The "cognitive speed-bump":
How world champion Tetris players trade milliseconds for seconds

John K. Lindstedt (john.lindstedt@oswego.edu)
Department of Computer Science
State University of New York at Oswego

Wayne D. Gray (gray@rpi.edu)
Department of Cognitive Science
Rensselaer Polytechnic Institute

Abstract

Tetris is a fast-paced puzzle solving game that requires players to rapidly maneuver falling blocks to clear rows and score points. Skilled Tetris players learn to execute moves in the game very quickly to keep up with the increasing time pressure. But world champion Tetris players employ more complex strategies that save precious milliseconds that enable them to reach even higher levels of play. Such strategies show mastery of the game's event structure, but also come with a startup cost— a "cognitive speed bump"— wherein they must momentarily decide whether to rotate a block left or right, even for scenarios where the distinction is not meaningful for performance. We present data showing both the world champions' superior overall action times, but also a preliminary "speed bump" that is consistent both within and between world champion players. Potential underlying memory structures are explored, and implications are discussed for both the Soft Constraints Hypothesis and the relationship between Hick's Law and expertise.

Keywords: cognitive strategies; extreme expertise; video games; real-time tasks; complex tasks; skill acquisition; Tetris; complex skills; interactive behavior

Introduction

To achieve high degrees of skill in complex, real-time tasks– for example, the video game Tetris– players must master strategies that best exploit the game's event structure (Zacks, Speer, Swallow, Braver, & Reynolds, 2007), taking actions bound by the rules of cognition (reaction times, uncertainty, decision-making, motor activation) and fitting them as best they can into the temporal dynamics of the task before them. The process of learning these strategies can be fraught with sub-optimal lures (Fu & Gray, 2004)– strategies that are objectively worse, but remain stable because of the excessive costs associated with implementing those strategies in the first place. For example, any strategy that requires additional decision-making steps may require time spent deliberating between a pool of possible choices (as follows from Hick's Law (Hick, 1952)), a process one can't typically afford in time-pressured tasks. As expertise in a task domain increases, however, the time taken to deliberate between possible actions is thought to decrease (Logan, Ulrich, & Lindsey, 2016), freeing up precious temporal resources to evaluate and implement better strategies (which, subsequently, may help break out of a "plateau" in performance gains (Gray & Lindstedt, 2017)).

The present work examines a case of two strategies for a sub-task in the complex, real-time video game Tetris:

one simpler strategy that may often cost additional button-presses; and a more cognitively complex strategy that saves button-presses, but with an apparent and remarkably consistent "cognitive speed bump" prior to taking actions, even among the very best players in the world. Data are presented from three categories of players: true novices, regional championship players, and world championship players.

The video game Tetris

Tetris is a puzzle-solving game wherein players arrange falling blocks called "zoids" into the "pile" at the bottom of the screen to fill and clear lines for points. As players clear lines, the game level will increase, increasing both the speed at which zoids fall and the points awarded for clearing lines. The time pressure begins at a leisurely rate of zoids falling 2.5 rows (out of 20) per second, to the full 20 rows per second, and beyond. At these highest speeds is when players are most challenged, engaged, and require the best strategies to ensure survival. Players have three kinds of actions available: translations, to adjust the horizontal position of the zoid; rotations, to adjust the orientation of the zoid; and dropping, to quickly advance when the zoid is in the desired state. For the purposes of the present study, we can consider a single "episode" of Tetris to be the events that take place between a zoid's initial appearance and the placement of that zoid into the pile. Figure1 illustrates the task and input device.

Rotation sub-task strategies

Tetris, like any complex task, can be broken down into several subtasks: visual search, evaluation of options, aligning the zoid vertically, recovering from (or avoiding) errors, etc. Of particular interest is the "zoid rotation" sub-task, wherein players must efficiently and accurately rotate zoids into the desired orientation in order to drop them properly. Figure 2 shows the rotation behavior of zoids in Tetris. There are three categories of zoids: "static" zoids, which do not rotate at all; "flipping" zoids, which do rotate, but only osciallate betwen two orientations tates; and "rotating" zoids, which have four distinct orientations. These final zoids are of most interest, because there are two viable methods for approaching them: the "mono-rotational" strategy, and the "bi-rotational strategy."
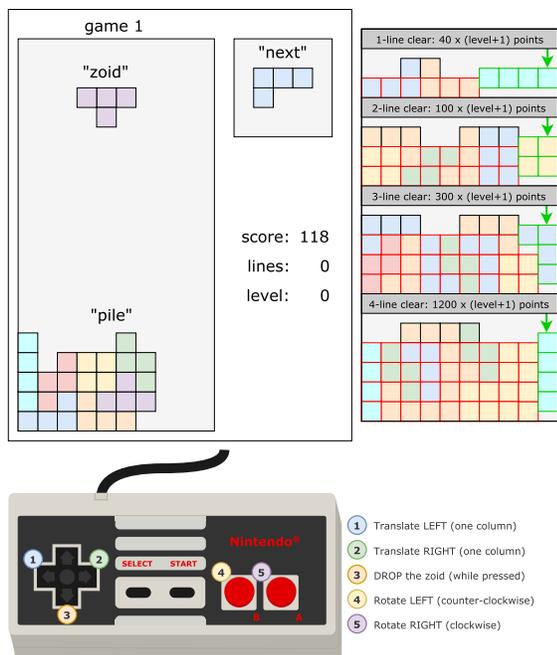
66

Figure 1: The anatomy of the game of Tetris. The upper left box illustrates the task environment, with a "zoid" falling down into the "pile". The upper right box outlines how line clears work, as well as the points awarded. The input device at the bottom is an illustration of an NES controller, labeled to indicate the actions available to players: translations (horizontal position), rotations (orientation of the zoid), and dropping (when the zoid is in the desired position and orientation).

Mono-rotational strategy   The mono-rotational strategy is simple: to get a zoid from its current orientation to the desired orientation, all a player needs to know is how many times to press the rotate button. This strategy ignores the option to rotate the zoid in two directions entirely– as far as this strategy is concerned, there is only one rotate button. This strategy comes with a critical inefficiency: sometimes the player will be required to press the rotate button three times (rather than once in the other direction), a dangerous cost for the break-neck pace of Tetris play.

Bi-rotational strategy   The bi-rotational strategy is more complex: the player utilizes both rotation buttons to rotate the zoid only in the direction that requires the fewest keypresses. As a result, the player must track more information, which may result in some cognitive slowdown, but as a reward for that additional effort, the number of keypresses required to rotate the zoid is kept to a minimum.

Cognitive implications of rotation strategies

So, why don't all Tetris players simply adopt the clearly "superior" bi-rotational strategy? The answer likely lies in the underlying cognitive costs (Gray, Sims, Fu, & Schoelles, 2006). Figure 3 illustrates a possible arrangement of the cognitive "pipeline" of the two rotational strategies. It is likely a
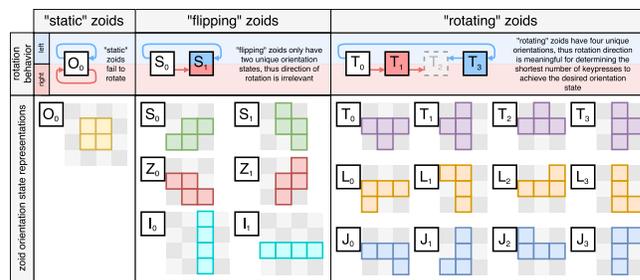

Figure 2: Internal representations of Tetris zoids and their rotation behavior. "Static" zoids do not rotate and require no rotation button presses; "flipping" zoids only have two orientations, so rotation is possible, but direction is irrelevant, as both direction result in the same final orientation; "rotating" zoids have four unique orientations, and thus require either multiple button presses, or use of both rotation directions to reach all possible states.

much simpler task– especially to a novice– to implement the mono-rotational strategy, as it requires no decision-making about which button to press, and likely much simpler memory structures to employ. Furthermore, the drawbacks of the mono-rotational strategy only become apparent when players attain sufficient skill to reach the highest levels of game speeds, wherein a single extra button-press can mean the difference between success and failure.

Method

We set out to examine how Tetris players of various skill levels make use of both the mono-rotational and bi-rotational rotation strategies.

Participants   The participants consisted of 30 players from both locally run Tetris tournament events on Rensselaer Polytechnic Institute's campus, and the 2016 Classic Tetris World Championship (CTWC) tournament: **novices** (N=10), **regional** champions (N=10), and **global** champion players (N=10).[1]

Procedure   All participants played a version of Tetris developed in-house known as "Meta-T" (Lindstedt & Gray, 2015), designed explicitly for research purposes. Data were collected either during the qualifying rounds of the tournament in question (for **novice** and **regional** champion players), or in private sessions arranged during the CTWC. To minimize the effect of the random factors of the game, all players played using the same fixed "random seed" meaning that zoids would appear in the same exact order for all players.

Results

The following analyses address three primary questions: (1) did players in each category actually differ in game scores,

---

[1]For a video example of CTWC players at peak performance during the tournament finals (with announcer commentary), refer to https://www.youtube.com/watch?v=DdfRQjb5o9k
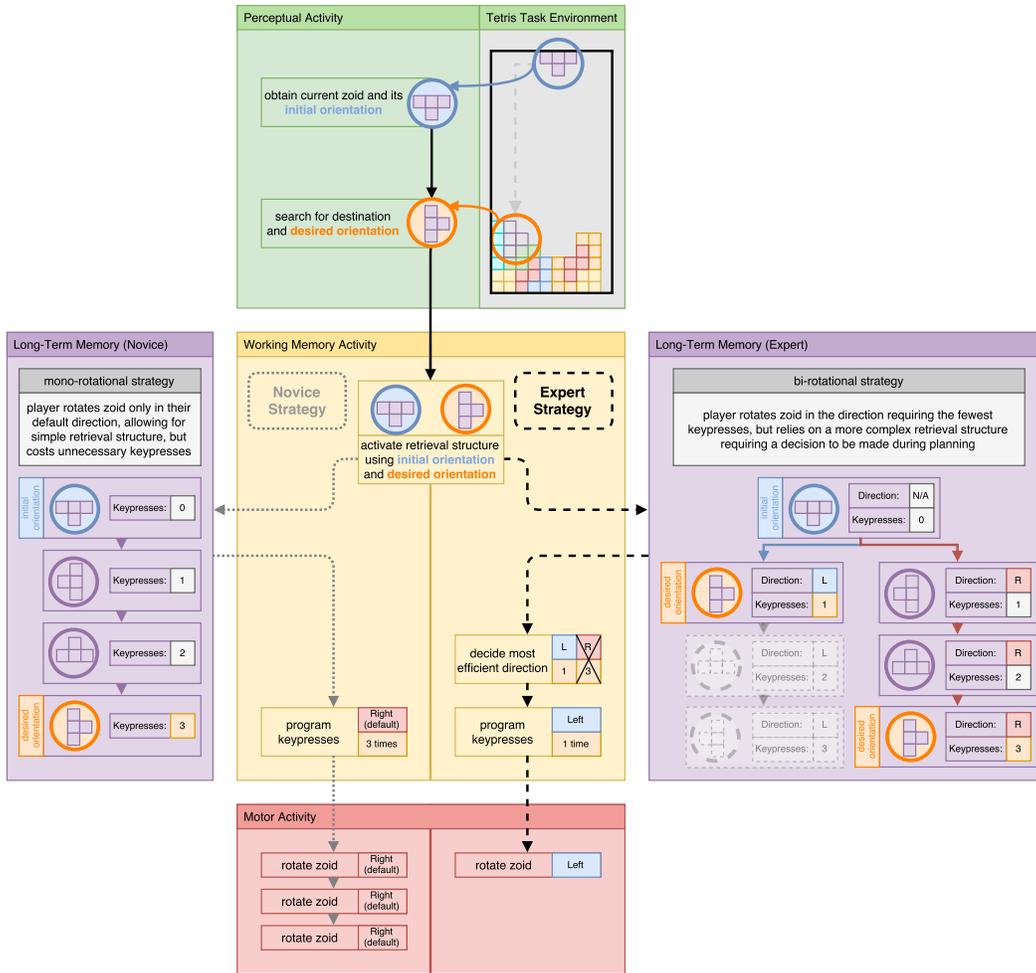
Figure 3: Diagram of the possible flow of information for each of two cognitive strategies involved in the zoid-rotation subtask. Players using the mono-rotational strategy rely only on simple memory retrieval structures to perform their task, but at the cost of sometimes executing two additional keypresses. Conversely, players using the bi-rotational strategy minimize number of rotations by employing both rotation directions, but at the cognitive cost of relying on a more complex retrieval structure. The retrieval structures presented here are intended to illustrate the relative complexity of the information involved in each strategy, not make strong predictions about the precise structure of memory.

(2) do they differentially adopt different rotational strategies, and (3) how do those strategies affect performance?

Differences in performance between skill groups

First, to verify that the players vary in their overall game performance, we conducted a one-way between subjects ANOVA to compare the effect of player skill on total game score (sixth root) for players in the novice, regional, and global skill conditions. There was a significant effect of player skill on game score (sixth root) for the three conditions [$F(2, 27) = 161.8$, $p < .00001$, $\eta^2 = 0.92$]. Figure 4 shows the trace of player scores over game difficulty levels, and the distribution of final game scores (sixth root). Post-hoc comparisons using a pairwise t-test with Bonferroni correction revealed that game score (sixth root) was significantly higher for global (M = 7.97, SD = .29) vs regional (M = 6.41, SD =

.33) players, both of which differed significantly from novices (M = 4.11, SD = .71). Regional players's scores also significantly differed from novices. These results indicate that the players do, in fact, differ significantly in their overall game performance.

Evidence of rotational strategy use

To rigorously compare players' strategy use, we measured the proportion of each player's rotation button-presses that were in that player's non-dominant rotation direction. First, we conducted a one-way between subjects ANOVA to compare the effect of player skill on log-proportion of non-dominant rotations for players in the novice, regional, and global skill conditions. Because the bi-rotation strategy is only meaningful for "rotating" zoids, we restricted the analysis to only include those zoids. There was a significant effect of player
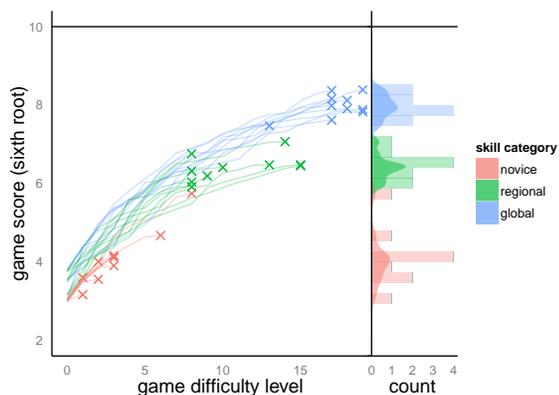
Figure 4: The trace of game scores across all game levels for all 30 participants. The 'x' marks the point at which the game terminated. The right panel shows the distribution of final game scores for players in each skill category.

skill on non-dominant rotations for the three conditions [F(2, 27) = 98.86, $p < .00001$, η^2=0.88]. Post-hoc comparisons using a pairwise t-test with Bonferroni correction revealed that the proportion of non-dominant rotations was significantly higher for global (M = .206, SD = .105) than for regional (M = .086, SD = .153) and novice (M = .005, SD = .011) players. Regional and novice players' proportions of non-dominant rotations did not differ. Figure 5 shows this effect.

These findings suggest that global champion players do, in fact, employ the bi-rotational strategy, while regional and novice players rely most heavily on the mono-rotational strategy, especially when the game's time pressure is highest.
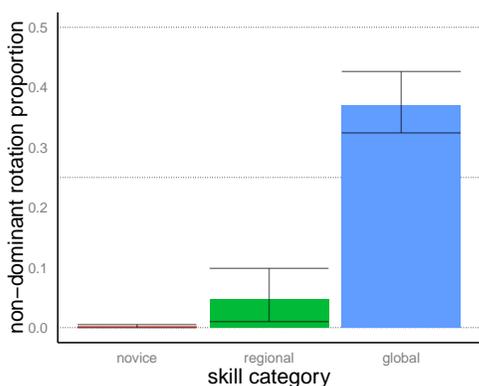

Figure 5: Comparison of non-dominant rotation usage for players of each skill category (x-axis). Error bars represent 95% confidence intervals.

Impact of rotational strategy use

To contrast the mono-rotational and bi-rotational strategies, we will examine both the motor costs (in terms of button-pressing efficiency) and the cognitive costs (in terms of initial reaction times) associated with each strategy. We consider regional players to by and large exhibit skilled usage of the mono-rotational strategy, while global players near-unanimously exhibit skilled usage of the bi-rotational strategy.

First, we examined the impact of rotational strategy adoption on motor efficiency by comparing the number of extra rotations for "rotating" zoids exhibited by players in each skill category. We conducted a one-way between subjects ANOVA to compare the effect of player skill on extraneous rotations for players in the novice, regional, and global skill conditions (collapsed across game levels and only for "rotating" zoids). There was a significant effect of player skill on non-dominant rotations for the three conditions [F(2, 27) = 37.18, $p < .00001$, η^2=0.73]. Post-hoc comparisons using a pairwise t-test with Bonferroni correction revealed that extra rotations were significantly lower for global (M = .022, SD = .018) than for regional (M = .756, SD = .262) players, and both were significantly lower than novices (M = 1.132, SD = .434) players. This result suggests that global players (who all appear to have adopted the bi-rotational strategy) are more efficient at rotating zoids than regional and novice players. Figure 6 shows the number of unnecessary button-presses observed in each player skill category.

Next, we examined the impact of each rotational strategy on players' initial reaction times. Initial reaction time for a given episode is considered to be the time between the zoid appearing on the screen and the player making their first button-press, time during which we assume any delays are due to additional cognitive processing. In addition, we want to examine those game episodes where players are close to the limit of their capability, when time pressure is high and a high degree of performance is critical, so we limit the analysis to only performance data from the highest level of game difficulty fully completed by each player.

We conducted two one-way ANOVAs, one for each player skill level (regional or global), to compare the effect of each zoid rotation type (static, flipping, or rotating) on players' initial reaction times. For the regional champion model (i.e., the skilled mono-rotational players), there was no significant main effect of zoid rotation type on initial reaction time [F(2, 27) = .20, $p = .821$, $\eta^2 = .01$]. For the global champion model (i.e., the skilled bi-rotational players), there was a significant effect of zoid rotation type on initial reaction time [F(2,27)= 82.55, $p < .00001$, $\eta^2 = .86$]. Post-hoc comparisons using a pairwise t-test with Bonferroni correction revealed that global players' initial reaction times for rotating zoids (M = 126.3ms, SD = 17.89) were significantly higher than for flipping zoids (M = 80.29ms, SD = 9.22ms) and static zoids (M = 57.33ms, SD = 6.55ms), and reaction times for flipping zoids were higher than for static zoids. Table 1 shows the means and standard deviations for each zoid rotation type for both regional and global champion players. This result suggests that global players (those employing the

bi-rotational strategy) are differentially affected by the zoids' number of possible rotations. Figure 8 shows this effect. Figure 7 expands to see how this effect plays out over the course of a game as difficulty increases.
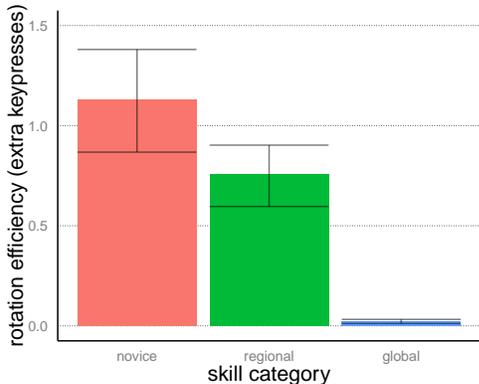

Figure 6: Rotation efficiency in terms of how many extraneous button presses are made to get the zoid into its ultimate orientation (a zoid could, at maximum, require two button presses to achieve a desired orientation, given the existence of the bi-rotational strategy). Error bars represent 95% confidence intervals.

Table 1: Means (and standard deviations) of reaction times for regional and global players for each zoid rotation type.

	regional champions		global champions	
static zoids	95.94	(47.16)	57.33	(6.55)
flipping zoids	84.01	(40.74)	80.29	(9.22)
rotating zoids	93.27	(45.22)	126.3	(17.90)

Discussion

The preceding analysis produced the following results. First, each skill category (novice, regional, and global) produced game scores and execution times that significantly differed from one another. Second, we demonstrated that novice and regional players adopted the mono-rotational strategy, and global champions adopted the bi-rotational strategy. Third, we showed that global champions (using the bi-rotational strategy) reduced the number of extraneous zoid rotations per episode to near-zero. Finally, fourth, we found expert players using the bi-rotational strategy had initial reaction times that were significantly impacted by zoid orientation type (static, flipping, and rotating, in increasing order of complexity), while there was no systematic difference for skilled users of the mono-rotational strategy.

We think the most striking finding in this study is that even extremely skilled players experience a cognitive "speed-bump" when adopting a strategy that involves more complex decision-making, and that this is true even at extreme levels of expertise. The global champions do, in fact, save time by adopting the bi-rotational strategy to avoid unnecessary zoid rotations, as evidenced by their decreased extraneous button presses (Figure 6). However, these global champion players appear to achieve their overall more efficient performance at the cost of considering more information, evidenced by their increased initial reaction times for more complicated zoid types (Figure 8). By contrast, the regional champions, who have adopted the simpler mono-rotational strategy, show no systematic slowdown when determining how much to rotate each type of zoid, regardless of how many orientation states that zoid may possess. The "bump" becomes even more evident when considering Figure 7, examining initial reaction times across all game difficulty levels. When comparing regional and global players at the same game difficulty level, global champion players consistently show increasing initial reaction times with the decision complexity of the zoid, while regional champions show no such differential increase.

With respect to "stable sub-optimal solutions" (Fu & Gray, 2004), it would seem that this "speed-bump" is somewhat damning– if even the best players in the world can't eliminate temporal costs associated with the bi-rotational strategy, then the novice would certainly not find this strategy initially palatable. Furthermore, the simple nature of the mono-rotational strategy may lead to a sort of "slippery slope" scenario wherein that simpler strategy gets practiced to the point of automaticity, further investing them in their chosen suboptimal solution and making it increasingly difficult to unlearn in favor of the superior strategy.

These findings also imply some interesting things about the decision structures involved in the two rotation strategies. While on the surface the bi-rotational strategy initially seemed to require a more complex decision process, it was also reasonable to expect that this complexity could be reduced with continued practice and chunking processes (Logan et al., 2016). This was evidently not the case. Even the best Tetris players in the world appear to remain susceptible to slowdowns associated with Hick's law (Hick, 1952), while those employing a simpler strategy show no such slowing. Or perhaps players at this level of skill may not, in fact, be merely plateaued, but have reached a "cognitive asymptote" for the adopted strategy in the task (Gray & Lindstedt, 2017). This leads me to a friendly amendment to Logan's (2016) suggestion that choice reaction times reduce with expertise: "Hick's law bends with practice, *except when it cannot.*"

Conclusion

In this study, we examined how novice, regional champion, and global champion Tetris players dealt with the zoid rotation subtask, adopting one of two rationally adaptive strategies: the mono-rotational strategy, which reduces cognitive complexity at the cost of motor inefficiency; and the bi-rotational strategy, which pays the cost of cognitive complexity for reduced motor activity, saving precious milliseconds required for high-level task performance.
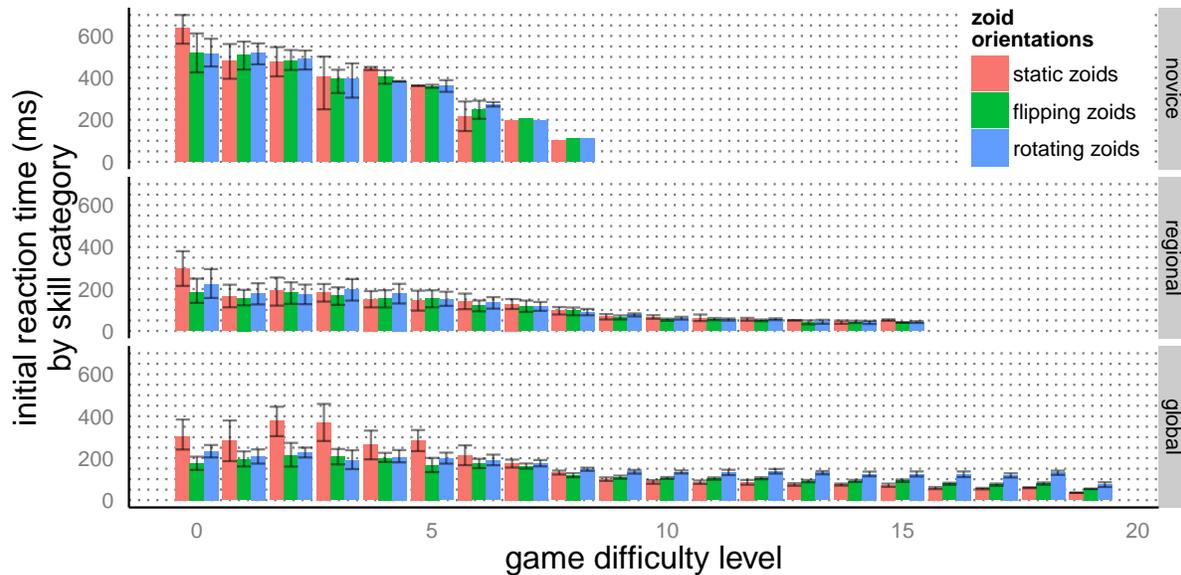
Figure 7: Initial reaction times for each zoid type for players in all three skill categories across the full breadth of game levels. Error bars represent 95% confidence intervals.
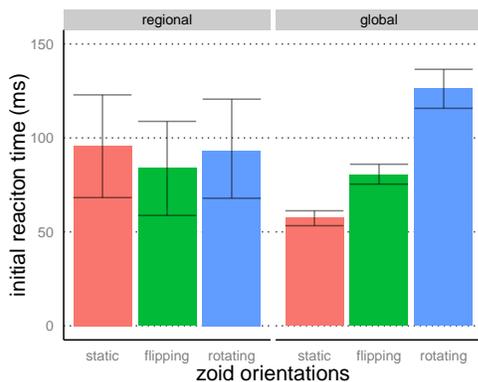

Figure 8: Mean initial reaction times per episode for champion players (regional and global) for the player's highest completed game level. Regional players adopt the mono-rotational strategy, while global players adopt the bi-rotational strategy. Error bars represent 95% confidence intervals.

The zoid rotation subtask is ultimately very simple: transform an object from state A to state D using 1 to 3 button-presses. Nevertheless, there exist two strategies for approaching it, a simple strategy which can be practiced to such an extent that it appears to require no decisions, and a more complex strategy which appears to possess some irreducible complexity. This implies that, despite the sometimes wondrous ability of the human cognitive machinery to adapt to any task, there may be limits– "cognitive asymptotes"– on just how cognitively streamlined a given strategy can be, even

for simple tasks.

References

Fu, W. T., & Gray, W. D. (2004). Resolving the paradox of the active user: Stable suboptimal performance in interactive tasks. *Cognitive Science*, *28*(6), 901–935. doi: 10.1016/j.cogsci.2004.03.005

Gray, W. D., & Lindstedt, J. K. (2017). Plateaus, dips, and leaps: Where to look for inventions and discoveries during skilled performance. *Cognitive Science*, *41*(7), 1838–1870. doi: 10.1111/cogs.12412

Gray, W. D., Sims, C. R., Fu, W.-T., & Schoelles, M. J. (2006). The soft constraints hypothesis: a rational analysis approach to resource allocation for interactive behavior. *Psychological review*, *113*(3), 461–82. doi: 10.1037/0033-295X.113.3.461

Hick, W. E. (1952). On the rate of gain of information. *The Quarterly Journal of Experimental Psychology*, *4*, 11 - 26.

Lindstedt, J. K., & Gray, W. D. (2015). Meta-T: Tetris as an experimental paradigm for cognitive skills research. *Behavior Research Methods*. doi: 10.3758/s13428-014-0547-y

Logan, G. D., Ulrich, J. E., & Lindsey, D. R. B. (2016). Different (key)strokes for different folks: How standard and nonstandard typists balance Fitts' law and Hick's law. *Journal of Experimental Psychology-Human Perception and Performance*, *42*(12), 2084-2102.

Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event perception: A mind-brain perspective. *Psychological Bulletin*, *133*(2), 273-293.