

Epistemic Action Increases With Skill

Paul P. Maglio

IBM Almaden Research Center
650 Harry Road, K54D-B2
San Jose, CA 95120
pmaglio@almaden.ibm.com

David Kirsh

Department of Cognitive Science, 0515
University of California, San Diego
La Jolla, CA 92093
kirsh@cogsci.ucsd.edu

Abstract

On most accounts of expertise, as agents increase their skill, they are assumed to make fewer mistakes and to take fewer redundant or backtracking actions. Contrary to such accounts, in this paper we present data collected from people learning to play the videogame Tetris which show that as skill increases, the proportion of game actions that are later undone by backtracking also increases. Nevertheless, we also found that as game skill increases, players speed up as predicted by the power law of practice. We explain the observed increase in backtracking as the result of an interactive search process in which agent-internal and agent-external actions are interleaved, making the cognitive computation more efficient (i.e., faster). We refer to external actions which simplify an agent's computation as *epistemic actions*.

Introduction

In this paper, we present experimental data which runs counter to an assumption that underlies most theories of skill learning: that more skilled agents take fewer redundant or backtracking actions than less skilled agents (e.g., Anderson, 1982; Logan, 1988; Newell & Rosenbloom, 1981). Intuitively, skilled agents ought to make fewer mistakes than unskilled agents and therefore ought to backtrack less and take fewer redundant actions. However, our studies of how people improve at playing the videogame Tetris reveal that sometimes getting better means backtracking more. Better players use the world better, even in the limited world of a Tetris board. Consequently, we explain the observed increase in backtracking as the result of interactive search in which agents reduce cognitive load by interleaving internal and external actions.

Previously, we introduced the term *epistemic action* to describe external actions that can be used to reduce the memory, time, and probability of error of agent-internal computation (Kirsh & Maglio, 1994). We justified our view by presenting data collected from Tetris players at all skill levels in which many examples of recurring backtracking behaviors could be found. We could not, however, prove that better players performed more epistemic actions. Here we present longitudinal data on the acquisition of Tetris skill which show that: (a) Certain sorts of backtracking increase as skill develops; and (b) despite this increase, Tetris skill resembles other skills in following the power law of practice (Newell & Rosenbloom, 1981). Taken together, these results support our claim that epistemic actions play a substantial role in skilled behavior.

In what follows, we first briefly describe Tetris, and then present analysis and discussion of our behavioral data.

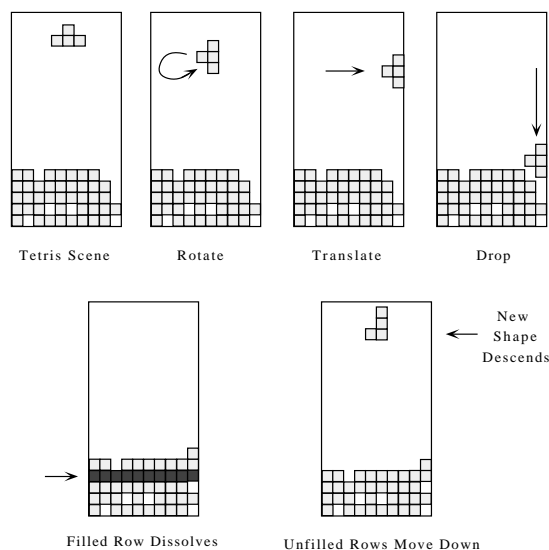









Figure 1: In Tetris, shapes fall one at a time from the top of the screen, eventually landing on the bottom or on top of shapes that have already landed. As a shape falls, it can be rotated, and moved to the right or left. The objective is fill rows of squares all the way across the screen. Completely filled rows dissolve and all partially filled rows above move down.

How to Play Tetris

Tetris is a popular videogame in which players maneuver falling shapes into specific arrangements on the computer screen. There are seven shapes (which we call *zoids*): , , , , , , . These fall one at a time from the top of a screen that is 10 squares wide and 30 squares high (see Figure 1). Each zoid freely falls until it lands on the bottom edge of the screen or on top of the squares of a zoid that has already landed. Once a zoid comes to rest, another begins falling from the top, starting the next Tetris *episode*. While a zoid falls, the player can control it, either *rotating* it 90° counterclockwise with a single keystroke, or translating it to the *right* or to the *left* one square with a single keystroke. To gain points, the player must carefully land zoids so that rows fill up with squares all the way across the screen. Filled rows then disappear and all filled squares above it drop down. This process is called *clearing rows*. As more rows are cleared, the game speeds up, and controlling how zoids land becomes more difficult. Filled squares pile up as unfilled, uncleared

rows become buried under poorly placed zoids. The game ends when the screen becomes clogged with incomplete rows and new zoids cannot descend. Thus, clearing rows serves the purposes both of scoring points and of delaying the game's end.¹

We recorded data from two players who practiced Tetris for about 20 hours each, encompassing approximately 40,000 keystroke interactions with the game. Neither participant had played the game before, and both agreed not to play except under computer observation during the course of the study. We now turn to these data.

How Players Improve With Practice

As players practice more, the number of rows they clear—their game score—increases. As Tetris players know, the rate of improvement is misleading, for the game speeds up as rows are cleared. Hence, players encounter different task demands during the course of a single game.

To ensure that we did not compare experts in high speed games with novices in slow speed games in our study, we controlled for the effects of game speed during analysis by separating episodes into three speeds, slow, medium, and fast, a division roughly following skill: everyone plays at slow speeds, better players attain medium speeds, and only the best players achieve fast speeds. Because both participants always played part of their games at slow speeds, we will compare behavior based solely on data gathered from the slow portions of the games.

Speed-ups Follow the Power Law of Practice

Typically, practice improves performance in accordance with a power function of practice time or practice trials (Newell & Rosenbloom, 1981) either by decreasing the time to react to stimuli by taking a single action (Seibel, 1963), or by decreasing the overall time it takes to perform a task that requires a sequence of actions (Crossman, 1959). In Tetris, the time to perform a sequence of actions can be measured within an individual episode as the interval between the time the falling zoid first becomes visible and the time of the last action that the player takes (see Figure 2). The time to take a single action can be measured as the interval between consecutive keypresses in episodes in which more than one action was taken (see Figure 3). In addition, another component of the overall time to place a zoid is the latency of the player's first action, that is, the interval between the time the falling zoid becomes visible in an episode and the time of the player's *first* keypress (see Figure 4).

In all three cases, our data indicate that performance speeds up according to the power law of practice. Figure 5 shows one example: when plotted on a log-log scale, the time between keypresses follows a straight line. To see that all our data are better fit by power curves than by other curves, consider Table 1. Following Newell and Rosenbloom (1981), we compared the correlations of the best-fit regressions for




Figure 2: The time to place a zoid is defined as the interval between the time the zoid first appears on the board and the time of the last action that the player takes to maneuver it. In this case, the zoid appears at time t_1 and the last action is taken at time t_2 . The overall time, then, is the interval $t_2 - t_1$.

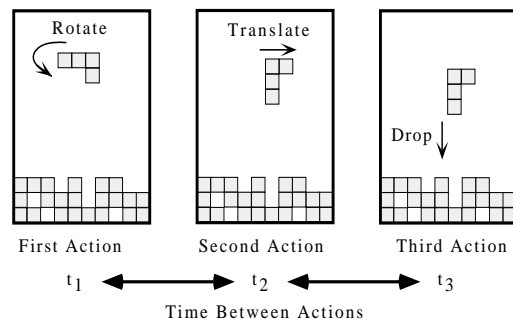


Figure 3: The time between actions is defined as the interval between consecutive keypresses. This figure illustrates two intervals: $t_2 - t_1$ and $t_3 - t_2$.




Figure 4: The time of the first action in an episode is measured from when the zoid first appears on the board. In this figure, the zoid first appears at time t_1 and the first action occurs at time t_2 . The time of the first action, then, is $t_2 - t_1$.

¹In addition to rotation and translation, the player can *drop* a falling zoid instantly to the bottom, effectively placing it in the position it would eventually land in if no additional keys were pressed. Dropping is an optional maneuver, and not all players use it. Dropping speeds up the pace of the game, creating shorter episodes without affecting the free-fall rate. We will not discuss dropping.

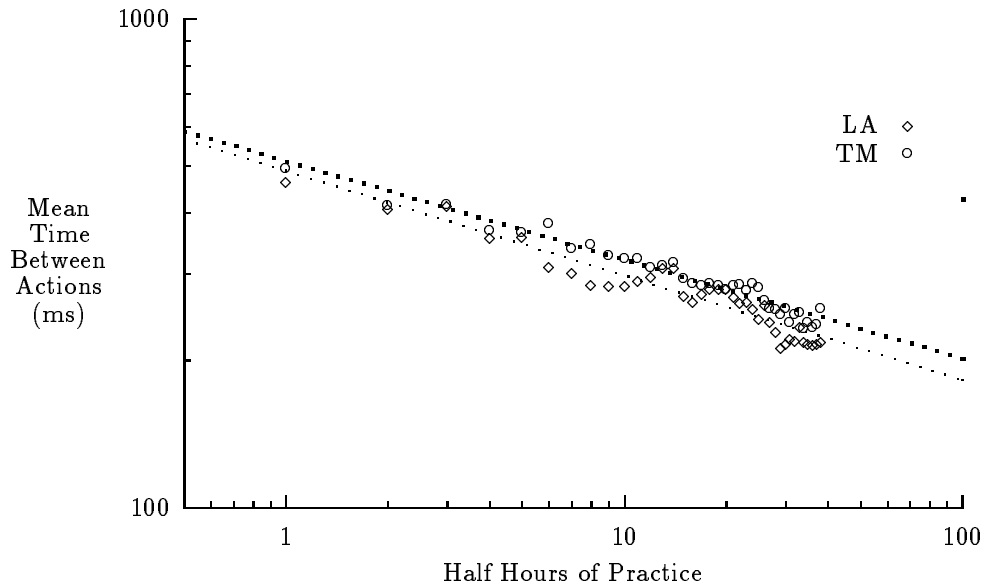


Figure 5: The time between consecutive actions decreases following a power function of practice. The data plotted on this graph are drawn from slow episodes only. The lines are the best-fit linear regressions of the data in log-log space.

	Linear			Exponential			Power Law		
	$T = A - BN$			$T = Be^{-\alpha N}$			$T = BN^{-\alpha}$		
	A	B	r^2	B	α	r^2	B	α	r^2
Time to Place a Zoid (s)									
LA	4.087	47.78	.7711	5.019	0.012	.8020	6.010	0.149	.8306
TM	5.195	53.82	.6781	5.201	0.012	.7162	6.508	0.171	.8481
Time Between Actions (ms)									
LA	370.2	4.693	.7721	374.2	0.016	.8439	488.1	0.214	.9108
TM	392.6	4.743	.8117	396.6	0.015	.8789	510.0	0.202	.9557
Time of First Action (ms)									
LA	936.8	12.23	.7234	975.2	0.016	.7424	1253	0.210	.7515
TM	1117	12.27	.5332	1106	0.013	.5224	1436	0.189	.6906

Table 1: Power curves fit the data better than lines or exponentials. To determine that the players speed up according to the power law, we followed Newell and Rosenbloom's (1981) method of comparing various regressions—a straight line, an exponential, and a simple power function—on the data. For the equations shown in the table, T is the performance measure (time), N is practice block (30-minute intervals), α is the rate of decrease determined by the regression, and B is a constant also determined by the regression. For both participants and for each measure, a power function fits the data better than a line or an exponential.

lines, exponentials, and simple power functions.² As the table shows, in each case, the data are fit best by a power function of practice. This means that *Tetris skill is like most other skills* because power-law improvement is universal (see Newell & Rosenbloom, 1981).

Because players become faster with practice, one might expect that players' actions also become more precise. That is, Tetris experts should not only take action faster than beginners, but they should take only the actions necessary to maneuver the falling zoid to its final position and orientation because experts make fewer mistakes, backtrack less, or simply see the solution sooner than beginners do. As we will show, our data indicate just the opposite: *the number of apparently extraneous actions increases with practice*. This result is surprising because theories that explain power-law improvement, for instance, by accumulating chunks (Newell & Rosenbloom, 1981) or cases (Logan, 1988), assume that behavior becomes more efficient and economical with practice (Crossman, 1959).

Backtracking Increases With Skill

As stated, we found that sometimes more skilled Tetris players actually take more *extra* actions—that is, actions that are later undone by backtracking—than less skilled players. To see that backtracking increases with skill, let us define *backtracking* or *extra* actions in Tetris to be actions that do not lie on the shortest path from the falling zoid's initial location and orientation to its final location and orientation (see Figure 6).

Using this definition of backtracking, we calculated the mean number of extra rotations. For analysis, we grouped the data into three consecutive six-hour intervals. We first calculated the mean number of extra rotations per episode for each game, and then used these averages as the raw scores for analysis. Figure 7 illustrates the results for LA.³ As shown, the average number of extra rotations per episode is significantly greater in expert games than in intermediate games. Backtracking increases with practice.

Now it may be objected that because the average number of extra rotations for LA is only around 0.2 per episode, extra rotations must occur relatively infrequently. Extra rotations occur in 7% of the episodes in which LA was an expert, in 5% of the episodes in which he was an intermediate, and in 4% of the episodes in which he was a beginner. These frequencies differ significantly ($p < .01$), but extra rotations are clearly the exception rather than rule, and therefore it might be illustrative to investigate the contexts in which they occur.

Figure 8 reveals that the percentage of episodes containing extra rotations varies by zoid type, and that the number of




Figure 6: *Backtracking* actions do not lie on the shortest path between a zoid's initial location and orientation and its final position. The trajectory shown in (b) is a shortest path. The trajectories shown in (c) and (d) contain backtracking.

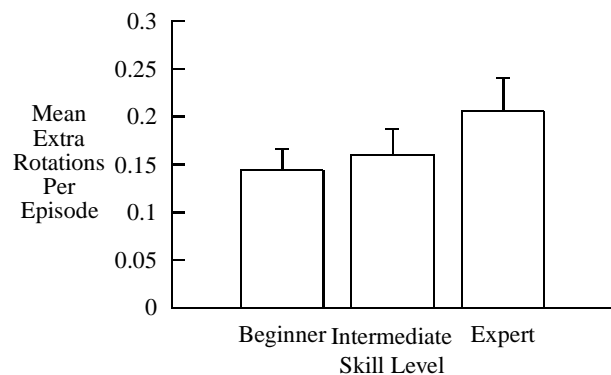


Figure 7: Extra rotations increase with expertise for LA. More precisely, the mean number of extra rotations was greater when LA played at the expert level than when he played at the beginner level.

²Newell and Rosenbloom (1981) also discuss fitting data to generalized curves by adding an additional parameter to account for prior experience. In particular, they consider fitting exponentials of the form $T = Be^{-\alpha N} + E$, and power functions of the form $T = BN^{-\alpha} + E$, where E is the additional constant used to represent prior experience. By incorporating additional parameters, better fitting regressions can always be found. But because our data contain only 38 points for each participant, there is the danger of overfitting the points. Therefore, we used the *simplest* functions in each case, that is, the functions containing two (rather than three) free parameters.

³Because of space limitations, the rest of the discussion will focus on LA's data. See Maglio (1995) for discussion of TM's data.




Figure 8: The percentage of episodes containing extra rotations varies both by skill and by zoid type. The data plotted in this graph show that the extra rotations occur more frequently for 1 and 2 than for other zoids at all skill levels, but especially at the expert level. Although the number of extra rotations increases with skill for all zoid types except 4, the number of extra rotations increases most for 1 and 2.

extra rotations increases most for 1 and 2. These data suggest that although extra rotations occur infrequently, they cannot be the result of simple motor errors. This follows because motor mistakes ought to affect each type of zoid equally. If extra rotations result from a baseline error in motor control processes (i.e., experts can recover from overshooting the desired orientation but beginners cannot), there is no a priori reason to suppose that the frequency of errors for 1 would be less than the the frequency of errors for 4. Errors should be distributed randomly. Because the percentage of extra rotations differs among zoid types, the conjecture that extra rotations are the result of recovering from simple motor mistakes must be ruled out.

Perhaps, however, extraneous rotations result from perceptual errors. For example, perceptual confusion might result when the falling zoid is 1 but there is a natural place to put 1 and not 2. More precisely, let us define a *mirror episode* to be a board configuration in which any placement of the falling zoid will create a hole, but in which some placement of the falling zoid's mirror image does not (see Figure 9). Obviously, mirror episodes can only occur for the zoids with mirror images: 4, 5 and 6, 7. In this case, the percentage of extra rotations in mirror episodes does not differ from the percentage of extra rotations in non-mirror episodes for the relevant types of zoids (see Figure 10). Thus, backtracking rotations are not the result of this type of perceptual error.

Early Rotations Aid Decision Making

We conjecture that rather than being the result of motor or perceptual errors, extra rotations are computationally efficient epistemic actions. For instance, extra rotations might help the player decide where to place the falling zoid. If this is the case,




Figure 9: In *mirror episodes*, there is a good place to put the falling zoid's mirror image but no good place to put the falling zoid itself. If players backtrack more because they make perceptual mistakes, extra rotations might be more frequent in mirror episodes.

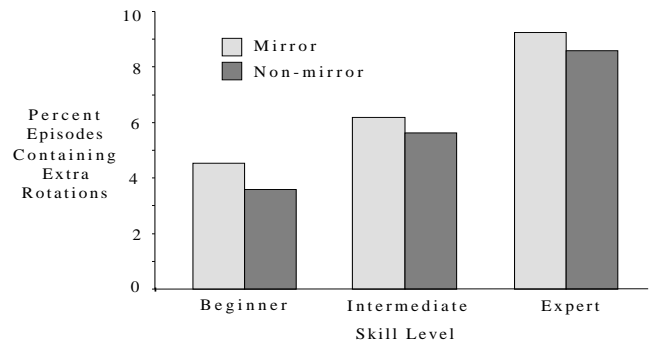


Figure 10: The frequency of extra rotations does not depend on mirror episodes for LA. Within each skill level, the percentage of extra rotations does not differ significantly between the mirror and non-mirror conditions ($\chi^2 < 1$ in all cases).

we would expect extra rotations generally to occur *before* the player has decided where to put the zoid. In particular, if a zoid is rotated into its final orientation before a player has made a final decision about orientation, then the player might continue to rotate the zoid to assist decision making. This implies that extra rotations should occur most often when the zoid is in its final orientation *before* the player is ready to judge whether the orientation is correct. This might happen in two ways: either because the player rotates the zoid very rapidly soon after it appears, or because the zoid appears on the board close to its final orientation. And in fact, the data for LA show that extra rotations do occur primarily when the zoid is in its final orientation early in the decision-making process (see Figure 11). In general, the final decision about orientation is not made much before 1130 milliseconds. For 1, the mean time to put the zoid into its final orientation is 1127 ms ($SD = 99$ ms), and for 2 and 3 the mean is 1122 ms ($SD = 71$ ms). As Figure 11 shows, however, the mean time that these zoids are rotated into their final orientation and then later unnecessarily rotated is 400–500 milliseconds earlier. For 4, the mean time is 676 ms ($SD = 149$ ms), and for 5 and 6 the mean time is 754 ms ($SD = 296$ ms). If extra rotations were the result of motor or perceptual mistakes, however, there is no reason to suppose that they would occur most often early in an episode. Thus, it seems that external

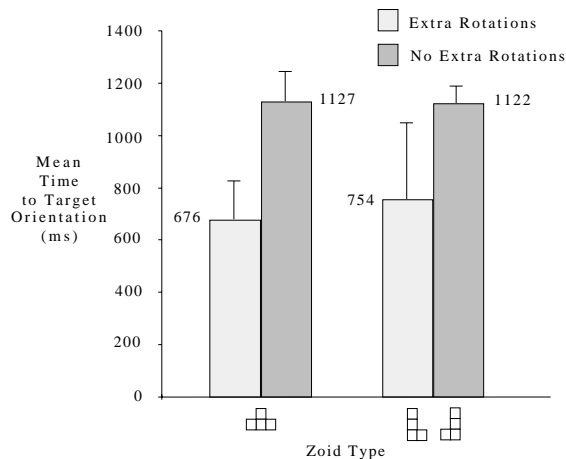


Figure 11: The mean time at which the falling zoid is first rotated into its target orientation is less when there are extra rotations than when there are no extra rotations. The data shown are from slow episodes in which LA was an expert. The mean time with extra rotations differs significantly from the mean time without in both cases: for I, $t(198) = 2.82, p < .005$; for II, $t(293) = 2.24, p < .026$.

rotation is being used to help make a placement decision.

To summarize, the number of extraneous rotations performed by LA increased with skill. This finding is counterintuitive because skilled players would be expected to head more precisely toward their goals. Yet because Tetris performance improves according to a power function of practice, Tetris skill must be like most other skills. Taken together, we believe these results support our hypothesis that redundant actions are epistemic actions which both simplify perceptual computation and play a natural role in skilled behavior (Kirsh & Maglio, 1994). We conclude with a brief discussion of some implications of this view.

Epistemic Actions Simplify Perception

It is no surprise, of course, that people offload symbolic computation (e.g., preferring paper and pencil to mental arithmetic; Hitch, 1978), but it is a surprise to discover that people offload perceptual computation as well. In Tetris, we conjecture that extra rotations are used to simplify the search for the best zoid placement by cueing retrieval from an orientation-specific index of zoids and board configurations (Maglio, 1995). In this way, we believe Tetris players set up their external environments to facilitate perceptual processing—much as gin rummy players physically organize the cards they have been dealt (Kirsh, 1995), and much as airline pilots place external markers to help keep track of appropriate speed and flap settings (Hutchins, 1995).

There is empirical evidence that people minimize their use of perceptual computational resources. For example, Ballard, Hayhoe and Pelz (1995) found that people performing a block-arranging task organized their actions and eye movements so as to minimize their working memory load. Rather than using memory of the visual scene to guide their actions, participants tended instead to move their eyes to gain just the information needed for their next action. These findings suggest that the

cost of moving the eyes to gather information is low relative to the cost of using short-term memory. It follows that serial processing (i.e., interposing eye movements between internal computations) is more computationally efficient than parallel processing (taking in all the information at once and calculating a plan) because of the high cost of internally holding and using partial results.

But we believe the idea that agents can rely on the world to provide an external memory which substitutes for internal memory (e.g., O'Regan, 1992) is only part of the story. Whereas eye movements are active from the point of view of the visual system—they serve to change the focus, and therefore to act on the agent's perceptual input—they are passive from the point of view of the task environment. By contrast, skilled Tetris players' extra rotations are *active* in the task environment yet change the perceptual input in much the same way that eye movements do. External rotations do more work than eye movements because rotation is a domain action. Rotating the zoid actually changes the stimulus, whereas moving the eyes does not. Thus, when physically rotating the zoid for its computational effect, the external world functions not as a passive memory buffer, simply holding information to be picked up by looking, but the world in interaction with the agent functions more like a working memory system, that is, like an *interactive visuo-spatial sketchpad*.

Acknowledgements

This research was supported in part by NIA Grant AG11851.

References

- Anderson, J. R. (1982). Acquisition of cognitive skill. *Psychological Review*, 89, 369–406.
- Ballard, D. H., Hayhoe, M. H., & Pelz, J. B. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience*, 7, 66–80.
- Crossman, E. R. (1959). A theory of the acquisition of speed skill. *Ergonomics*, 2, 153–166.
- Hitch, G. J. (1978). The role of short-term working memory in mental arithmetic. *Cognitive Psychology*, 10, 302–323.
- Hutchins, E. (1995). How a cockpit remembers its speed. *Cognitive Science*, 19.
- Kirsh, D. (1995). The intelligent use of space. *Artificial Intelligence*, 73, 31–68.
- Kirsh, D. & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18, 513–549.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 9, 492–527.
- Maglio, P. P. (1995). *The computational basis of interactive skill*. Doctoral dissertation, University of California, San Diego.
- Newell, A. & Rosenbloom, P. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: LEA.
- O'Regan, J. K. (1992). Solving the “real” mysteries of visual perception: The world as outside memory. *Canadian Journal of Psychology*, 46, 461–488.
- Seibel, R. (1963). Discrimination reaction time for a 1,023 alternative task. *Journal of Experimental Psychology*, 66, 215–226.