The Need for Speed?: Exploring the Contribution of Motor Speed to Expertise in a Complex, Dynamic Task

Catherine Sibert (c.l.sibert@rug.nl)
Faculty of Science and Engineering, University of Groningen
Broerstraat 9, Groningen, 9712 CP, Netherlands

Roussel Rahman (roussel.rahman@gmail.com)
Department of Photon Science, SLAC National Accelerator Laboratory
2575 Sand Hill Rd, Menlo Park, CA 94025, USA

Abstract

Experts, across tasks, tend to be faster and more accurate than novices. While the speed and accuracy of decision-making are critical elements of expert performance in dynamic tasks, their interaction over the course of learning remains less explored. Can we determine when poor performance is due to failures of execution – knowing the correct action but being unable to carry it out – or failures of insight – not identifying the optimal action? In this work, we explore how time pressure affects accuracy at various stages of learning in a complex, dynamic task: the game of Tetris. We emulate human decision-making processes under time pressure in several reinforcement learning models by training them under time pressures present on humans. Subsequently, we compare the performance and the behavior of human players against the ones demonstrated by AI players of equivalent skill. At the surface level, the AI models are able to achieve human-like performance levels at different stages of expertise. However, when probed at lower levels, we find that their behavior and strategies are considerably different from the ones employed by human experts. Examining why and how the models differ from humans highlights the promise of using AI models to study the nuances of human decision-making in dynamic tasks, along with the need to explain both human and AI behavior at multiple performance levels for accurate understanding.

Keywords: Cognitive Modeling, Expertise, Skill Acquisition

Introduction

The speed-accuracy trade-off is considered a foundational limit of human performance (Ratcliff & Tuerlinckx, 2002; Heitz, 2014). While this limit broadly holds true for simple perceptual-motor tasks, as tasks increase in complexity, the relationship between speed and accuracy also becomes more complex. In fact, in many cases, expertise in such tasks is identified by an individual's ability to be both fast and precise (e.g., Seamster, Redding, Cannon, Ryder, & Purcell, 1993; Lindstedt & Gray, 2019; Gray & Banerjee, 2021). The ability to perform accurately in complex and dynamic environments is believed to be the result of the development of skills and strategies that allow experts to be accurate within the time demands of the task (Gigerenzer & Goldstein, 1996; Gigerenzer & Selten, 2002; Gigerenzer, 2008).

However, this expertise develops over time in different ways, and understanding where in the skill development trajectory a particular individual is currently located will be vital to developing systems to train and support experts more efficiently. As expertise is the ability to perform quickly and accurately, one indicator of a lack of expertise is the presence of errors in the task. We consider errors to come in two forms: failures of **execution** and failures of **insight**. Failures of execution occur when an individual intends to take a particular action but is unable to do so, perhaps due to time or physical limitations. Conversely, failures of insight occur when an individual correctly carries out their chosen action but does not realize that the action is suboptimal in some way.

Understanding the types of errors a particular individual is making would go a long way towards understanding how their skill is developing, but in most cases, this can be very challenging to identify. Performance in many complex tasks is measured by the final outcome, meaning that the contribution of any individual action within the task is hard to isolate. Additionally, many tasks lack a concrete ground truth for what the "right" action at each moment should be, further complicating efforts to identify errors.

While human performance can be messy and difficult to quantify, cognitive models provide a potential tool for making analyses more concrete. In contrast to AI models that seek optimal solutions for the given task, cognitive models attempt to emulate human cognitive processes and capture the mechanisms of human decision-making, providing vital insights into the methods and strategies that humans use when making decisions in complex environments. Moreover, cognitive models provide a way to investigate AI behavior and strategies under different environmental constraints. In this work, we emulate human decision-making in cognitive models to play the game of Tetris, a complex, dynamic game that requires both speed and accuracy for expert performance.

Tetris as an Expertise Domain

The video game Tetris has been used as a research task in a variety of domains and is of particular interest as a way of exploring skill acquisition and expertise (Gray, 2017). In Tetris, players make decisions about where to place a series of geometric shapes that fall from the top of the screen in a pile that accumulates at the bottom. Whenever the pieces are placed such that they form one or more complete, unbroken rows, these rows disappear, making the pile smaller. The game becomes progressively more difficult by increasing the speed at which the pieces fall and ends when the pile reaches the top of the screen. To perform well, players must make a series of fast and accurate decisions about the best position for each of several hundred pieces in a constantly changing environment. Across skill levels, players employ a range of strategies and

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

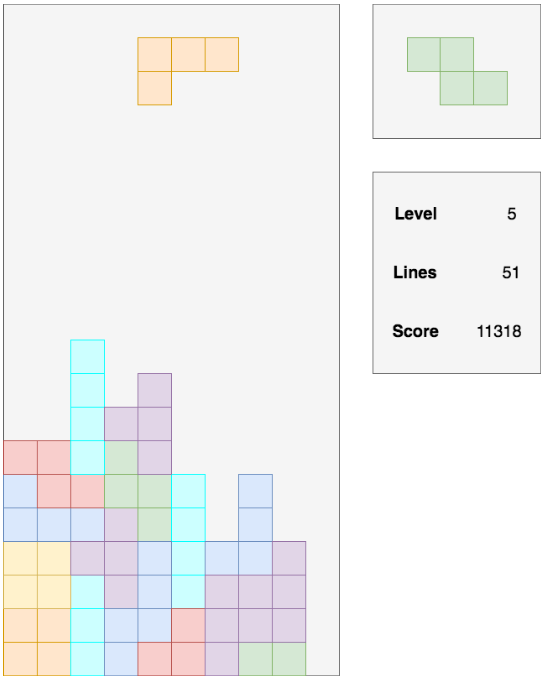

Figure 1: A typical episode in Tetris

behaviors in an effort to accumulate the highest score possible before the pieces fall too quickly to position well.

Figure 1 shows a typical gamestate in Tetris. Planning or strategizing is an important component of Tetris expertise. The players gain points at a higher rate when multiple lines are cleared together, with the maximum corresponding to four-line clears with a single piece. The points awarded for each type of line clear increase with the game level. While higher levels provide more points per line clear, the speed of gameplay also increases (the zoid takes 16 seconds to fall from top to bottom at level 1, 2 seconds at level 9, and only 1/3 of a second at level 29). As such, players must balance the desire to set up for high scoring moves with the speed with which the piece can be moved into place. The main information to use in planning consists of (a) the current board state, (b) the current zoid, and (c) the next zoid. There are seven different zoid types that enable a large search space of zoid placements. For example, in Figure 1, the incoming 'L'-zoid can be placed in the right-hand well to earn two line clears, but an expert player is likely to wait for an 'I'-zoid to clear four lines together.

Models of Tetris

There is a long history of testing machine learning models with Tetris due to its computational complexity (Robertson, 2012; Boumaza, 2009; Carr, 2005). Despite the complexity of the observed human behavior, relatively simple machine learning models have been able to achieve high levels of performance. At first glance, these models perform so far beyond the level of human experts that comparisons seem futile. However, upon closer inspection, these models turn out to have a secret advantage: where human players must make and execute their chosen placements under increasing time pressure, the models very often remove any kind of time pressure. This makes sense when the models are viewed as proof that a particular machine learning algorithm is effective at finding an optimal set of weights within a large feature space, but by removing the factor of time, the models are tackling a distinct task from human players, and as a result, solve it differently. Having effectively unlimited time fundamentally changes the behavior of the models, who rack up high scores not by planning and executing high-scoring moves (as human experts do) but instead by continually making safe, low-scoring moves and extending the game far past the point that any human could reasonably reach when constrained by time pressure.

While accurately incorporating all elements of the time costs faced by humans would be challenging, a fairly simple measure - limiting the game length to a reasonable human scale (525 pieces) - was able to produce models that much more closely resemble human performance and behavior (Sibert, Gray, & Lindstedt, 2015). These more human-like models have provided a helpful point of comparison for human players and have been a useful tool for exploring the fine-grained details of human performance in Tetris (Sibert & Gray, 2018; Sibert, 2019).

To be truly valuable as a means of understanding and explaining human behavior as it develops across skill levels, it is also important to have a clear picture of the full space of possible human performance. In this work, we combine the time pressure of finite game length on models with time pressure at the execution level – that is, executing the moves required for desired placements in each episode within the available time. The time pressure parameters and their degrees were determined from the analysis of a large corpus of human gameplay data. After establishing in our first analysis that the models emulate human-like performance at the surface level, we probe into their behavioral strategies that led to the scores and the effects of time pressure on the strategies.

General Methods

Defining Human Expertise in Tetris

Population Study Between 2015 and 2018, a long-term study was carried out at Rensselaer Polytechnic Institute, where subjects were recruited to play an hour of Tetris in a controlled environment. All subjects were given the same set of game seeds (which determines the sequence of pieces players must place) in random order and played as many games as they could complete within the time limit. The goal of the study was to build up a large corpus of behavioral data that could be used for analyses in the future. In order to collect data from the expert end of the skill spectrum, additional participants were recruited from Tetris tournaments, both local and international.

Criterion Score Performance in Tetris is generally measured by the final game score; the ultimate result of all of the hundreds of individual placement decisions. However, the

Table 1: High and Low score cut-off points for membership in each expertise group, along with their sizes in our dataset.

Skill Level	Low Score	High Score	Group Size
Extreme Expert	260001	999999	6
Expert	95001	260000	11
Advanced	45001	95000	31
Intermediate	15501	45000	168
Proficient	5801	15500	176
Beginner	1501	5800	100
Novice	601	1500	19
Extreme Novice	0	600	4

random nature of the piece sequence means that a single game is not always reflective of a player's actual skill level. To reduce the effects of a single bad game, players in the population study were evaluated by their **Criterion Score**: the average score of the best four games completed during the hour time limit. The criterion score of players recruited at tournaments was approximated by the highest score between two standardized qualifier games (with consistent game seeds) or another standardized test game.

Expertise Classification As the Population Study progressed, it was determined that having skill labels for groups of players would be useful, rather than just the numerical criterion scores. A hierarchical clustering analysis was performed on the criterion scores of the subjects, grouping similar scores together. Based on the clustering, the subjects were divided into eight skill bins, with the majority falling into the Intermediate or Proficient level and the fewest in the Extreme Novice and Extreme Expert groups.

Rather than repeating the clustering each time new data was added to the set, the initial clustering was used to set thresholds for criterion scores, and additional subjects were sorted based on these. A summary of the expertise groups is provided in Table 1. According to this classification, the models described in the previous section are Experts.

The Role of Time Pressure

Time pressure plays a significant role in shaping the performance and behavior of human players, but most machine learning models do not incorporate any kind of time-based restrictions, limiting their ability to predict or explain human decisions. A rough approximation of time pressure, in the form of a hard limit of pieces available per game, was able to provide insights into how humans develop strategies in response to environmental constraints by prioritizing higher-scoring but riskier placements over safer, lower-scoring placements (Sibert et al., 2015). While the game length limit provides an excellent proof of concept that the machine models are capable of capturing elements of human behavior, it does not provide a very nuanced implementation of how time pressure impacts individual decisions, nor how time pressure is handled at different levels of skill.

Table 2: Time parameters for each skill bin. The Initial and the Average latency are measured in milliseconds, and Efficiency is measured in key presses.

Skill Level	Initial Latency (ms)	Average Latency (ms)	Path Efficiency (keypresses)
Extreme Expert	79.28	162.11	0.50
Expert	70.56	173.79	0.57
Advanced	96.12	221.72	0.98
Proficient	272.25	433.31	1.64
Intermediate	151.24	311.02	1.16
Beginner	411.29	536.51	1.88
Novice	448.76	668.98	2.06
Extreme Novice	571.37	804.4	2.62

Time Parameters To incorporate time pressure into the models at the placement level, three parameters were extracted from the human behavioral data that determine the time taken to decide on the optimal position for the current zoid and execute the sequence of button presses required to move it into place.

- **Initial Latency** - Time between the beginning of the episode (i.e., when the active piece appears at the top of the board) and the first button press. This is likely analogous to decision time.
- **Average Latency** - Average time between all subsequent button presses, representing the average button pressing speed of a player.
- **Path Difference** - Difference between the number of button presses needed to execute the optimal path to the player's chosen move and the number of presses actually made by the player.

Time Parameters Across Expertise Groups For each of the eight expertise groups, group time parameters were determined by taking the average value for each parameter for every decision made by all players in the group. These time parameters are presented in Table 2.

Incorporating Time Pressure into Models

The base Tetris model makes placement decisions by generating all possible positions for the current zoid and board state and calculating a score for each based on the associated feature weights. The chosen move is the position with the highest score. The time parameters were incorporated by adding another stage to this process. After the possible moves are ranked by score, the model uses the time parameters to determine if its desired move is possible within the time available using the following equations:

$$Time = InitLat + AvgLat * NumPress \tag{1}$$

$$NumPress = OptimalPathPress + ExtraPress \tag{2}$$

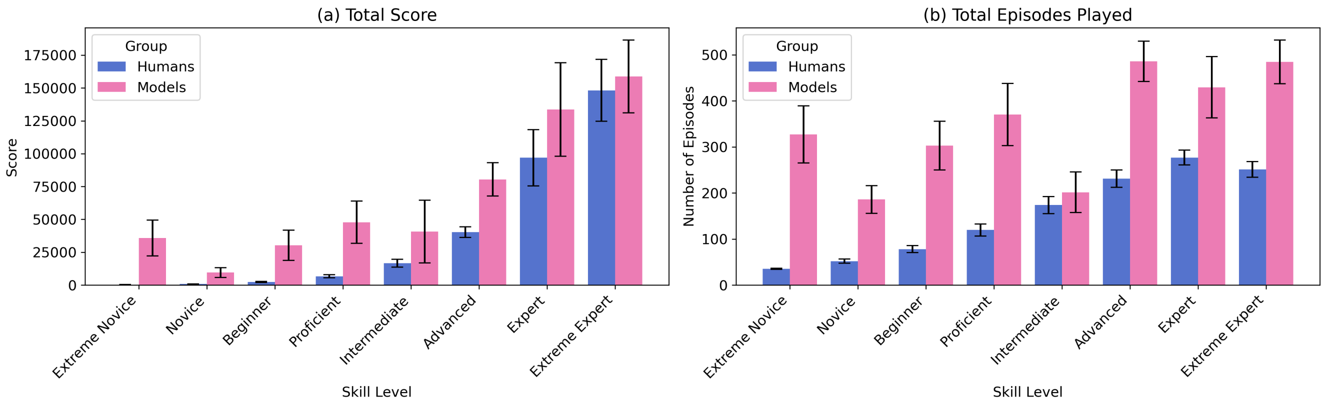

Figure 2: Overall performance by human players at different skill levels and by AI players emulating the human experts. The error bars represent 95% confidence intervals. The total scores (Fig. a) increase similarly with expertise across the two sets of players, albeit at slightly different rates. The episodes per game (Fig. b) also show consistent increases in episodes played with expertise. These similarities indicate that the models trained under the time pressure of human experts emulate the overall performance quite well. A notable difference is that the human players play fewer episodes than the models do to achieve similar score levels.

Where Time is the time cost in ms) for a placement, InitLat is the time between when a piece first appears on the board and the first keypress, and is roughly analogous to decision time. AvgLat is the average time needed to make any individual keypress, and NumPress is the number of keypresses made in order to move the piece into the desired place. OptimalPathPress is the minimum number of presses, and ExtraPress represents how many extra key presses, on average, the player is likely to make. An ExtraPress value of 0 represents perfect efficiency (with NumPress = OptimalPathPress) where a player always navigates along the optimal path. Positive ExtraPress values represent taking longer paths or adding extra rotations and translations.

The time cost, Time, is then compared to the drop time at the current game level – that is, the time (in ms) it would take a zoid to naturally fall from the top of the screen to the bottom with no intervention from the player. For this initial analysis, Time was compared to the maximum drop time for each level, neglecting the shorter drop distances as a result of zoids already on the board. If the time cost is less than the time available, the model places the zoid and moves on to the next decision point. Otherwise, the chosen move is eliminated as a possibility, and the time comparison is repeated for the next highest-ranked move. This process is repeated until a possible move is found. If no moves are determined to be executable in the time available, the game ends.

Training Time Pressure Models

Models were trained with the Dellacherie feature set (Fahey, 2015) using Cross-Entropy Reinforcement Learning, a technique for finding optimal feature values within a large search space over successive generations (Thiery & Scherrer, 2009a, 2009b; Szita & Lorincz, 2006). At each generation, 100 candidate models were generated within the search area, and each played a test game of Tetris. The feature weights of the ten best performers, in this case optimized for high score, were averaged together to form the center point of the search space for the next generation. The size of the search space is also adjusted at each generation to reflect the spread of values represented by the best models. Over time, the search space narrows in on the optimal feature weights.

In previous studies, human-like behavior was encouraged by setting a constant game limit. Here, this was replaced by the inclusion of time parameters described above, with games ending when the model is no longer able to execute moves. In order to reduce computation time, instead of using a predetermined learning time (80 generations), training was automatically halted when the features were considered to have converged (the variance in weights for the top ten models fell below 0.001).

Analysis 1: Performance of Time Pressure Models

To explore how the speed limitations of players at different skill levels impact their behavior and performance, a total of eight models were trained using the averaged time parameters from each of the eight expertise groups.

Methods

Each of the models was trained using the Cross-Entropy Reinforcement Learning method described above. The models with the final weights produced by training played 100 test games. Performance was measured by the average score of these test games, as well as the game length (total episodes or decision points).

Results

The performance of the models is compared to their related expertise group in Figure 2. As we see, the performance of models trained under different levels of time pressure matches reasonably well with that of their human counterparts. For both humans and models, the total game score increases with expertise (Fig. 2a). The game scores achieved at each skill level become increasingly similar at higher levels of expertise. The episodes played per game (Fig. 2b) represent the lengths of games the players needed to achieve the scores. We see that the length of games increases with expertise for both sets of players. However, a notable difference is that, at each skill level, the human players play a considerably lower number of episodes than the models to achieve similar scores. Additional performance metrics reflecting overall game performance replicated these patterns.

Discussion

Slowing down the models produces performance, measured by score, that is roughly analogous to the expertise group with similar speeds, suggesting that execution speed plays a significant role in the expertise level of a player. However, the other primary performance metric, game length, does not line up as well between the humans and models. This suggests that while the final outcome is similar, the methods used to achieve that outcome are not, and a closer look should be taken at the behavior of both the models and the humans.

Analysis 2: Behavior of Time Pressure Models

To better understand the actions that achieve the final game score, we investigated gameplay behavior in models and humans at game and decision levels.

Methods

Using the same set of 100 simulated games, the behavior of each of the eight time pressure models was categorized by examining the proportion of types of line clears made within a game. A high proportion of single-line clears is characteristic of a lower-skill, survival-first strategy, while higher-order line clears are associated with a higher-skill strategy that emphasizes planning to maximize rewards.

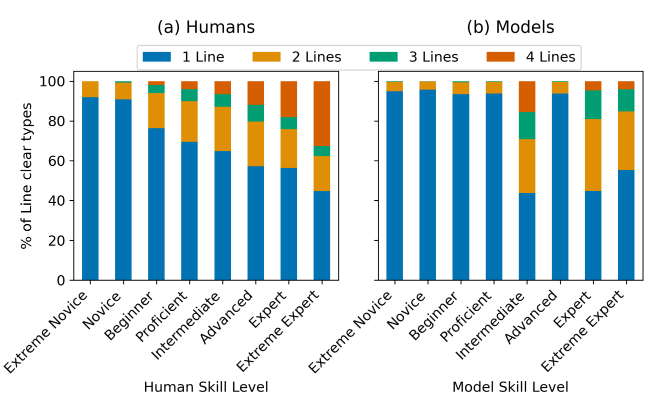

Figure 3: Types of line clears by (a) humans at different skill levels and (b) the models emulating different skill levels.

In addition, we investigated the agreement of moves chosen by humans and models. For all episodes played by humans, the models were presented with the board state and the active piece and asked to predict the optimal placement. The model choices were then compared to the move actually chosen by the player. The proportion of correctly predicted pieces over all decisions is defined as the Best Move Match Rate. [1]. In each episode, there may be multiple good moves that are similarly effective. Therefore, we also calculate the Good Move Match Rate, as the proportion of episodes where the human chose a move within an acceptable space of the best move for that episode.

Results

Figure 3 shows the line clear distributions for human players of different skill levels and their AI counterparts, who demonstrate distinct patterns of line clear types. In humans, as expertise increases, players gradually reduce their reliance on single-line clears and pursue higher-order line clears with greater frequencies. In models, while both behavior patterns are present, there appears to be a more abrupt shift around the intermediate level rather than a slow development.

The match rates for moves in episodes are presented in Table 3). Both match rates improve with expertise, with the good move match rate increasing beyond 90%. However, the match rates are much lower for the best move, indicating some disagreements between human players and models on what they consider to be the best move in each episode.

Discussion

While implementing time parameters into the model seems to provide a rough approximation of player skill, the behavior of the time-pressured models does not match well with the behavior of players in their analogous skill group. Where human behavior shows a gradual increase in the use of higher

[1] For models trained without time pressure, the Best Move Match Rate correlates with the player's skill level, ranging from about a 30 percent match for novice players, to a 65 percent match for experts (Sibert et al., 2015)

Table 3: Model match rates for humans across all skill bins. Values are taken by finding the average value for each player and then the average value across all players in the group. Standard deviations are reported in parentheses.

Skill Level	Best Move Match Rate	Good Move Match Rate
Extreme Expert	59.8% (1.5%)	94.5% (0.9%)
Expert	59.4% (2.3%)	95.1% (1.0%)
Advanced	47.1% (5.2%)	74.9% (4.5%)
Intermediate	54.6% (3.4%)	91.5% (1.9%)
Proficient	45.4% (7.0%)	72.7% (7.2%)
Beginner	40.8% (7.1%)	63.2% (7.8%)
Novice	32.7% (6.8%)	57.7% (8.9%)
Extreme Novice	17.1% (5.3%)	40.1% (9.4%)

order line clears as skill develops, the models show a clear split, with higher skill (i.e., faster) models relying on multiple line clears, but the slower models switching abruptly to single line clears. This suggests that the strategy of the models is not changing as they get faster, but when the model is too slow, it is no longer able to pursue the high order line clears, which is the best way to a high score. It appears that the Intermediate level is when this switch occurs.

The match rate results confirm that the time-pressured models are not accurately capturing what the equivalent humans consider to be a good move. If the time pressure is the only factor influencing behavior, the ability to execute a move in the available time should have a huge impact on how "good" a move is considered to be. The failure of slow models to predict the individual moves of slow players demonstrates that slowing the models down has not significantly changed which move the model considered to be the best.

In part, this may be due to how the time pressure was implemented: the models rank all available moves and then apply the time pressure to determine if the move is possible to execute. If there is not enough time, that move is eliminated from consideration, and another is chosen. In theory, if particular types of moves are frequently not possible, the model will adjust its features over time to rank executable moves more highly. However, most of the time, the time pressure does not impact the move that the model chooses. Particularly for fast models, their chosen move is always possible for the bulk of the game, with pressure only eliminating possibilities at very high levels. This likely means that the time pressure is not getting the opportunity to shape the model's weights, and its conception of what a good move is assumes no limitation on execution time. This does not fully explain the lack of ability to match slow players, where the time pressure is present for a larger proportion of decision points, suggesting the poor performance of low-skill players is not driven by speed alone.

General Discussion

The inclusion of time parameters into the model clearly demonstrates that speed plays an important role in expertise, but the failure of slow models to match the individual decisions of slow players shows that speed is not the only factor.

While we have performance data from players that perform better than the models, the models are capable of reasonably high scores and can generally be said to be making good decisions about where to place pieces, especially when used for comparison against lower-level players. Disagreements between the model and the player choices can, therefore, be considered failures either of execution or of insight. When allowed to play without time pressure, the models (by their own judgment) make no errors of any kind. The time-pressured models have perfect insight, but by using time parameters to determine the model's ability to move a piece to a desired place, they become capable of failures of execution.

In theory, training the models under different time constraints will shape the model's judgment about placement quality, and thus, compared to unconstrained models, could be said to have failures of insight. However, as discussed above, the time parameters do not seem to have enough influence on the weights of most of the models. The failure of these slow models to well match the individual decisions of players at low skill levels, therefore, suggests that these players are most often suffering from failures of insight and are selecting a suboptimal move because they are not yet aware of some higher-level strategy. This distinction is important because targeted training to improve performance would look different for each type of failure. Failures of execution in Tetris could be improved by speed training or improved route planning, while failures of insight would be better addressed by feedback about individual placement decisions or studying examples of "good" placements.

Of course, these models are still far from perfect. With the Time Parameter Models alone, it is not yet possible to distinguish which types of errors players are making. In particular, cognitively plausible models of the decision-making processes contributing to latency could delineate the changes in decision-making processes with expertise (Van Opheusden et al., 2023). A promising approach is to train models directly on human players, producing a set of feature weights that capture the insight level of each skill group. By creating models that isolate each aspect, it may be possible to then identify the degree to which insight or execution is shaping a particular player's behavior and identify the type of training that would be most likely to cause improvement.

Despite several mechanical advantages (e.g., the ability to perform a feature analysis on all 30+ possible placements without time cost), truly expert players still outperform the models, meaning that extreme skill is not explainable by the models. Still, the results demonstrate the power of even imperfect models in revealing the nuances of human behavior in a complex, dynamic task.

Conclusion

While speed is well understood to play a significant role in expertise across many tasks, simply slowing down models of Tetris only approximates the performance of different expertise groups at a surface level and does not produce models that make similar choices at the level of individual placement decisions. This dichotomy helps to reveal the distinction between failures of insight and failures of execution, the fine-grained identification of which could have implications for improving training in wider domains, where the ground truth quality of an individual choice within a complex, dynamic task is otherwise extremely difficult to determine.

Acknowledgments

Please address all correspondence to Catherine Sibert (c.l.sibert@rug.nl). This work was supported, in part, by grants N00014-13-1-0252 and N00014-17-1-2943 to Wayne D. Gray from the Office of Naval Research, Dr. Ray Perez, Project Officer. We wholeheartedly thank Wayne D. Gray for his invaluable knowledge, guidance, and support in this work. A special thanks to Ropa Denga for her valuable insights and discussions along the way.

References

Boumaza, A. (2009). On the evolution of artificial Tetris players *. Retrieved from https://hal.archives-ouvertes.fr/hal-00397045/file/boumaza-2009.pdf

Carr, D. (2005). Applying reinforcement learning to Tetris. Retrieved from http://www.cs.ru.ac.za/research/g02C0108/files/litreviewfinalhandin.pdf

Fahey, C. P. (2015, 09). *Tetris AI.* Retrieved from http://www.colinfahey.com/tetris/ ([Online; accessed 2015-Jan-30])

Gigerenzer, G. (2008). Why heuristics work. *Perspectives on psychological science*, *3*(1), 20–29.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological review*, *103*(4), 650.

Gigerenzer, G., & Selten, R. (2002). *Bounded rationality: The adaptive toolbox*. MIT press.

Gray, W. D. (2017). Game-XP: Action games as experimental paradigms for cognitive science. *Topics in Cognitive Science*, *9*(2), 289–307.

Gray, W. D., & Banerjee, S. (2021). Constructing expertise: Surmounting performance plateaus by tasks, by tools, and by techniques. *Topics in Cognitive Science*, *13*(4), 610–665.

Heitz, R. P. (2014). The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Frontiers in neuroscience*, *8*, 150.

Lindstedt, J. K., & Gray, W. D. (2019). Distinguishing experts from novices by the mind's hand and mind's eye. *Cognitive psychology*, *109*, 1–25.

Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic bulletin & review*, *9*(3), 438–481.

Robertson, J. (2012). *A brief history of tetris ai.*

Seamster, T. L., Redding, R. E., Cannon, J. R., Ryder, J. M., & Purcell, J. A. (1993). Cognitive task analysis of expertise in air traffic control. *The international journal of aviation psychology*, *3*(4), 257–283.

Sibert, C. (2019). *Unpuzzling tetris: exploring the mechanisms of expertise in a complex, dynamic task with simple machine learning models* (Publication No. 27663905) [Doctoral dissertation, Rensselaer Polytechnic Institute]. ProQuest Dissertations and Theses database.

Sibert, C., & Gray, W. D. (2018). The tortoise and the hare: Understanding the influence of sequence length and variability on decision-making in skilled performance. *Computational Brain & Behavior*, *1*, 215–227.

Sibert, C., Gray, W. D., & Lindstedt, J. K. (2015). Tetris-^{TM}: Exploring human performance via cross entropy reinforcement learning models. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 2188–2193).

Szita, I., & Lorincz, A. (2006). Learning Tetris using the noisy cross-entropy method. *Neural Computation*, *18*(12), 2936-2941.

Thiery, C., & Scherrer, B. (2009a). Building controllers for Tetris. *International Computer Games Association Journal*, *32*, 3–11. Retrieved from http://hal.archives-ouvertes.fr/inria-00418954/

Thiery, C., & Scherrer, B. (2009b). Improvements on learning Tetris with cross-entropy. *International Computer Games Association Journal*, *32*(1), 23-33.

Van Opheusden, B., Kuperwajs, I., Galbiati, G., Bnaya, Z., Li, Y., & Ma, W. J. (2023). Expertise increases planning depth in human gameplay. *Nature*, *618*(7967), 1000–1005.