

Some Epistemic Benefits of Action: Tetris, a Case Study

David Kirsh Paul Maglio

dkirsh@ucsd.edu pmaglio@ucsd.edu

Cognitive Science Department

University of California, San Diego

La Jolla, CA 92093-0515

Abstract

We present data and argument to show that in Tetris—a real-time interactive video game—certain **cognitive** and **perceptual** problems are more quickly, easily, and reliably solved by performing actions in the world rather than by performing computational actions in the head alone. We have found that some translations and rotations are best understood as using the world to improve cognition. They are not being used to implement a plan, or to implement a reaction. To substantiate our position we have implemented a computational laboratory that lets us record keystrokes and game situations, as well as allows us to dynamically create situations. Using the data of over 30 subjects playing 6 games, tachistoscopic tests of some of these subjects, and results from our own successful efforts at building expert systems to play Tetris, we show why knowing how to use one's environment to enhance speed and robustness are important components in skilled play.

Introduction

In this paper we present data and argument to show that in Tetris—a real-time interactive video game—certain **cognitive** and **perceptual** problems are more quickly, easily, and reliably solved by performing actions in the world rather than by performing computational actions in the head alone.

In Tetris, there are only four actions a player can take: translate right, translate left, rotate, drop. Tetrazoids—henceforth zoids—enter from the upper boundary of a rectangular playing field at a fixed speed which increases as the game proceeds, leaving the player with less and less time to make the judgments involved in choosing and executing a placement. Because all actions move the current zoid one way or another, every action the player takes has the effect of bringing the current zoid closer to its final

position or farther from it. See figure 1.

Owing to the pace of the game it is not surprising that players in the earliest phase make moves before they can know where they wish to place the current piece. We have found that often these moves are best understood as having an *epistemic* function. They are not intended to achieve the *pragmatic* end of bringing a piece closer to its goal position. They are being used to change the world so as to help the agent acquire vital information early on.

Surprisingly, such epistemic functions are not confined to the earliest phases. Some translations and rotations occurring in later phases of decision and execution are also best understood as using the world to improve cognition. Some, for instance, seem designed to help the player *identify* a piece, or to *verify* that a particular action is a good one to take, or to *minimize the mental rotation* necessary to decide on a placement. We see the general function of these actions to be that of improving cognition by:

1. reducing the space complexity of mental computation;
2. reducing the time complexity of mental computation;
3. reducing the unreliability of mental computation.

These are not easy claims to defend in a game as complex as Tetris. We have implemented a computational laboratory that lets us record keystrokes and game situations, as well as allows us to dynamically create situations. Using the data of over 30 subjects playing 6 games, tachistoscopic tests of some of these subjects, and results from our own successful efforts at building expert systems to play Tetris, we will try to defend our conclusion that knowing how to use one's environment to enhance the speed and robustness of mental computation are important components in skilled play.




Figure 1: Translating for verification

In figure 1 we have an example of the basic game of Tetris. As each piece descends from the ceiling, the player must choose a region in which to place the piece. When a row of squares fills up, it disappears and all the rows above it drop down. As the game speeds up, achieving good placements becomes increasingly difficult.

One technique most players strike on to reduce error is to translate zoids to the wall. In figure 1, we see an instance of how $\square\square\square$ is regularly translated to the outer right wall and back again before it is dropped. The explanation we prefer is that the subject confirms that the column of the zoid is correct, relative to his or her intended placement, by quickly moving the zoid to the wall and then with eyes on the contour tapping out the number of squares to the edge of the intended placement.

The idea that external actions can simplify mental computation is commonplace when symbol manipulation is involved. The activities of adding, accounting, composing, navigating (Hutchins, 1990), etc., are more difficult if agents must rely on their own memory without aid from external supports. Writing reduces the space complexity of the mental computations involved. When symbol manipulation is not involved, however, especially in tasks requiring quick response, it is less widely appreciated that certain non-perceptual actions can simplify mental computation.

For instance, there is a tacit belief among planning theorists that intelligent behavior is either *reactive* or *planned* (Tate, Hendler & Drummond, 1990). In environments where an agent has time for reflection or forethought, planning can occur, and the agent may benefit from the advantages of previewing possibilities, hence mental backtracking is possible and local minima can be avoided. In rapidly changing environments, where there is not enough time to formulate a planned response—as is typical of arcade video games—the advantage lies with agents who have precompiled plans into reactions (Agre & Chapman, 1987). Where time is scarce, reactive systems, based on reliable statistical models of contingencies plus rapid sensing of environmental conditions, can be expected to score higher than systems which plan, unless, of course, there is enough time between actions to combine elements of both planning and reaction (Georgeff & Lansky, 1987). In each case, though, the assumption is that actions either are perceptual or should, if possible, bring the system physically closer to its goals.

A significant percentage of non-perceptual actions in Tetris actually take the agent physically farther from its ultimate goals. These costs are worth incurring because they are more than made up for by the epistemic or computational benefits they provide. They are rational actions if seen to be directed at transforming the agent's state, rather than the world's.

The idea that real-time systems must act so as to intelligently regulate their intake of environmental information is, at present, a topic of considerable interest (Simmons et al., 1992). But whereas existing inquiries have tended to focus on *control of attention*—the selection of elements within an image for further processing—or *control of gaze*—the orientation and resolution of a sensor—as the means of selecting information, our concern in this paper is with *control of activity*. We wish to know how an agent can use ordinary actions—not sensor actions—to unearth valuable information that is currently unavailable, hard to detect, or hard to compute.

Early and Late Epistemic Actions

Let us call the phase spanning the period when pieces are being identified, and the onset of the phase when plans are implemented, the *identification phase* (see figure 3). The duration of this phase varies with piece, subject, and the speed of the game. For instance, for $\square\square\square$ and $\square\square$, when the game is proceeding at average speed, the identification phase probably begins around 600 ms and ends around 1800 ms, whereas for $\square\square$ and \square the identification phase begins around 800 ms and ends around 2200

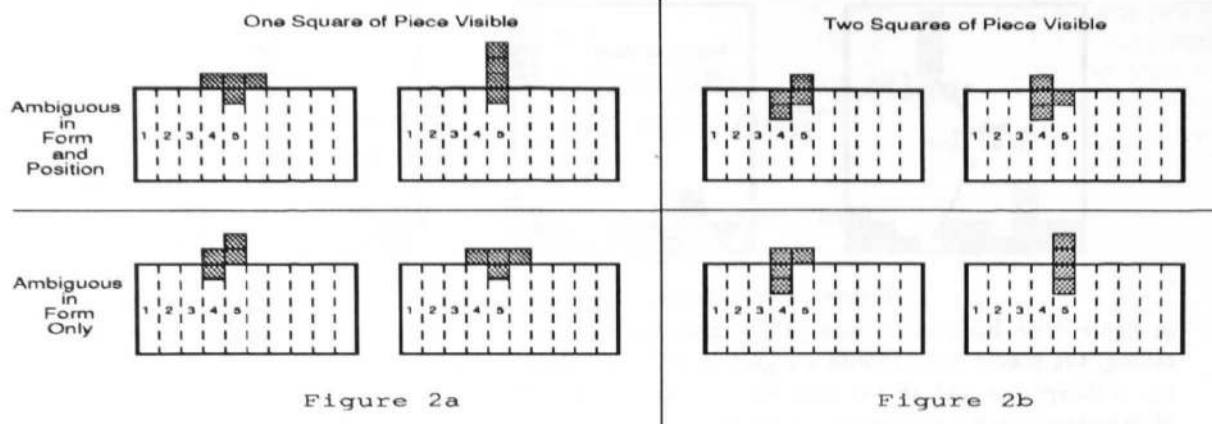


Figure 2: Rotating for early discovery

Rotating certain zoids very soon after they appear is a practical method for getting extra information about pieces. Here we see zoids as they first enter the playing field, in 2a they are one square in, in 2b they are two squares in. The upper portion of both 2a and 2b show zoids that look identical at this stage, both in position and in form. The bottom portions show zoids that look identical in form alone. Careful examination reveals that they are in different columns. Players are not explicitly aware of this column difference. The data show that players do not come out rotating, as we originally thought, but rather have a great burst once they are two rows out. At this point they show considerable sensitivity to column difference. Players have a much greater tendency to rotate zoids ambiguous in both form and position (such as those seen in the upper portion of 2b) than they have of rotating zoids that are ambiguous in form alone. By rotating ambiguous zoids early, players are able to make faster identifications, thereby either setting up the conditions for testing candidate placements early or setting early constraints on a candidate generator.

ms. The period before this phase we call the *pre-identification phase*, and the period after it, the *post-identification phase*.

As figures 2 and 1 illustrate, players at the intermediate and expert level regularly perform unambiguously epistemic actions in the pre- and post-identification phases.

1. Very early in the pre-identification phase players often rotate certain zoids before they have completely emerged, as if trying to disambiguate the zoid from all others as soon as possible.
2. In the post-identification phase players often drop certain zoids only after translating them to the nearest outer wall and then back again, as if to verify the column of placement.

The value of these actions is easy to appreciate. The first procedure, **rotate early**, serves to unearth facts otherwise hidden until later. When a zoid first enters the playing field and only a fraction of its total form is visible, the player must rely on subtle clues to disambiguate it. See figure 2. To be sure, players need not follow a strategy that requires them to disambiguate zoids as quickly as possible. But, in fact, we have noted that the more perfectly ambiguous a piece is, the more it is rotated early. The simplest explanation is that early rotation is for fact finding. By rotating a partially hidden piece, a player

un-occludes part of it, thereby scaring up new information. The faster this may be done, the sooner an unambiguous image of the piece can be formed. The value of gaining this information early presumably outweighs the cost of possibly rotating a piece beyond its "goal" position.

The second procedure, **translate zoid to an edge and translate back again**, serves to confirm the column which the piece is currently in. After having chosen a spot to place the current piece, and having implemented a plan to direct the piece to that spot, a player may wish to confirm that he or she has succeeded in moving the piece to its intended column. This further phase is most useful when the piece is still high above the active contour and about to be dropped. See figure 1. This action cannot be confused with a pragmatic action, for by definition it requires moving the piece away from the currently intended column. On the occasions when it is performed, the pragmatic cost is more than offset by the benefit of reducing possible error.

Epistemic Actions During the Identification Phase

Actions performed in the identification phase are more difficult to classify, particularly since an ac-

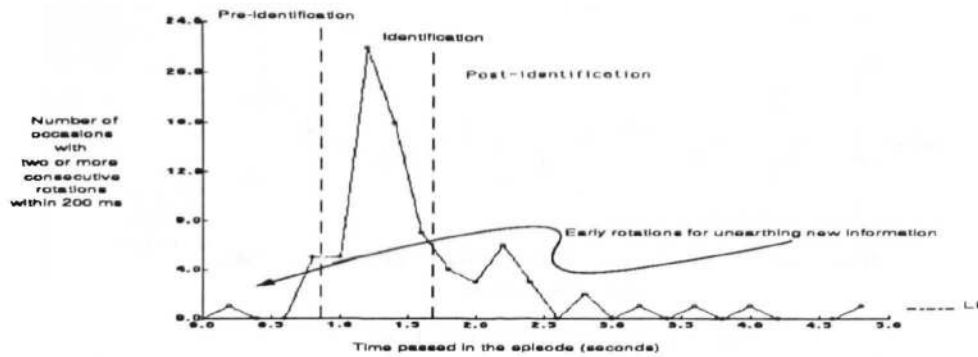


Figure 3: Rotating to help identification

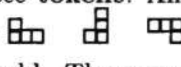
When are the fastest rotations performed? Here we see subject PD's time course of double rotations for L's. By a double rotation we mean two rotations in very quick sequence (i.e., 200 ms or less). Two points should be noted. First, the greatest number of double rotations occurs in the region of 800 ms to 1600 ms: a period that corresponds nicely with the period in which pieces are being fully identified, as shown in figure 4. Second, PD had one double rotation at 200 ms, well before he could have identified the piece. PD also had 6 single rotations before 200 ms, a fact we interpret as confirming our conjecture that very early rotations serve to dig up information that otherwise would be hidden for another 400 ms. Similar results hold for all the subjects we have examined so far.

tion may serve both epistemic and pragmatic functions simultaneously. For instance, a zoid rotated in the direction needed for final placement, may, at one and the same time, help the player make an identification, while advancing the cause physically. The two functions—epistemic and pragmatic—are logically distinct, though it is hard to prove which function a given action subserves. Three epistemic functions an action may perform in the identification phase are:

1. help to *identify* a piece's type;
2. help to *verify* the identity of a piece once it is typed, i.e., reduce probability of misidentification;
3. help to *generate* candidate placements.

As can be seen in figure 3, subjects are more prone to have a burst of rapid rotations in the identification phase than at any other time. These actions of rotating pieces in the world take far less time than rotating pieces mentally. A natural conjecture is that they are being used to either facilitate identification or to reduce the probability of misidentification during this crucial period. This becomes more convincing when we consider how players decide where to place pieces.

Although we do not yet know exactly how a player selects where to place a piece, we have good reason to believe that some matching of piece shape to potential placement location must occur. To make this idea more precise, we need to introduce some terminology.

Each piece **type**, except the square, has two or four different orientations, called **piece tokens**. An L type, for instance, has four tokens: . Tokens are structures in the world. The mental image corresponding to a token is an **icon**; and the time required to create an icon is the **iconifying period**. We assume the iconifying period lasts 50-80 ms, the time required to flood V1, primary visual cortex, with a retinal image (Hillyard, 1985). The pieces already sitting on the board have an upper boundary called the **current contour**. The process of comparing an icon to small regions of the current contour we call running an **iconic mask** over the contour and **envisioning** a placement. The measure of how snugly a placement fits into its neighboring pieces is called its **local fitness**. On the basis of experiments with Robotetris, see figure 5, we have discovered that aiming to maximize the local fitness of placements is an important factor in player longevity.¹

With these terms in mind we state two different methods (with variants) for determining placement, and consider how permitting epistemic actions can reduce their space-time complexity, and probability of error.

¹In Robotetris the decision concerning where to place a piece is determined by a judicious weighting of such features as "how many holes would this placement create", "how many rows would this placement eat", "how flat is the resulting contour", in addition to "how close to the globally maximum fitness is this placement".




Figure 4: Mental Rotation Task

In this figure we display the findings of PD's mental rotation tests for L's. The horizontal line marks the mean times for successful recognitions; the vertical lines mark the 90% confidence interval. The region within this interval we call the identification phase. We found that the identification phase for an L under its various rotations was nearly linear, as suggested by Tarr & Pinker (1989). More precisely, given our data at this stage we can only report that we haven't disconfirmed the constant time hypothesis.

Type-based Method

1. Iconify token;
2. identify piece type, automatically creating a stack of appropriate icons;
3. (a) *computationally intensive version*:
for each icon in the stack
 - i. run its corresponding mask over the current contour,
 - ii. envision the result of a placement,
 - iii. compute that envisioned placement's local fitness, and
 - iv. store the information about (place, score) in a list;
 or
 (b) *memory intensive version*:
for each icon in the stack
 - i. run an orientation-neutral mask over the current contour,
 - ii. look-up in an associative list the best orientation for each point on the current contour, and
 - iii. store the information about (place, score, orientation) in a list;
4. choose the placement that best maximizes local fitness as well as certain other weighted features.

Token-based Method

1. Iconify token;
2. create icon mask;
3. same as the steps in 3a and 3b above (i.e., without iteration);
4. generate a new icon by

(a) *physically* rotating the current token (go to 1); or

(b) *mentally* rotating current icon, (go to 2)

5. choose the placement that best maximizes local fitness as well as certain other weighted features.

If the token-based method resembles the human process of selecting placements, physical rotation is likely to be valued as a means of reducing both the time and effort of mental computation occurring in step 4. Pieces can be physically rotated in less than 100 ms whereas we estimate that mental rotation takes in the neighborhood of 800 to 1200 ms, based on pilot data, such as that displayed in figure 4. This may be misleading if we assume that because of priming effects, second and subsequent rotations are faster than first rotations.

If the type-based method resembles the human process, on the other hand, physical rotation is not especially helpful in enhancing the speed of computing local fitness. We assume that once a piece has been correctly identified, one may have access to its shape under all rotations, since it may be stored in this multiple perspective form. If physical rotation is useful in this type-based method, it will be because it abbreviates the time needed for Step 1: identify type of piece. For example, suppose it takes 1200 ms to identify a piece type from a presentation of a single token, whereas it takes 1000 ms to identify a type if shown one token for 600 ms immediately followed by another token for 400 ms. In such cases, it seems natural to conclude that rapid presentation of multiple perspectives of a piece stimulate retrieval of all perspectives faster than presentation of a single perspective.




Figure 5: Global Fitness

In this figure we compare the tendency of human players and Robotetris to choose locations on the basis of global fitness. As the results show, human subjects vary in how strongly they weight global fitness. Intermediates (mean score 36 rows) place pieces in the globally fit place 41% of the time; experts (mean score 93 rows) 39%. In the version of Robotetris considered here (moderate performance with mean score of 876 rows) 41% of placements are globally fit, lumping Robotetris with intermediates in global fitness, though far above experts in performance.

Conclusion

We have argued that standard state transforming actions are, at times, best understood as serving an epistemic rather than a pragmatic purpose. The point of a particular action may seem to be that of bringing an agent physically closer to its goals, yet upon more careful analysis the real point of that action may be to increase the reliability of a judgement, or to reduce the space-time resources needed to compute it. Most thoughtful theorists of action now agree that a natural part of planning and acting is gathering information. Characteristically, however, this has been interpreted to mean that planners should have an active hand in controlling sensor actions. The thrust of our account of epistemic actions in the game of Tetris is that the scope of epistemic activity is much wider than sensor related activity. Verification and experimentation are the simplest of epistemic functions. There are countless others in every natural form of intelligent activity. It is axiomatic that adaptive creatures would strike on such strategies for augmenting their cognitive abilities.

References

Agre, P. & Chapman, D. (1987). Pengi: An implementation of a theory of activity. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 196–201.

Georgeff, M. & Lansky, A. (1987). Reactive reasoning and planning. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 677–682.

Hillyard, S. (1985). Electrophysiology of human selective attention. *Trends in Neuroscience*, 8, 400–405.

Hutchins, E. (1990). The technology of team navigation. In J. Galegher, R. Kraut, & C. Egido (Eds.), *Intellectual teamwork: Social and technical bases of collaborative work*. Hillsdale, NJ: Lawrence Erlbaum.

Simmons, R., Ballard, D., Dean, T. & Firby, J. (Eds.). (1992, March). *Control of selective perception*. AAAI Spring Symposium Series. Stanford University.

Tarr, M. & Pinker, S. (1989). Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology*, 21, 233–282.

Tate, A., Hendler, J., & Drummond, M. (1990). A review of AI planning techniques. In J. Allen, J. Hendler, & A. Tate (Eds.), *Readings in Planning*, pages 26–49. Morgan Kaufman.