A Novel Board Game Paradigm for Studying the Acquisition of Overlapping Schemas

AUTHOR

Hannah Tongxin Zeng

ADVISOR

Christopher A. Baldassano
Assistant Professor of Psychology

A thesis submitted in partial fulfillment of the requirements for the degree of

Bachelor of Arts in Cognitive Science

Barnard College, Columbia University

May 6, 2024

Abstract

How does the brain acquire multiple overlapping schemas, the structured bodies of knowledge supporting our categorization and anticipation about what will come? We investigated this question using a novel board game paradigm where participants learned two games, Four-in-a-row and Knobby, with shared board configurations but differing win conditions requiring distinct optimal strategies. We conducted analysis on both game-level and move-level to compare effectiveness of schema acquisition in blocked and interleaved learning conditions (N=26). Several key findings emerged: 1) Blocked learning facilitated better overall performance, suggesting minimal interference between schema. 2) A primacy effect, where the first schema learned gained a substantial and persistent advantage. 3) Blocked learning renders better separation between schema, which is reflected as better moves to the specific game. Additionally, the study highlights the utility of game-based learning environments in cognitive research, offering a dynamic and engaging platform that closely mirrors real-world learning scenarios. This research contributes to our understanding of how different learning strategies affect the retention and application of new knowledge, particularly in settings where multiple schemas are involved.

1 Introduction

Our perception of the world is constantly scaffolded by the knowledge we gained from the past. The brain extracts structured bodies of knowledge, or schemas, about how events are typically unfolded in a given context, from prior experience to generate predictions about what is likely to occur in the future [10, 2, 4, 5, 22]

Schema plays various roles in our day-to-day life: schemas about social behaviors help you foresee the social repercussions of a behavior; motor schemas for producing and evaluating muscle coordination and movement [24]; story schemas that underpins how we predict and comprehend narrative structures [16] and many more. Specifically, our study is focused on event schemas that assist the anticipation of events over time, particularly the next possible actions in a given game state.

1.1 Learning overlapping schemas

To understand the world, we need to learn and maintain many new schemas with very little interference to existing ones.

A new schema is often created to resolve any inconsistencies between new information and existing schema, also called restructuring [27]. Consider the following example: suppose that a tea-lover who would make a cup of tea every morning starts learning to make pour-over coffee. Initially, they would see both tea-making and coffee-making as a processes of pouring hot water over raw materials to produce a beverage. Over time, however, they would realize the unique specificities of coffee, such as the influence of water temperature, speed and timing at which water is poured, and other techniques on the flavors of the final product. A new body of knowledge is thus needed for anticipating how the coffee will taste differently in a certain water temperature, pressure, and so on. Formation of a new schema is not always necessary when new information can be integrated into existing schema without fundamentally altering it (accretation [27]), or only demand minor tweaks in the schema (tuning [27]). For instance, a Spanish speaker starts learning Portuguese will notice similar patterns in the grammatical principles and vocabularies of the two languages, and thus naturally extend what they already know about Spanish to understand and produce Portuguese in a conversation. In this case, the existing schema about Spanish is modified to capture the domain of Portuguese, rather than forming an entirely separate schema. Our study probes into the prior case when distinct but overlapping schemas do exist.

While maintaining multiple schemas is essential to human intelligence, it presents a challenge to artificial neural networks, which are especially prone to catastrophic forgetting or catastrophic interference (CI) [13, 19, 18]. This means that when trained on multiple tasks sequentially, these models forget how to perform previously learnt tasks. CI occurs primarily because of re-appropriation of networks that are important for previously learnt task for a new task, causing disturbances in weights or complete lost of representations.

Several approaches exist to tackle this problem. One approach is interleaved-learning, which ensures that tasks are trained in a mixed fashion. In this case, forgetting can be prevented due to the joint optimization of network weights for all tasks. Successful results include training a single agent able to play multiple Atari games [23, 20]. To make sure that past knowledge interleaves with new tasks, an episodic memory system is needed for replaying previous episodes to the network [14, 18]. However, this approach has been shown to slow down learning [17] and it is impractical to execute for a high amount of tasks with large-scale data. A second approach is blocked-learning or split-learning. This approach is often combined with additional mechanism for deciding the point of splitting, for example, splitting off weights when high prediction errors present [3]. This allows the model to completely master one task before transitioning to another. Essentially, models learn to allocate network resources—specific sets of weights—in a way that prevents competing representations from interfering with each other.

1.2 Blocked-learning vs. Interleaved-learning

A natural question arises as how the brain implements its solution for preventing catastrophic interference while acquiring multiple schemas. Studies about memory encoding revealed that the brain is selective about choosing neural ensembles for storing new memory. When a new memory is encoded, the brain selects a neural ensemble that is less active or excitable, reducing the likelihood of interference with previously stored memories. This process, known as "engram competition," helps ensure that distinct memories are allocated to separate neural populations [21].

Recent research suggests that the effectiveness of Blocked-learning over Interleaved-learning in preventing catastrophic interference extends beyond memory encoding to the acquisition of complex schemas. Beukers et al [3] investigated schema learning using narrative event prediction tasks, where participants learned the transition structures of two different narrative "chains" with shared states but different transition probabilities. They found that participants who learned the schemas in a blocked fashion (multiple instances from one schema before switching to the other) outperformed those who learned in an interleaved fashion (schemas alternating on each trial).

The authors proposed a Bayesian model that incorporates representational splitting when there are large prediction errors. In the blocked condition, the large error when switching between narratives triggers the creation of a separate schema, preventing the neural network from reusing the same set of weights. This mechanism aligns with the concept of engram competition in memory encoding, suggesting a common principle of allocating distinct neural resources to mitigate interference between schemas or memories. Furthermore, Beukers et al. [3] demonstrated that the timing of Blocked vs Interleaved training matters. Inserting a block of training on each schema early in an otherwise interleaved sequence led to better performance than inserting blocks late, as early blocks allow proper segmentation of the schemas before interference occurs. In other words, the effectiveness of blocked-learning is likely due to successfully prompting the initial formation of separate schemas, which may not have occurred in Interleaved-learning to begin with.

Building on these findings, we consider competing theories regarding the effectiveness of blocked and interleaved learning in environments where schemas are distinct yet complex. Our study aims to determine: 1) the extent to which each learning strategy supports the acquisition of multiple complex schemas without triggering catastrophic interference; 2) how the first-learned schema might interfere with the learning of a second, through either catastrophic interference or incomplete schema separation; 3) the ability of individuals to alternate decision-making effectively between two overlapping contexts.

To explore these theories, we employed a novel board game paradigm, providing a dynamic and ecologically valid setting to examine schema acquisition. By manipulating the game rules and the learning sequence, we assess how different learning strategies impact schema acquisition and interaction between schemas. The interactive nature of board games allows us to gather rich behavioral data, such as move-level decisions and reaction times, offering deeper insights into the cognitive processes involved in learning and applying schemas.

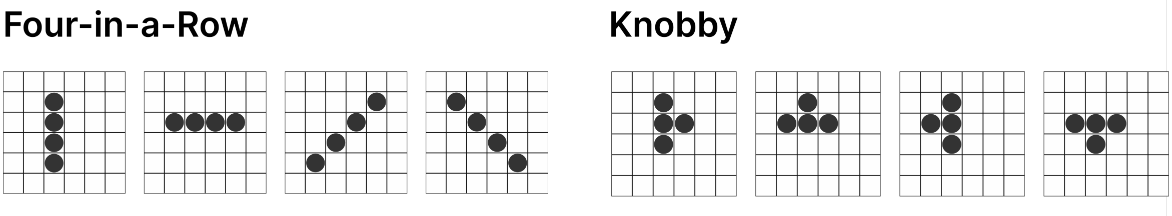

Figure 1: Winning conditions for Four-in-a-row and Knobby.

The two versions of the game in our study, "Four-in-a-row" and "Knobby", shared the same game board configuration and both involved getting four pieces into a particular shape in order to win. 1 But the difference in the winning conditions (creating a straight line versus a specific shape) created a large difference in the optimal strategies and moves in each game type, which we expected to lead to the formation of distinct schemas for each game. This assumption of schema dissimilarity is crucial for our hypotheses:

We hypothesized four possible outcomes for learning two distinct board games, Four-in-a-row and Knobby, under blocked and interleaved learning conditions. Each hypothesis about the outcome reflects a different interaction between the schemas:

H1: **Independent Learning**: Performance metrics (e.g., win rates) are similar for both games, suggesting that learning one game does not affect the learning of the other. This scenario would indicate complete separation of schemas.

H2: **Reduced Learning for the Second Game**: If learning the second game is less efficient than the first, this would imply incomplete separation and partial interference between the schemas.

H3: **Facilitated Learning for the Second Game**: An improvement in learning efficiency for the second game would indicate that the first game's schema facilitates the second. In other words, transfer learning happens between the games.

H4: Inability to Learn Either Game: This would suggest complete interference between the schemas, an outcome predicted by Beukers et al [3] but expected only under interleaved conditions due to increased competition for representational resources.

Given the similarities in the game boards and action spaces, we proposed that neither complete independence (H1, or Hypothesis 1) nor facilitated learning of the second game (H3) were likely. Because both the number of required pieces to win and the shape of the board are intentionally designed to be exactly the same, it is impossible to observe complete independence. Facilitated learning of the second game was considered unlikely given the distinct strategies required for each game, which can be quantitatively measured by the move evaluation by AI players (see Section 3.2). Our experiments aimed to test these predictions and assess the extent of interference or facilitation between the schemas.

In the following sections, we first describe the design of our board game paradigm and the experimental procedure. We then present our findings on how blocked and interleaved-learning affect game performance, move-level decisions, and reaction times. Finally, we discuss the implications of our results for theories of schema acquisition and catastrophic interference, and suggest directions for future research.

2 Methods

2.1 Experiment Design

Motivated by whether blocked or interleaved training curricula facilitate better schema learning, we developed a board game paradigm based on previous studies in the lab [10, 9]. On a 6×6 board, participants play two different games against distinct AI opponents

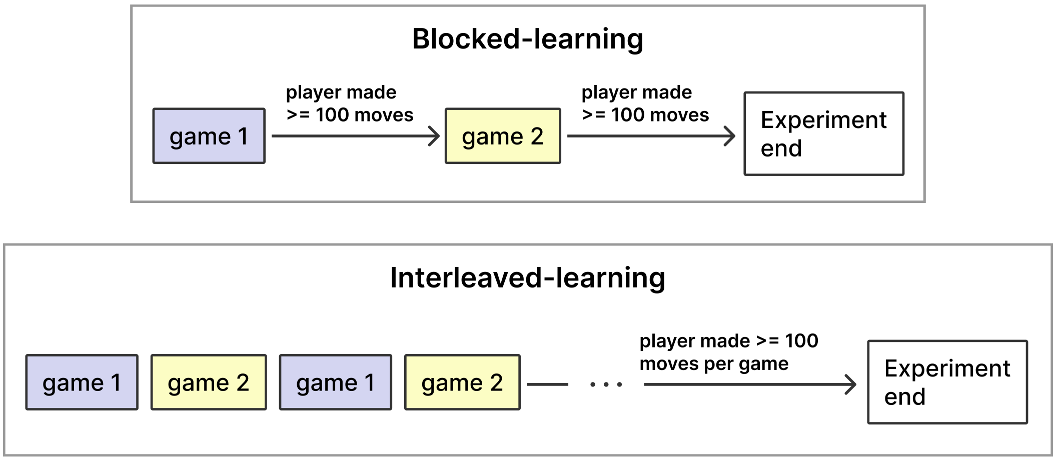

Figure 2: Two learning paradigms.

trained with neural networks based on Q-learning, a reinforcement learning paradigm.

The first game, **Four-in-a-row**, is a more advanced version of tic-tac-toe that asks players to construct four pieces in a consecutive sequence horizontally, vertically, or diagonally in order to win. In comparison, the winning condition of the second game **Knobby** is aligning four pieces so as to construe a T-shape with the stem directed toward left, right, up or down for winning. The two games are designed to ensure that the format of the game for learning schemas, that is the size of the board and numbers of pieces necessary for winning, remain the same, while the distinct winning condition leads to different optimal move strategies.

The study employs a between-subjects design. Preceding the experiment, participants were randomly allocated to two groups, the block-learning group and the interleaved-learning group, in a blinded manner. Within the block-learning group, participants engage in the game under one rule for 100 moves before transitioning to another rule for an equivalent number of moves. In contrast, participants in the interleaved-learning group alternate between rule types after every game they complete, i.e. the winning rule changes with each game. Before participants start the first game, they are instructed clearly on the winning condition in each game. The initiation order of each rule within both groups is equally probable, and the current game rule is consistently communicated explicitly to participants. Either the participant or the AI was randomly chosen to make to first move in each game.

The experimental procedure was automated through a Python script, and all participants conducted the experiment on a Windows computer. Participants were instructed to self-pace, and asked to take breaks only between games, and no more than twice.

2.2 Neural Network Model Configuration and Training

In our experiment, we used an artificial neural network model called *AlphaZero Gomoku* [26], which was originally designed for the game of *Gomoku* or five-in-a-row. We adapt it for our games Four-in-a-row and Knobby and train two separate models for playing against human participants during the experiment as well as evaluating participants' move quality later in the analysis stage. *AlphaZero Gomoku* is a specific implementation of *AlphaZero* [25] and uses Q-learning, a model-free reinforcement learning algorithm that seeks to learn a policy telling an agent what action to take under what circumstances. It does not require a model of the environment or a database of example games, instead learning through self-play which strategies are effective for winning the game. This approach allows the model be trained on its own data for developing a strategy from scratch, learning from each game played without prior knowledge of the game dynamics.

In the section below, we documented the technical design of the model and the experimental as well as analytical pipeline of the study. The codes and data collected can be found on GitHub Repository at `https://github.com/zhannahz/AlphaZero_Gomoku`.

2.3 Experiment Pipeline

2.3.1 Game logic handling

The experiment procedure including inputting participant information and experiment conditions is mostly implemented in `experiment.py`, with some steps such as getting the score of an action from a MCTS player for evaluating player move handled in `human_play.py`.

The game environment is represented by the `Board` class, which is initialized with a width w, height h, and the number of pieces in a row required for a win `n_in_row`. While it is as simple as setting both w and h as 6, and `n_in_row` as 4 to get a Four-in-a-row model, we wrote a separate game logic for training the Knobby model. The board state is stored as a dictionary, with keys representing moves as locations on the board and values representing the player who made the move. The game logic, including move validation, win condition checks, and game flow management, is handled by `game.py` and `human_play.py`.

2.3.2 Neural Network Architecture

Implemented in the `PolicyValueNet` class, the policy-value network used in this study (`policy_value_net_pytorch.py`) follows the architecture described in the AlphaGo Zero paper [25]. It takes the current board position s as input, represented by a $4 \times 6 \times 6$ tensor encoding the board position using 4 binary feature planes, and outputs a policy vector $\mathbf{p}$, representing the probability of each action, and a scalar value v, estimating the expected outcome of the game from the current player's perspective.

In each self-play game, the current best model plays against itself, starting from an empty 6×6 board. At each move, the model uses Monte Carlo Tree Search (MCTS) to select actions, with 400 simulations per move. The search is guided by the policy and value outputs of the neural network. To ensure diversity in the self-play data, Dirichlet noise with $\alpha = 0.3$ is added to the prior probabilities of the root node, and actions are sampled based on the visit counts of each branch of the root node. The self-play game ends when a winner is determined or the board is full. The game data, consisting of the board positions s, MCTS action probabilities π, and game winner $z \in 1(humanplayer), 2(AIplayer), 3(tie)$, is saved from the perspective of the current player at each move.

The MCTS algorithm is implemented in the `MCTS` class in `mcts_alphaZero.py`. It uses the policy-value network to guide the tree search and evaluate leaf nodes. Each node in the MCTS tree (`TreeNode`) maintains its own value Q, prior probability P, and visit-count-adjusted prior score u. The node expansion is guided by the action priors, which are tuples of actions and their prior probabilities according to the policy function.

The search process is described as below:

1. **Expansion**: Each tree node can be expanded based on the potential actions and their associated prior probabilities. This expansion is performed once an action is selected that has not been explored yet, leading to the creation of a new child node for each unexplored action.
2. **Selection**: During the tree search, actions are selected among the children nodes based on a combination of the action value Q and a bonus $u(P)$, which incorporates the exploration-exploitation trade-off controlled by a parameter c_{puct}.
3. **Update**: After simulating a game through the leaf node, the node values are updated based on the simulation results. This update is recursively applied to all ancestors of the leaf node to ensure that all relevant paths are updated with the new information.

2.3.3 Training Pipeline

The training process in `train.py` follows the self-play reinforcement learning paradigm. In each iteration, the current best model plays against itself to generate training data. The generated data, consisting of board states, MCTS move probabilities, and game outcomes, is then used to train the neural network. The trained model is evaluated against the previous best model, and if its win rate is better than previous win rate, it becomes the new best model. Both the Four-in-a-row model and the Knobby model achieved a win rate > 80 percent against itself. It took 5 to 10 hours to train an AI model for our experiment.

Prior to training a model, we set up relevant hyperparameters as below:

- c_{puct}: The exploration coefficient which determines the weight given to the prior probability P in the bonus term. Due to a high branching factor in the board games, this value is set to 5 for a higher tendency to explore.
- temp (τ): The temperature parameter for the softmax function applied to the visit counts during move selection, set to 0.75 for training and 10^{-3} for playing against participants.
- n_{playout}: The number of MCTS playouts per move, set to 400.
- learn rate (α): The learning rate for the neural network optimizer, set to 2×10^{-3}.

2.4 Participants

On the Columbia SONA platform, we recruited 32 students seeking research participation credit as part of their introductory psychology classes. All participants were over 18 years of age and gave informed consent for the study. The experimental protocol was approved by the Institutional Review Board of Columbia University (AAAS0252).

Following the exclusion of data from six participants due to technical issues, the revised participant count for each experimental condition was as follows: in the blocked condition, N=13 (7 participants started with Four-in-a-row, and 6 with Knobby); in the interleaved condition, N=13 (6 participants start with Four-in-a-row and 7 with Knobby). The age of the remaining 26 participants ranged from 19 to 37 years, with a mean age of 22.4 years. The average educational attainment was 13.5 years. Gender distribution was balanced, including 11 females, 11 males, and 3 non-binary individuals. Ethnically, the cohort included 13 White, 6 Asian, 2 Hispanic or Latino, 2 Black or African American individuals, and 2 participants of mixed ethnicity, specifically White and Black or African American.

2.5 Limitations in Methods

This study, constrained by the scope of a senior project, did not include monetary incentives or attention checks. The absence of incentives may have influenced participant motivation and engagement levels, potentially affecting the learning outcomes. Similarly, without attention checks, it's challenging to gauge continuous engagement, which is crucial for tasks requiring sustained attention. These limitations highlight areas for future research to enhance experimental design and data integrity.

3 Results and Findings

3.1 Game-level analysis

Overall, participants are adequately challenged in both the game of Four-in-a-row and Knobby, with a slightly higher average win rate in Knobby (0.421) 3, aligning with the the-

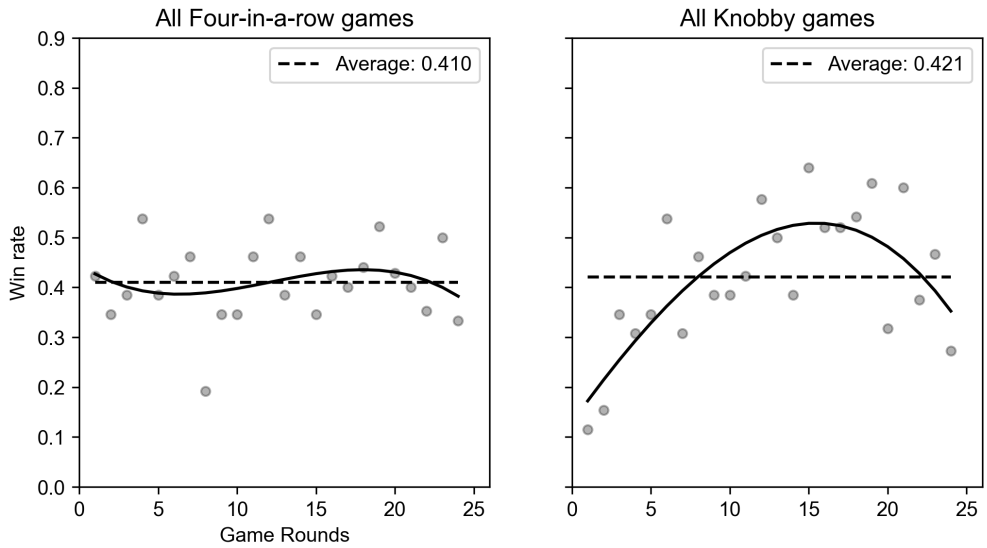

Figure 3: Win rates of Four-in-a-row versus Knobby. Each dot represents the average win rates across participants at a specific round.

oretically lower difficulty of the latter [7]. However, as a novel game that no participants should have played before, Knobby's win rates exhibits as a rising learning curve. Four-in-a-row, on the other hand, maintains a static win rates over the course of the experiment. One explanation for this surprising observation is that the total length of the experiment does not offer enough time for participants to improve significantly on Four-in-a-row. Instead, players remain an certain level of performance that they inherits from the game of Tic-tac-toe (three pieces in a row).

The order of play, specifically who goes first, significantly influences the chances of winning in both games. Participants achieve an average of approximately 70.9% wins when they make the first move in a game, 27.4% higher than the average win rates in the games where the AI player goes first. The result shows that both Four-in-a-row and Knobby are games sensitive to play order due to the possibility to force a win.

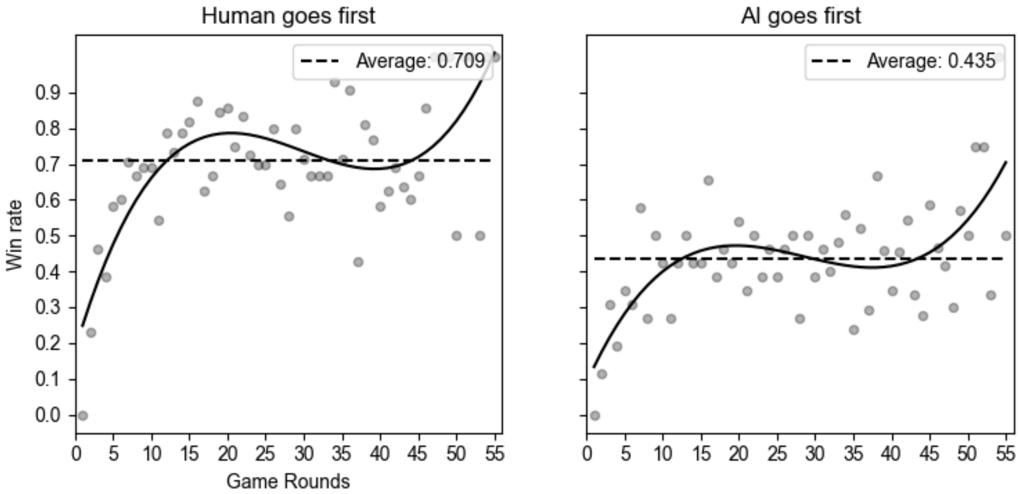

Figure 4: Win rates of games with players going first versus AI going first.

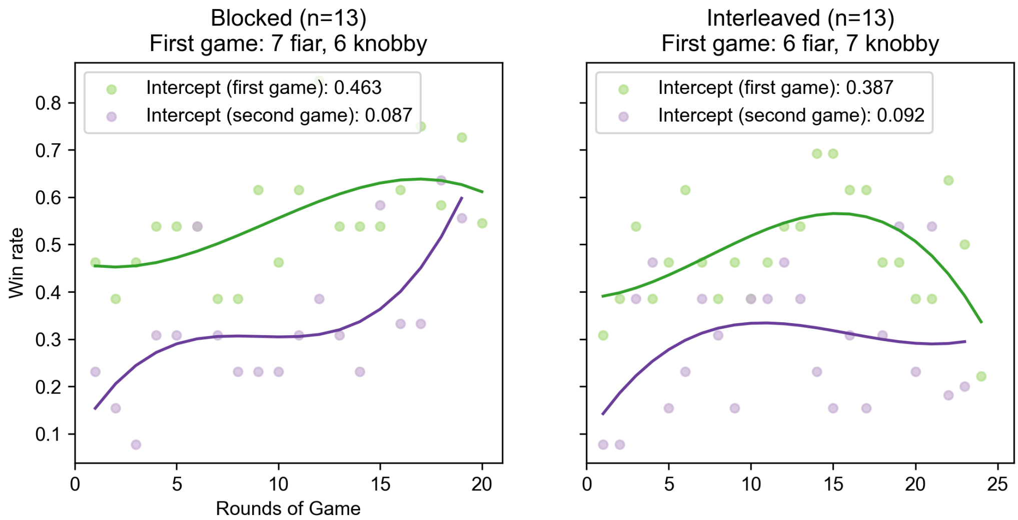

Figure 5: Win rates. Win rates are calculated as the ratio of rounds won to total rounds played by all participants within both learning paradigms—block and interleaved. Note that win rates are computed distinctly for the first game and the second game, which are defined by the order by which the participant encounter them.

(1) The first schema learned gains significant initial advantage

The win rate comparisons between the first and second game types are presented in 5, which illustrates the win rates over rounds of gameplay for the Blocked and Interleaved learning conditions. In the Blocked condition, the initial win rate for the first game type is significantly higher, with an intercept of 0.463, compared to the second game type, which has an intercept of only 0.087. This pattern indicates a strong "primacy effect," where learning and mastering the first schema offers a distinct initial advantage, which can be crucial in early performance.

In the Blocked condition, win rates of the first game type increase over time, suggesting that participants are consolidating their learning and applying it effectively as they progress. However, the second game type does not show a similar level of improvement, maintaining a relatively flat trajectory until the late stage of the experiment, suggesting a hindrance by the first game type in the learning of the second game type.

In the Interleaved condition, the primacy effect persists, albeit with slightly lower intercepts: 0.387 for the first game type and 0.092 for the second. The win rates in the Interleaved condition peak at around the midpoint of the rounds and then slightly decline. This could indicate an intensification of schema interference or a lower ceiling of performance in the interleaved learning paradigm.

(2) Blocked-learning results in better game performance at the end

The Blocked condition's learning trajectory is consistent with conventional models, where cumulative experience correlates with progressive mastery of the schema. The higher initial win rate, coupled with a positive learning slope, suggests effective strategy consolidation leading to improved performance over successive rounds.

Conversely, the Interleaved condition exhibits an initial increase in win rate followed by a subsequent decline in both games, suggesting an earlier attainment of the ceiling. Notably, the win rate for the second game type in the Interleaved condition demonstrates a slight improvement initially, but this is followed by a plateau or even a slight decline as the rounds progress. This pattern indicates that consistent performance improvement for the second game type is severely hindered, in addition to the eventual decline in performance of the first game type, suggesting that interference is more severe when two games are played in an interleaved fashion.

(3) Overall faster gameplay over time despite diverse gameplay speed

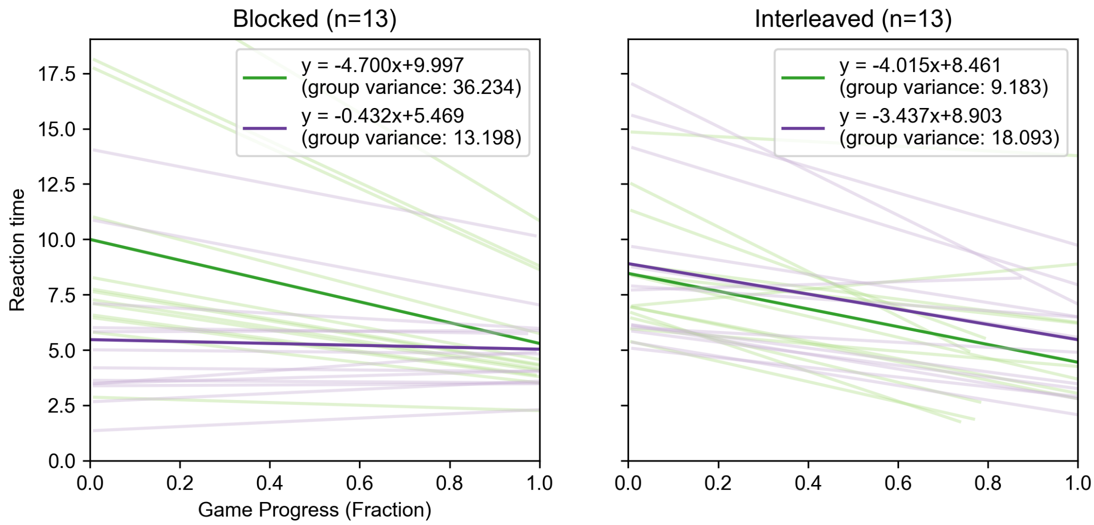

Figure 6: Move-level reaction time. Reaction times are calculated as the time between the onset of a participant's turn and the execution of their move are recorded to access the learning process. A Mixed linear effect regression is conducted across group (n=13 for Blocked and for Interleaved respectively).

While win rates reflect the outcome of the game, reaction times capture the speed of decision-making processes and may suggest the extent of patience. Both Blocked and Interleaved-learning groups showed improvement in gameplay speed over time, as seen in 6. The Blocked condition (left) shows a significant decrease in reaction time in the first game (coef. = -4.700, P < 0.001, group variance = 36.234) and a non-significant shallower slope in the second game (coef. = -0.432, P = 0.590, group variance = 13.198). The relatively static trend of the second game type in blocked condition shows that participants did not slow down when they encounter the second game type, suggesting that the decrease in reaction time is due to the experimental progress in general. Similarly, the Interleaved condition (right) exhibits significant decreases in both the game type (coef. = -4.015, P < 0.001, group variance = 9.183) and second game type (coef. = -3.437, P < 0.001, group variance = 18.093). One explanation is that participants are going faster regardless of game order and regardless of conditions.

3.2 Move-level analysis

To find out how well participants are learning the game, we need to access move-level performance. The quality of a move can be accessed by the extent to which the move adheres to the optimal strategy of playing the game given a particular circumstance.

Each move's quality was assessed by comparing the estimated probability scores generated by two distinct neural network models. These models estimate the probabilities of possible moves by simulating the game and recording the number of visits to each available board position. Because it is possible for a move to have 0 probability if a board position was never samples, we smoothed all probability estimates by taking a weighted average of the model probabilities with a uniform distribution over unoccupied squares (with a weight of $\alpha = 0.01$ for the uniform distribution and $1 - \alpha = 0.99$ for the model distribution). Then, we take the difference between the model probabilities of the first and second game type. Finally, we applied a mixed linear effect model to analyze this difference. The model was set up as a function of game type and game progress, in addition to accounting for the random effects of game progress within each participant. Game progress was defined as the fraction of the current round number divided by the total number of rounds in the given game type. The mixed linear effect model was structured as follows:

$$\text{smoothed\_model}_{\text{m1}} = (1-\alpha)\cdot p_{\text{raw\_m1}} + \alpha \cdot \left(\frac{1}{\text{\# empty squares}}\right) \tag{1}$$

$$\text{smoothed\_model}_{\text{m2}} = (1-\alpha)\cdot p_{\text{raw\_m2}} + \alpha \cdot \left(\frac{1}{\text{\# empty squares}}\right) \tag{2}$$

$$\text{model}_{\text{diff}} = \text{smoothed\_model}_{\text{m1}} - \text{smoothed\_model}_{\text{m2}} \tag{3}$$

$$\text{model}_{\text{mle}} = \text{model}_{\text{diff}} \sim \text{game\_type} \times \text{game\_progress} + (\text{game\_progress} \mid \text{participant}) \tag{4}$$

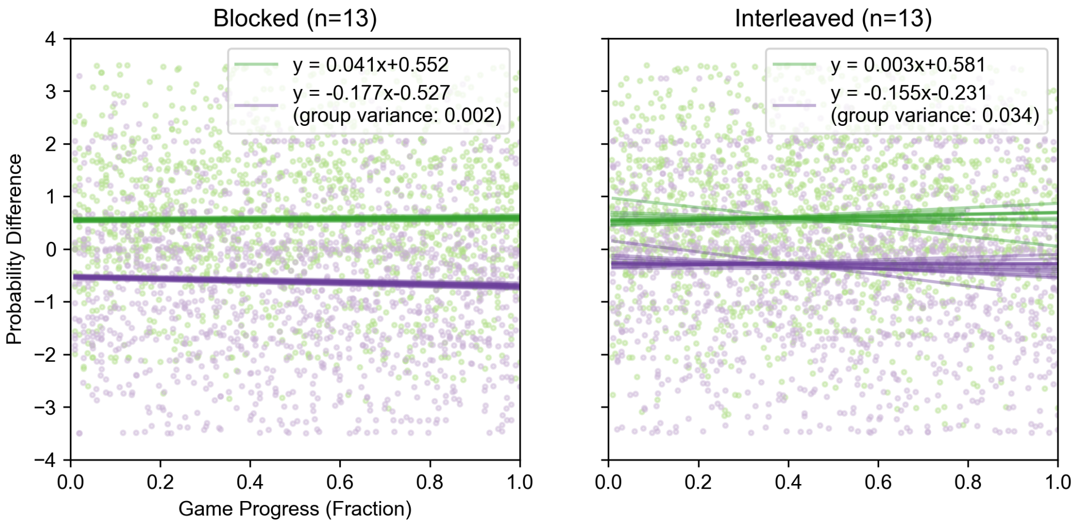

Figure 7: Probability difference. Probabilities associated with each possible move preceding a participant's action is extracted from the two distinct AI models, each trained on the separate games. A mixed linear effect regression is conducted across group (n=13 for Blocked and for Interleaved respectively) to reveal the trend of probability difference across game progression.

To assess the effectiveness of each learning condition, we computed a mixed-effects model to analyze the probability differences between the moves in each game type. This model included fixed effects for game type (first game versus second game), game progress (fraction of games played so far), and their interaction, and random effects for subjects to account for individual differences in learning and performance.

(4) Blocked-learning better help establish and stabilize distinct schemas

Both conditions showed a significant difference between the first and second game, with a more pronounced effect observed in the Blocked condition (coef. = -1.079, $P < 0.001$) than in the Interleaved condition (coef. = -0.813, $P < 0.001$). This indicates that, overall, participants were more (correctly) playing moves that were more aligned with the first-game model when playing their first game and more aligned with the second-game model when playing their second game. The interaction between game type and progression fraction was not significant in either condition, but was numerically larger in the Blocked (coef. = -0.218, $P = 0.253$) versus the Interleaved condition (coef. = -0.158, $P = 0.400$), suggesting that there was more learning occurring in the Blocked condition (since participants were becoming more aligned with the correct game models over the course of the experiment).

4 Discussion

How are overlapping schemas learned consecutively? We proposed a novel board game paradigm to investigate this systematically, focusing on whether blocked- or interleaved-learning facilitates better schema acquisition. The key finding of this study that blocked learning facilitates better schema separation and performance aligns with previous work, while extending it to a more naturalistic, interactive setting.

One contribution of this work is the use of a novel board game paradigm. One specific advantage of using a 6×6 board as the base upon which game rules are design for this study is that it allows us to customize the games for this experiment, ensuring they overlap in appearance and actions but differ in optimal strategies. However, our motivation of using games as experimental paradigm extends beyond this study. Traditionally, memory research has oscillated between highly-controlled discrete item tests, which allow for precise mathematical modeling but lack ecological validity, and narrative-based approaches that incorporate real-life complexities but struggle with reproducibility and formal modeling constraints [11, 12, 6]. While these methods have advanced our understanding of cognitive processes, they often fail to capture the dynamic interplay of memory encoding and retrieval that characterizes everyday cognitive activities [8, 15]. We believe that games offer an ideal intersection between overly simplistic domains and uncontrolled naturalistic environments. With careful experimental design, games provide just the right amount of complexity [1] for people to learn abstractions, predict possible upcoming scenarios, and plan out actions, while being feasible to be captured computationally. In addition, gameplay as a task facilitates a more internally motivated learning process and affords participants to have active controls over event sequences, offering a unique opportunity to study cognitive processes in contexts that resemble everyday event dynamics.

The findings of our study align with Beukers et al. [3], who emphasize the benefits of Blocked-learning in separating representational resources early for different schemas.This is particularly evident in the significantly better performance of the first game that participants play, indicative of faster and more stable schema formation. Furthermore, we found that the initially learned schema gains a profound and persistent advantage, as participants in both condition performed better at the first game type they played, suggesting a primacy effect in consecutive schema learning.

However, it is important to note that our conclusions should only extend to consecutive learning of schemas that are more distinct than similar. In other words, Four-in-a-row and Knobby are games that demands different mental models for predicting and planning moves, despite that they both operate on the same board configuration. Therefore, the evident interference that occurs in interleaved-learning can be explained from the perspective of competitions between the representational resources of the two distinct schemas. On the other hand, two similar schemas are more likely to be captured by a single, more generalize internal model, leading to assimilation than competition of schema representations. Future studies are needed to understand the role of task similarity in block and interleave-learning paradigms, such as varying the degree of similarity between tasks to examine how this influences schema interference and transfer effects.

Future studies could also build on our paradigm and investigate aspects of multiple schema learning that were beyond the scope of this study. An unanswered question is whether schemas learned in the blocked learning condition are better stabilized and retained over the long term, which could be explored by adding a final examination session to test participants on both games after a period of time.

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my thesis advisor, Prof. Chris Baldassano, for his invaluable mentorship during my time as a research assistant in his lab and for his guidance throughout my thesis project. My gratitude goes to Prof. John Morrison for his steadfast support since I set out to create my own cognitive science major until earning this degree. I am also grateful to Prof. Seth Cluett and Prof. Mark Santolucito who, though not in my field, have in their own ways empowered me to fearlessly pursue what I value intellectually.

I express love to my friends whose belief in my pursuits constantly motivated me. I cherish my future partner's immense emotional support this demanding year, despite long-distance. I owe gratitude to my family for their genuine interest and unconditional backing, though cognitive science lies outside their realms.

Ultimately, the cognitive science discipline itself provided an intellectual "hometown" that no physical space has ever provided. The questions and ways of pursuing them in the field have changed me as a person.

May 5, 2024

New York

References

[1] Marc Malmdorf Andersen, Julian Kiverstein, Mark Miller, and Andreas Roepstorff. Play in predictive minds: A cognitive theory of play. *Psychological Review*, 130(2):462, 2023.

[2] Christopher Baldassano, Janice Chen, Asieh Zadbood, Jonathan W Pillow, Uri Hasson, and Kenneth A Norman. Discovering event structure in continuous narrative perception and memory. *Neuron*, 95(3):709–721, 2017.

[3] Andre O Beukers, Silvy HP Collin, Ross P Kempner, Nicholas T Franklin, Samuel J Gershman, and Kenneth A Norman. Blocked training facilitates learning of multiple schemas. *Communications Psychology*, 2(1):28, 2024.

[4] Andy Clark. Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, 36(3):181–204, 2013.

[5] Karl Friston. The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11(2):127–138, 2010.

[6] Orit Furman, Nimrod Dorfman, Uri Hasson, Lila Davachi, and Yadin Dudai. They saw a movie: long-term memory for an extended audiovisual narrative. *Learning & memory*, 14(6):457–467, 2007.

[7] Martin Gardner. *Colossal Book of Mathematics: Classic Puzzles Paradoxes And Problems*. WW Norton & Company, 2001.

[8] Andrew C Heusser, Paxton C Fitzpatrick, and Jeremy R Manning. Geometric models reveal behavioural and neural signatures of transforming experiences into memories. *Nature Human Behaviour*, 5(7):905–919, 2021.

[9] Jiawen Huang, Eleanor Furness, Yifang Liu, Morell-Jovan Kenmoe, Ronak Hasan Elias, Hannah Tongxin Zeng, and Chris Baldassano. Making accurate prediction facilitates robust episodic memory formation beyond learned stimulus probability. 2024. Work In Progress.

[10] Jiawen Huang, Isabel Velarde, Wei Ji Ma, and Christopher Baldassano. Schema-based predictive eye movements support sequential memory encoding. *Elife*, 12:e82599, 2023.

[11] Michael J Kahana and Jonathan F Miller. Memory, recall dynamics. *Encyclopedia of the mind. Thousand Oaks, CA SAGE Publications Inc*, 2013.

[12] Walter Kintsch. Models for free recall and recognition. *Models of human memory*, pages 331–373, 1970.

[13] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.

[14] 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.

[15] Hongmi Lee, Buddhika Bellana, and Janice Chen. What can narratives tell us about the neural bases of human memory? *Current Opinion in Behavioral Sciences*, 32:111–119, 2020.

[16] Jean M Mandler. On the psychological reality of story structure. *Discourse Processes*, 10(1):1–29, 1987.

[17] James L McClelland. Incorporating rapid neocortical learning of new schema-consistent information into complementary learning systems theory. *Journal of Experimental Psychology: General*, 142(4):1190, 2013.

[18] James L McClelland, Bruce L McNaughton, and Randall C O'Reilly. 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, 1995.

[19] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.

[20] Emilio Parisotto, Jimmy Lei Ba, and Ruslan Salakhutdinov. Actor-mimic: Deep multitask and transfer reinforcement learning. *arXiv preprint arXiv:1511.06342*, 2015.

[21] Asim J Rashid, Chen Yan, Valentina Mercaldo, Hwa-Lin Hsiang, Sungmo Park, Christina J Cole, Antonietta De Cristofaro, Julia Yu, Charu Ramakrishnan, Soo Yeun Lee, et al. Competition between engrams influences fear memory formation and recall. *Science*, 353(6297):383–387, 2016.

[22] David E Rumelhart. Notes on a schema for stories. In *Representation and understanding*, pages 211–236. Elsevier, 1975.

[23] Andrei A Rusu, Sergio Gomez Colmenarejo, Caglar Gulcehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. Policy distillation. *arXiv preprint arXiv:1511.06295*, 2015.

[24] Richard A Schmidt. A schema theory of discrete motor skill learning. *Psychological review*, 82(4):225, 1975.

[25] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*, 2017.

[26] Jun Xiao Song. Alphazero_gomoku: Implementation of alphazero algorithm for playing gomoku. `https://github.com/junxiaosong/AlphaZero_Gomoku`, 2023. Accessed: 2023-09-15.

[27] Sharon Alayne Widmayer. Schema theory: An introduction. *Retrieved December*, 26:2004, 2004.