N_t^{\circ}, N_t^{\bullet\circ})$:

$$m_t = 2^{-\Delta_t/h_t}, \quad \text{where} \quad h_t := 2^{\Theta \cdot x_t} \tag{1}$$

where $N_t^{\bullet}, N_t^{\circ}$ and $N_t^{\bullet\circ}$ represent the number of recalled ($\bullet$), forgotten ($\circ$), and total ($\bullet\circ$) interactions up to time t. The free parameter $\Theta \in \mathbb{R}^3$ modulates the influence of historical statistics N_t on future performance.

Similarly, PPE uses a learning rate β and a forgetting parameter α along with an additional parameter a representing prior knowledge:

$$m_t = (a + N_t^{\bullet\circ})^{\beta}\, \Delta_t^{-\alpha} \tag{2}$$

As shown in Figure 1, both models assume that performance on a task improves with experience $N_t^{\bullet\circ}$, but is offset by forgetting over time Δ_t. The specific forms of the dual process of learning and forgetting differ between the two models, but they share the fundamental assumption that both are important factors when modeling human performance.

Hierarchical Ornstein–Uhlenbeck Model (HOUM). We propose a Hierarchical Ornstein–Uhlenbeck Model (HOUM) as a new framework for modeling learning and forgetting under uncertainty, where hierarchical Bayesian inference allows us to amortize parameter learning across learners. At its core, HOUM uses an Ornstein-Uhlenbeck (OU) process, a mean-reverting process common in physics and neuroscience, to describe dynamics of a memory trace m_t:

$$\mathrm{d}m_t = \alpha\,(\mu - m_t)\,\mathrm{d}t + \sigma\,\mathrm{d}W_t, \tag{3}$$

where α represents the rate of reversion to a long-term baseline μ, where μ indicates the strength of long-term memory. The OU process also accounts for random fluctuations, where W_t represents a Wiener process of variance σ. The strength of a memory trace as a function of time is given by the solution of Eq. 3, depending only on the parameters $\Theta := \{\alpha, \mu, \sigma\}$:

$$m_t = \underbrace{\mathrm{e}^{-\alpha\Delta_t}\, m_{t-1}}_{\substack{\text{short-term}\\ \text{memory dynamics}}} + \underbrace{\mu\left(1 - \mathrm{e}^{-\alpha\Delta_t}\right)}_{\substack{\text{long-term}\\ \text{memory convergence}}} + \underbrace{\sigma \int_{t-\Delta_t}^{t} \mathrm{e}^{-\alpha(t-s)}\,\mathrm{d}W_s}_{\text{random fluctuations}} \tag{4}$$

Structured Memory Traces: GraphHOUM. We can extend the model by accounting for the effect of the structured interdependence of knowledge components on memory. We incorporate learning dependencies into the long-term memory strength via parameters a_{ij}, quantifying how learning knowledge component i empowers the learner to tackle j:

$$\mu_t := (1-\omega) \underbrace{\sum_{j \neq i} a_{ij} m_t^{(i)}}_{\substack{\text{structure}\\ \text{empowerment}}} + \omega \underbrace{\log(N_t^{\bullet}/N_t^{\bullet\circ})}_{\substack{\text{feedback}\\ \text{empowerment}}} \tag{5}$$

The dependencies among knowledge components model the spreading activation of memory traces across the knowledge graph. An additional term describes the effect of positive feedback; the parameter ω quantifies their relative importance.

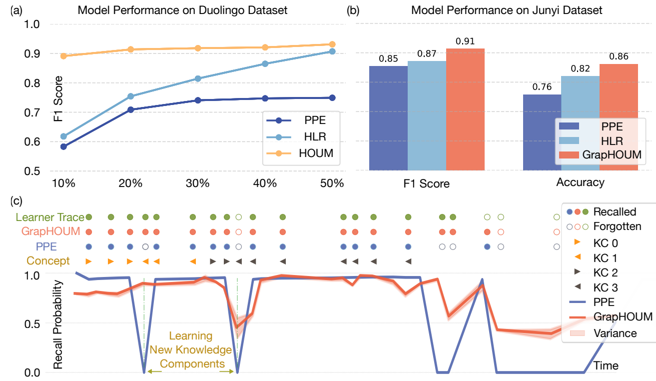

Figure 2: Experimental results. **a**) F1 score of predictions from HLR, PPE, HOUM with increasing fractions of training data. **b**) F1 score and accuracy of predictions from HLR, PPE, GrapHOUM on 50% training data. **c**) Predicted recall probability of knowledge components (KCs) from PPE vs. GrapHOUM. Green dots represent actual learner performance (●/○: recall/forget); other colors respective model predictions. Colored triangles ►, ◄ mark presentation times of KCs.

Results

Experiments. We use learning histories shared publicly by online educational platforms Duolingo (Settles & Meeder, 2016) and Junyi (Chang, Hsu, & Chen, 2015). The Duolingo dataset on second-language learning comprises 13,854,226 interactions with 19,279 knowledge components while the Junyi dataset consists 25,925,992 interactions with 722 mathematics knowledge components organized in a knowledge graph A, where each $a_{ij} \in A$ represents the dependency of knowledge component i on j. At time t, a learner l is presented with a knowledge component c_t, and we observe their binary performance y_t. The interaction histories of each learner can be represented as $\mathscr{H}^T = (c_t, y_t)_{t=1}^T$.

Prediction Performance. In education, it is important to predict a learner's future performance based on $\mathscr{H}^T$ to select appropriate further teaching material. For PPE and HLR, we use gradient descent to obtain model parameters. For HOUM and GraphHOUM, the tractable likelihood of the OU process $p(\mathscr{H}^t|\Theta)$ allows for hierarchical Bayesian inference $p(\Theta|\mathscr{H}^t) \propto p(\mathscr{H}^t|\Theta)p(\Theta)$. We use variational inference for the posterior distribution of the parameters $p(\Theta|\mathscr{H}^t)$. Conditioned on the inferred parameters, we can then make predictions about the distribution $p(y_{t+1}|\Theta)$ of future performance. Figure 2a compares performance across models. We selected 100 sequential interactions of learners from the Duolingo dataset, trained on the first 10-50% of the data, and predicted the rest. HOUM robustly outperformed HLR and PPE. HOUM's advantage comes from using hierarchical Bayesian inference, thus allowing it to better generalize across learners.

In Figure 2b, we compare GraphHOUM with HLR and PPE on the structured Junyi dataset (trained on 50% of the data). By incorporating the connected structure of knowledge components, GraphHOUM makes better predictions than HLR and PPE.

Latent Memory Trace Dynamics. Figure 2c illustrates how GrapHOUM leverages the structure of knowledge to model the dynamics of memory, compared against PPE which lacks this capability. In the example, vertical dashed lines denote times at which learners are presented with a new knowledge component (KC 1, ◀; Fig. 2c). Standard, unstructured regression models struggle to predict performance on a novel component. For instance, in the first presentation of ◀, learners are studying the *midpoint* knowledge component after having learned about *distance*. The PPE model predicts a recall probability of 0 for the new knowledge component because there is no prior exposure. In contrast, GrapHOUM leverages the dependence of *midpoint* on *distance*, and is able to generalize recall probability (with uncertainty) based on it.

Conclusion

We present an innovative framework for modeling the dynamic and structured nature of human learning and memory. In two experiments on real-world online-learning datasets, our HOUM and GrapHOUM models outperform previous HLR and PPE models on their own datasets. Taken together, our results suggest that leveraging knowledge structure can significantly enhance the performance and interpretability of memory models. This framework has the potential to advance our understanding of cognitive processes and contribute to the development of more effective strategies, tools, and interventions in educational settings. We intend to extend our scope in future work by inferring knowledge graphs in settings where the ground truth is not known.

Acknowledgments

This work is supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A and funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy–EXC2064/1–390727645. We thank our Robert Bamler for helpful insights and feedback on this work.

References

Chang, H.-S., Hsu, H.-J., & Chen, K.-T. (2015). Modeling exercise relationships in e-learning: A unified approach. In *Edm* (pp. 532–535).

Ebbinghaus, H. (1885). *Über das Gedächtnis: Untersuchungen zur experimentellen Psychologie*. Leipzig: Duncker & Humblot.

Karuza, E. A., Thompson-Schill, S. L., & Bassett, D. S. (2016). Local patterns to global architectures: influences of network topology on human learning. *Trends in cognitive sciences*, *20*(8), 629–640.

Lee, M. D., & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge university press.

Lynn, C. W., & Bassett, D. S. (2020). How humans learn and represent networks. *Proceedings of the National Academy of Sciences*, *117*(47), 29407–29415.

Piaget, J. (1970). Science of education and the psychology of the child. trans. d. coltman.

Rumelhart, D. E. (2017). Schemata: The building blocks of cognition. In *Theoretical issues in reading comprehension* (pp. 33–58). Routledge.

Settles, B., & Meeder, B. (2016). A trainable spaced repetition model for language learning. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 1: long papers)* (pp. 1848–1858).

Walsh, M. M., Gluck, K. A., Gunzelmann, G., Jastrzembski, T., & Krusmark, M. (2018). Evaluating the theoretic adequacy and applied potential of computational models of the spacing effect. *Cognitive science*, *42*, 644–691.