On Distinguishing Epistemic from Pragmatic Action

David Kirsh and **Paul Maglio**
Department of Cognitive Science
University of California, San Diego

Address correspondence to:
David Kirsh
Department of Cognitive Science
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0515

Running Head: Epistemic and Pragmatic Action

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 than by performing computational actions in the head alone. We have found that some of the translations and rotations made by players of this video-game are best understood as actions that use the world to improve cognition. These actions are not used to implement a plan, or to implement a reaction; they are used to change the world in order to simplify the problem solving task. Thus, we distinguish pragmatic actions—actions performed to bring one physically
closer to a goal—from epistemic actions—actions performed to uncover i
formation that is hidden or hard to compute mentally.

To illustrate the need for epistemic actions, we first develop a
information-processing model of Tetris-cognition, and show t
explain performance data from human players of the game—
relax the assumption of fully sequential processing. Star
regard many actions taken by players because they app
superfluous. However, we describe many such actions th
by players that are far from superfluous, and that p
proving human performance. We argue that tradi
because they regard action as having a singl
By recognizing a second function of ac
explain many of the actions that a tra
argument is supported by numerous
outline how the new category of
theories of action more gener

ON DISTINGUISHING EPISTEMIC FROM PRAGMATIC ACTION

duce the general idea of an *epistemic action*, and
Tetris, a real-time, interactive video-game. Epistemic
ical actions that make mental computation easier, faster, or
able—are *external* actions that an agent performs to change its own
ational state.

The biased belief among students of behavior is that actions create physical states which physically advance one towards goals. Through practice, good design, or by planning, intelligent agents regularly bring about goal-relevant physical states quickly or cheaply. It is understandable, then, that studies of intelligent action typically focus on how an agent chooses physically useful actions. Yet, as we will show, not all actions performed by well-adapted agents are best understood as useful physical steps. At times, an agent ignores a physically advantageous action and chooses instead an action that
seems physically disadvantageous. When viewed from a perspective which in
cludes epistemic goals—for instance, simplifying mental computation—s
actions once again appear to be a cost-effective allocation of the ag
and effort.

The notion that external actions are often used to simplif
putation is commonplace in tasks involving the manipula
symbols. In algebra, geometry, and arithmetic, for inst
diate results—which could, in principle, be stor
recorded externally to reduce cognitive loads (
position (Lerdahl & Jackendoff, 1983), ma
and a host of expert activities too m
strably worse if agents rely on t
putational abilities without
research on representati
lights the need to u
structures (No
Less wid
simplifyi
clearl
h

the best way to interpret the actions is not as moves intended to improve board position, but rather as moves that simplify the player's problem-solving task.

More precisely, we use the term *epistemic action* to designate a physical action whose primary function is to improve cognition by:

1. reducing the memory involved in mental computation, i.e., space complexity;
2. reducing the number of steps involved in mental computation, i.e., time complexity;
3. reducing the probability of error of mental computation, i.e., unreliability.

Typical epistemic actions found in everyday activities have a
time-course than those found in Tetris. These include familia
saving actions such as reminding, e.g., placing a key in a shoe,
around a finger; time-saving actions, such as preparing
partially sorting nuts and bolts before beginning
reduce later search (a similar form of comple
under the rubric "amortized complexity
gathering activities such as explor
help decide where to camp for
Let us call actions wh
to its physical goal
actions. As sugge
& Drummond
on decisi
on prag
i

broadened to include perceptual as well as pragmatic actions (see for example, Simmons et al., 1992). However, these inquiries have tended to focus on the control of gaze (the orientation and resolution of a sensor), or on the control of attention (the selection of elements within an image for future processing, Chapman, 1989), 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.

One significant consequence of recognizing epistemic action as
gory of activity is that if we continue to view planning as state-
we must redefine the state-space in which planning occurs.
of interpreting the nodes of a state-space graph to be p
to interpret them as representing both *physical* a
this way, we can capture the fact that a se
same time, return the physical world to i
alter the player's informational st
who moves a piece to the left o
nal position, performs
game unchanged.
something or suc
time lost by the
us to cont
search

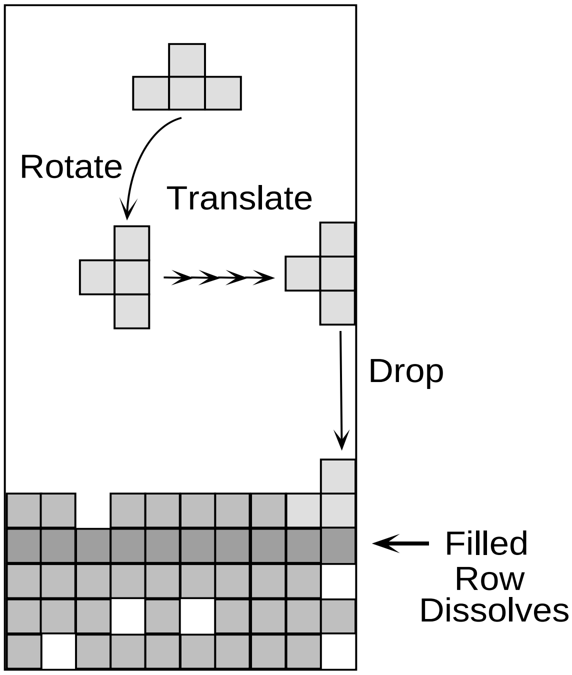

Figure 1. In Tetris, shapes, which we call zoids, fall one a time from the top of the screen, eventually landing on the bottom or on top of shapes that have already landed. As a shape falls, the player can rotate it, translate it to the right or left, or immediately drop it to the bottom. When a row of squares is filled all the way across the screen, it disappears and all rows above it drop down.

cause Tetris is fun to play, it is easy to find advanced subjects willing to play under observation, and easy to find novice subjects willing to practice until they become experts.

Playing Tetris involves maneuvering falling shapes into specific arrangements on the screen. There are seven different shapes, which we call **Tetrazoids**, or simply **zoids**: 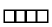, , 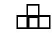, , , , . These zoids fall one at a time from the top of a screen that is 10 squares wide and 30 squares high (see Figure 2). Each zoid's free-fall continues until it lands on the bottom edge of the screen or on top of a zoid that has already landed. Once a zoid hits its resting place, another zoid begins falling from the top, startin

next Tetris episode. While a zoid is falling, the player can **rotate** it 90° counterclockwise with a single keystroke, or translate it to the **right** or to the **left** one square with a single keystroke. To gain points, the player must find ways of placing zoids so that they fill up rows. When a row fills up with squares all the way across the screen, it disappears and all the rows above it drop down. As more rows are filled, the game speeds up (from an initial free-fall rate of about 200 ms per square to a maximum of about 100 ms per square), and achieving good placements becomes increasingly difficult. As unfilled rows become buried under poorly placed zoids, the squares pile up, creating an uneven contour along the top of the fallen sqaures. The gam
ends when the screen becomes clogged with these incomplete rows, and ne
zoids cannot begin descending from the top.

In addition to the rotation and translation actions, the player
a falling zoid instantly to the bottom, effectively placing it in
would eventually come to rest in if no more keys were pr
an optional maneuver, and not all players use it. Dro
to speed up the pace of the game, creating shorter
the free-fall rate.

There are only four possible actions
right, translate left, rotate, and
so small, the game is not
newcomer can play at a
for experts, becaus
score, leaving
and execut
ceptual
mu

3. We have designed and implemented an expert system to play Tetris and have compared human and machine performance along a variety of dimensions.

In what follows, we use these data to argue that standard accounts of practiced activity are misleading simplifications of whatever processes actually underlie performance. For instance, standard accounts of skill acquisition explain enhanced performance as the result of chunking, caching, or compiling (Newell, 1990; Newell & Rosenbloom, 1981; Reason, 1990; Anderson, 1983). Although our data suggest that Tetris-playing is highly automated, we cannot properly understand the nature of this automaticity unless we see how
closely action is coupled to cognition. Agents do not simply cache associa
rules describing what to do in particular circumstances. If caching
source of improvement, efficiency would accrue from following
same cognitive strategy used before caching, only doing it fa
behavioral routines are compiled. If chunking were t
ment, efficiency would accrue from eliminating inte
sometimes to more far-reaching strategies, bu
basic style. Our observations, however, i
ferent behavioral tricks. Agents
to prime themselves to recog
checks or verifications t
such epistemic proce
cedures; they are
to make the m
To ma
proc

on a classical information-processing model of expertise that supposes Tetris-cognition proceeds in four major phases:

1. Create an early, bitmap representation of selected features of the current situation.
2. Encode the bitmap representation in a more compact, chunked, symbolic representation.
3. Compute the best place to put the zoid.
4. Compute the trajectory of moves to achieve the goal placement.

Figure 2 graphically depicts this model.

Phase One: Create Bitmap

Light caused by the visual display strikes the retinal cortex and initiates e
visual processing. Elaborate parallel neural computation extract
dependent features and represents them in a brief sensory r
called an iconic buffer (Sperling, 1960; Neisser, 1967).
iconic buffer are similar to maps, in which importan
as contours, corners, colors, etc., are present
That is, the memory regions that carry
segments are not *labelled* by symbol
of line segment, or any other at
Rather, such information is
to encode it in an ex

Pha

By attending to su
tracted and e
of RoboTetri
cludes

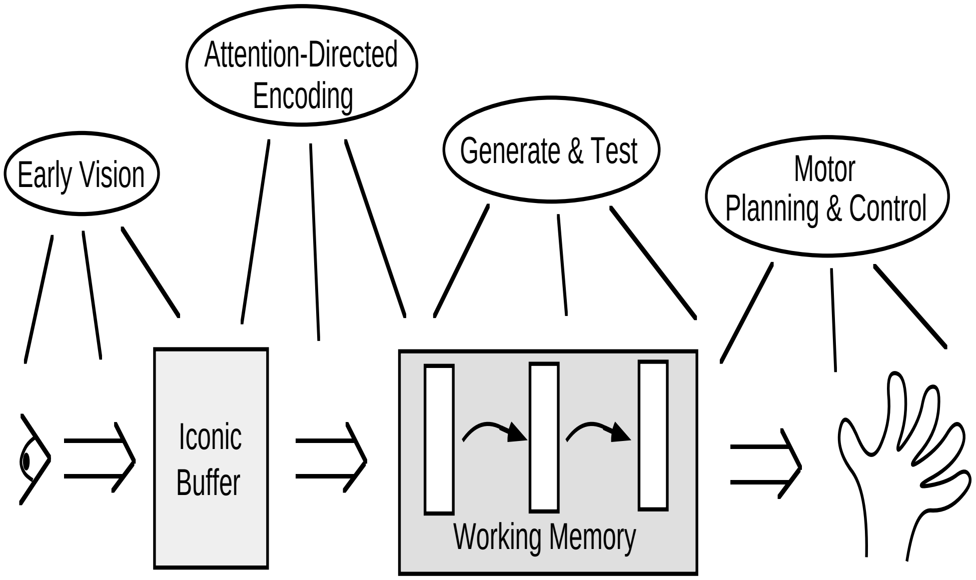

Figure 2. In our classical information-processing model of Tetris-cognition, first
a bitmap-like representation floods the iconic buffer, then attention selectively
examines this map to encode zoid and contour chunks. These chunks accumulate
in working memory, providing the basis for an internal search for the best place
to put the zoid. This search can be viewed as a process of generating and
luating possible placements. Once a placement has been chosen, a motor
r reaching the target is computed. The plan is then handed off to a motor
or regulating muscle movement.

concave corners, convex corners, and **T-junctions** (see Figure 3). Such
presentation has advantages, but our argument does not rely critically
on this choice. Another set of symbolic features might serve just as well,
provided that it too can be computed from pop-out features—such as line
segments, intersections, and shading (or color)—by selectively directing attention to conjunctions of these (Treisman & Souther, 1985), and that it
facilitates the matching process of Phase Three.

As yet, we do not know if skilled players encode symbolic features more
quickly in working memory than less skilled players. Such a question is
worth asking, but regardless of the answer, we expect that absolute speed of
symbolic encoding is a less significant determinant of performance than the
size of the chunks encoded. *Chunks* are organized or structured collectio
of features which regularly recur in play. They can be treated as lab
rapidly retrievable clusters of features which better players us
both zoids and contours (see Figure 4). As in classical
that much of expertise consists in refining selectiv
larger chunks of features to be recognized rap

Given the importance of chunking, a key
language—one provably satisfied by our li
it is expressive enough to uniquel
and to allow easy expression
determining whether a part
contour (see Figure 5).

Ph

Once zoid and contour a
be compared in worki
which to place the
ways this matc
search for t
and to
of

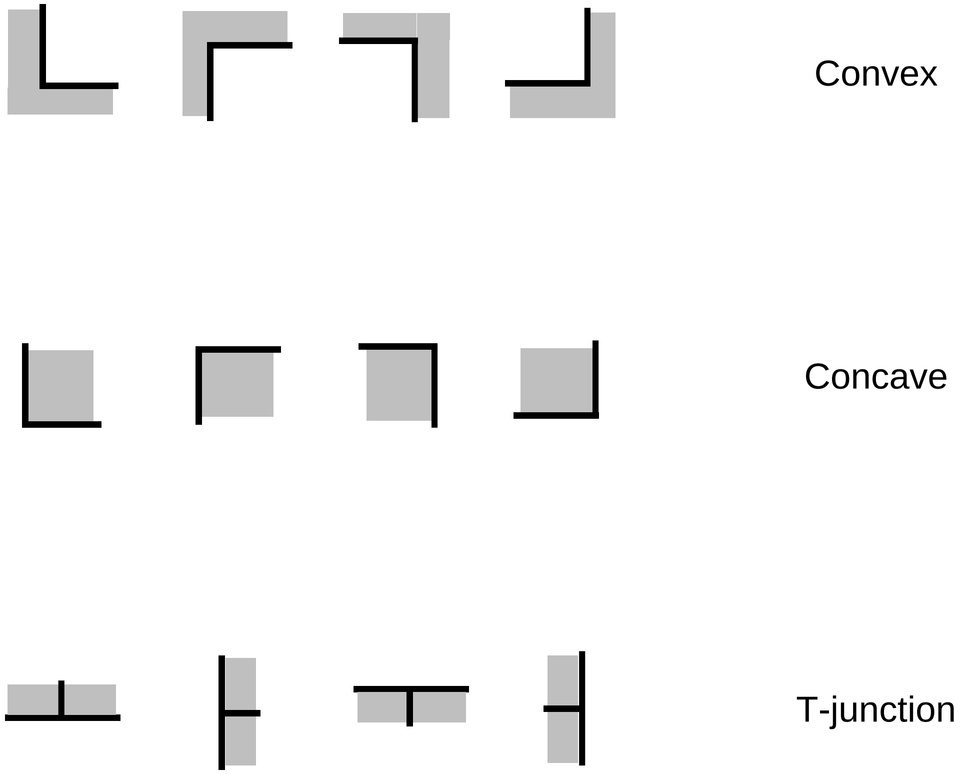

Figure 3. Three general features—concave, convex, T-junction—in each of their orientations create twelve distinct, orientation-sensitive features. These features are extracted by selectively attending to conjuctions of the more primitive features: lines, intersections, and shading.

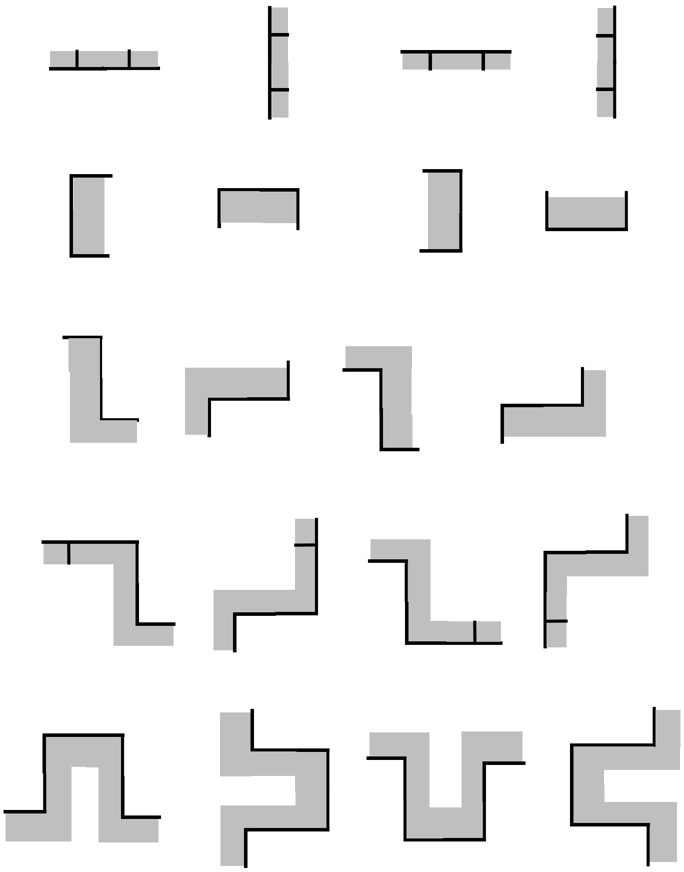

Figure 4. The greater a player's expertise, the more skilled the perception. This is reflected by the size and type of the chunked features which attention-directed processes are able to extract from iconic memory. This figure shows chunks of different sizes and types. Each chunk is a structured collection of primitive features.

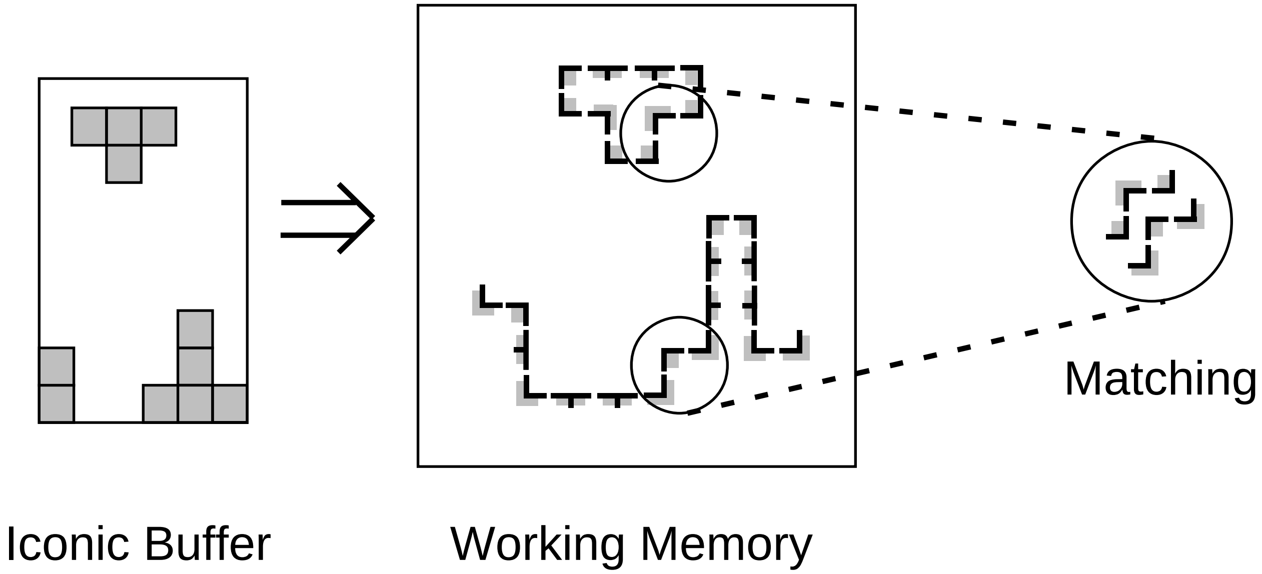

Figure 5. A good representation must make it easy to recognize when zoid and contour fragments match. In this figure, a zoid chunk matches a contour chunk when concave corners match convex corners and straight edges match straight . This simple complimentarity is probably computed in the visuo-spatial of working memory (Baddeley, 1990).

placement involves matching chunks to generate candidate locations. To test the candidates, actual placements are simulated in an internal model of the Tetris situation.

Phase Four: Compute Motor Plan

Once a target placement is determined, it is possible to compute a sequence of actions (or equivalently, keystrokes) that will maneuver the zoid from its current orientation and position to its final orientation and position. The generation of this motor plan occurs in Phase Four. We assume that such a motor plan will be minimal in that it specifies just those rotations and translations necessary to appropriately orient and place the zoid.

After Phase Four, RoboTetris carries out the motor plan by c
affecting the ongoing Tetris game, effectively hitting a sequenc
take the planned action.

This completes our brief account of how a classical inf
theorist might try to explain human performance, and
RoboTetris on these principles.

How Realistic is this Model?

As we have stated it, the model is fully sequential: Phase T
before Phase Three begins, and Three is completed befor
cause all processing within Phase Four must also be
ecution begins, the muscle control system cannot
movements until a complete plan has been forn
occuring before the processing of Phase Four
unplanned; they cannot be under *rational* c
to be no better than random actions.

This is patently not what we see
occur in abundance, almost from
If players actually wait until
number of rotations shou
be performed on the z
each zoid emerge
be expected t
Thus, a

rotated three times before repeating an orientation, ought to average out to 1.5 rotations. As can be seen in Figure 6, each zoid is rotated more than half its possible rotations. And as Figure 7 shows, rotations sometimes begin extremely early, well before an agent could finish thinking about where to place the zoid.

If we wish to save the model within the classical information-processing framework, one obvious step is to allow Phase Four to *overlap* with Phase
Three. Instead of viewing Tetris-cognition as proceeding serially, we can vi
it as a *cascading* process in which each phase begins its processing be
has been given all the information it will eventually receive. In
agent will regularly move zoids before completing delibera
way to capture this notion is to suppose that Phase Thr
Phase Four with its *best estimate* of the final choic
computing a path to that spot and the agent ini
Phase Four produces its first step.

In the AI planning literature, the
terleaving (Ambros-Ingerson & Steel
ecution before they have settl
orthodox planner execute
subgoals, and hence
executes its fi
has bui
is

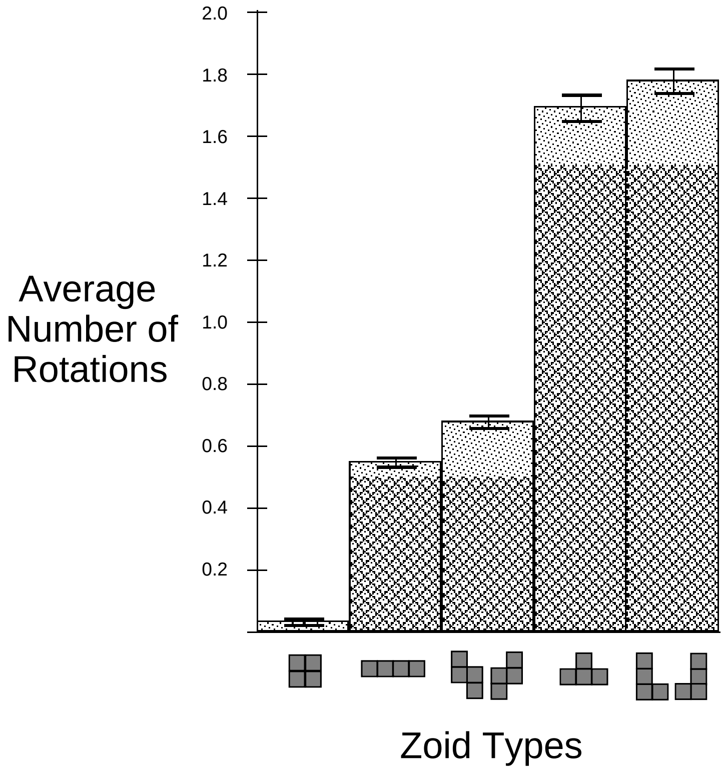

Figure 6. This bar graph shows the average number of rotations for each type of zoid from the moment it emerged to the moment it settled into place. Zoids such as are rotated significantly more than , and both types are rotated more an the expected number of rotations, shown by the crosshatched portions of the Similarly, zoids such as are rotated more than , and both exceed the umber required for purely pragmatic reasons. The error bars indicate intervals.

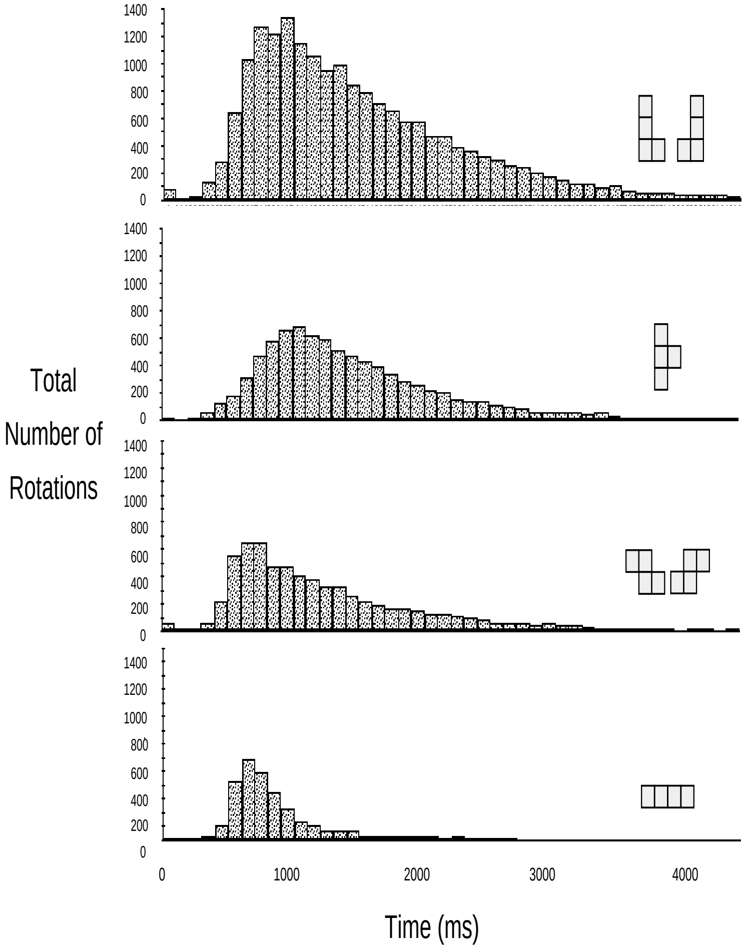

Figure 7. These histograms show the time-course of rotations for ⊟'s, ⊢'s, ⊞'s, and □□□□'s. Each bin contains the total number of rotations performed within its time-window. Note that rotation begins in earnest by 400-600 ms, and on occasion, at the very outset of an episode. The implication is that planning cannot be completed before rotation begins.

then, presumably, he or she ought to start out early toward that location and make corrections to zoid orientation as plan revisions are formulated. Early execution, on average, ought to save time.

In theory, such an account is plausible. That is, we would expect to find extra rotations in interleaving planners because the earlier an estimate is made, the greater the chance it will be wrong, and hence the more likely the agent will make a *false start*.

In fact, however, given the time course and frequency of rotations we observe in Tetris, particularly among skilled players, an explanation i
of false starts makes no sense. First, the theory does not explain
might start executing before having *any* estimate of the final
a zoid. We have observed that occasionally a zoid will b
(before 100 ms), well before we would expect an a
idea of where to place the zoid. This is particul
ms, the zoid is not yet completely in view,
even reliably guess the zoid's shape.[2] Si
is hardly reasonable that Phase Four
ought to act on.

Second, there is a significan
has reasonable grounds for a
target orientation, the ag
more times, dependi
depend on how lon
between keyst
keystrokes
is ne

a well-adapted agent.

In our view, the failure of classical and interleaving planners to explain
the data of extra rotations is a direct consequence of the assumption that
the point of action is always pragmatic: that the only reason to act is for
advancement in the physical world. This creates an undesirable separation
between action and cognition. If one's theory of the agent assumes that
thinking precedes action, and that, at best, action can lead one to re-evalua
one's conclusions, then action can never be undertaken *in order to alt*
way cognition proceeds. The actions controlled by Phase Four can n
for the sake of improving the decision-making occurring in Phas
for improving the representation being constructed in Phase
view, cognition is logically prior: cognition is necessar
but action is never necessary for intelligent c

To correct this one-sided view, we ne
of an action is to put one in a better
more quickly identify the curren
information; to more effe
action of rotating
easier to compu
suddenly n
to perf
the

Pragmati
th

4. make it easier to identify a zoid's type,
5. simplify the process of matching zoid and contour.

Each of these epistemic actions serves to reduce the space, time, or unreliability of the computations occurring in one or another phase of Tetris-cognition. We are not claiming, however, that every player exploits the full epistemic
potential of rotation. From a methodological standpoint, it is often hard to
prove that an agent performs a particular action for epistemic rather
for pragmatic reasons because an action can serve both epistemic a
matic purposes simultaneously. Rotating a zoid in the directio
final placement may also help the player identify the zoid.
makes it difficult to quantify the relative influence of epi
functions. Nonetheless, the two functions are logi
there are clear cases in which the only pl
choice of action is epistemic.

Early Rotations for

When a zoid first enters at the top of the scre
form is visible. At medium speed, a zoi
every 150 ms. Therefore, it takes al
instance, to emerge. It is c
complete shape as soo
zoid is consisten
image could d
Given
a stra
s

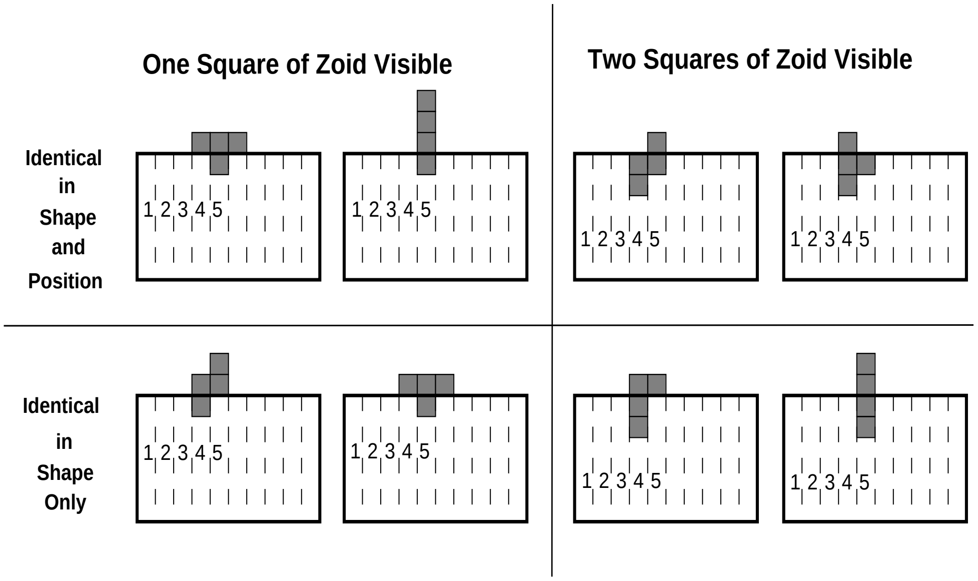

Figure 8. This figure shows zoids as they first emerge at the top of the screen. To the left, they are one square in, and to the right, two squares in. At the top, the visible portions of the zoids are identical both in position and in shape. At the bottom, zoids are identical in shape alone; careful examination reveals that the images are in different columns. Players have a much greater tendency to otate partially hidden zoids ambiguous in both shape and position than they of rotating partially hidden zoids that are ambiguous in shape alone.

ape and position produces an early image such that no matter how player knows, it is impossible to tell which zoid is present solely on he basis of the early image.

Our data show that a player is more likely to rotate a partially hidden zoid that is ambiguous in both shape and position than one ambiguous in shape alone. Partially hidden zoids ambiguous in shape only are *not* rotated more than completely unambiguous ones. This suggests that players are sensitive to information about column because, in principle, zoids ambiguous in shape alone are distinguishable by column. Hence early rotation would add no new information. Yet, when interviewed, no player reported noticing that zoids begin falling in different columns. Thus, although players are sensitive to column, and are more likely to rotate in those cases where it is truly informative to do so, they do not realize they have this knowledge.

Early rotation is a clear example of an epistemic action. Nonetheless, one might try arguing against this view by suggesting that there is pragmatic value in orienting the zoid early, and so its epistemic function is not decisive.
Such an explanation, however, fails to explain why partial displays that
ambiguous in shape and position are rotated more often those that a
ambiguous shape and position. Nor would such an explanation make
we believe that an agent has yet to formulate a target *orientation*
this early stage. It is certainly possible that a player begi
set of target spots on the board where he or she would
zoid. Some players do report having hot spots
begins. And some of these players do *translat*
that whatever shape emerges, they are like
But such early intentions explain ear
does not know the shape of a z
in the right orientation
that the point of
later, and
potent
P

In Phase Thre
a useful repr
tatio
fin

Method One: The player identifies the type of the zoid before looking for
possible placements, using knowledge of all orientations to search for snug
fits. This means that the player extracts an abstract, *orientation-independent*
description of the shape, or chunk, before checking for good placements.
Method Two: The player does not bother to compute an orientation-
independent representation of the zoid or chunk. Leaving the representation
in its orientation-sensitive form, the player redirects attention to the contour,
looking for possible matches with the orientation-specific chunk. In this
second method, contour checking can begin earlier than in the first method,
but to be complete, the process of contour checking must be repeated fo
the same zoid or chunk in all its different orientations. Needless to s
may discover players who use some of each method, possibly with
running concurrently.
When we look more closely at these methods, we see seve
where epistemic actions would be useful.
Consider method two first. Somehow a player must compar
of a zoid in all its possible orientations to fragments o
do this, the player may compare the zoid in its curr
contour, then use *mental imagery* to recreate how
rotated (see Figure 9).[3] Another possibility—f
time—is that the player may rotate the zoi
orientation-specific comparison.

The clearest reason to doubt that
mental rotation is that zoids can
ms, whereas we estimate that
ms to mentally rotate a zoi
in Figure 10.[4] We obta

[3] Possibly, the play
reasoning to ju
is made c
l

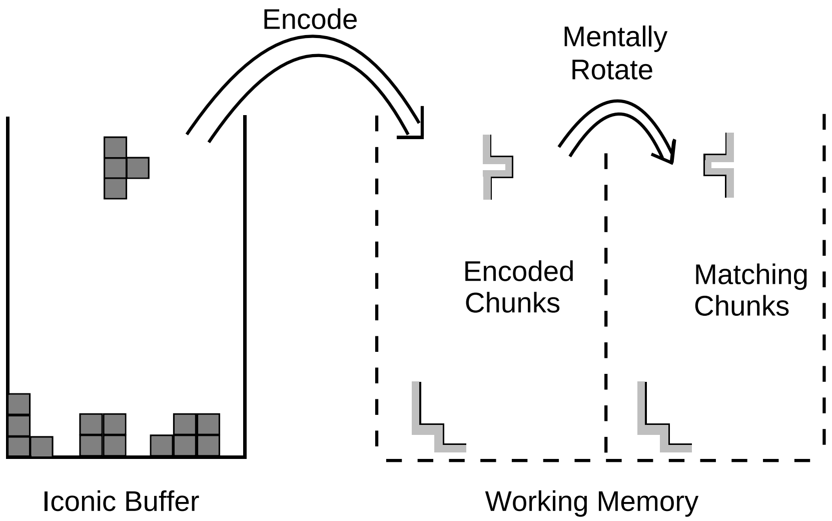

Figure 9. A chunk extracted from the image of a zoid is normalized by internal processes and compared to a chunk extracted from the image of a contour. A computationally less intensive technique of comparing zoid and contour would y on physical rotation of the zoid to take the place of the internal normalization esses.

similar to the one used by Shepard and Metzler (1971). In our experiment, two zoids, either S-shaped (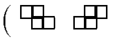) or L-shaped (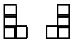), were displayed side-by-side on a computer screen. The zoids in these pairs could differ in orientation as well as handedness, but in all cases, both items were of the same type. To indicate whether the two zoids matched or whether they were mirror images, subjects pressed one of two buttons. Three Tetris players participated: one intermediate, one advanced, and one expert. Each subject saw eight presentations of each possible pair of zoids. The results, as graphe
in Figure 10, show reaction time as an increasing function of the ar
difference between the orientations of the two zoids (from 0° to 18

Even allowing an extra 200 ms for subjects to select the r
the time saving benefits of physical over mental rotation
time is not all that is saved. There are also costs asso
and memory needed to create and sustain mental i
instance, suppose that matching proceeds by c
zoid with chunks of the contour. Even if ch
faster than we expect, there are still sig
a record of the chunks that have al
test process requires repeatedly
new chunk to check. The net r
would soon fill up with a)
is the target for mat
and d) a marker
from It seems
the extra step
compare
accou

In method one, play
tion of the zoi
pay the pro

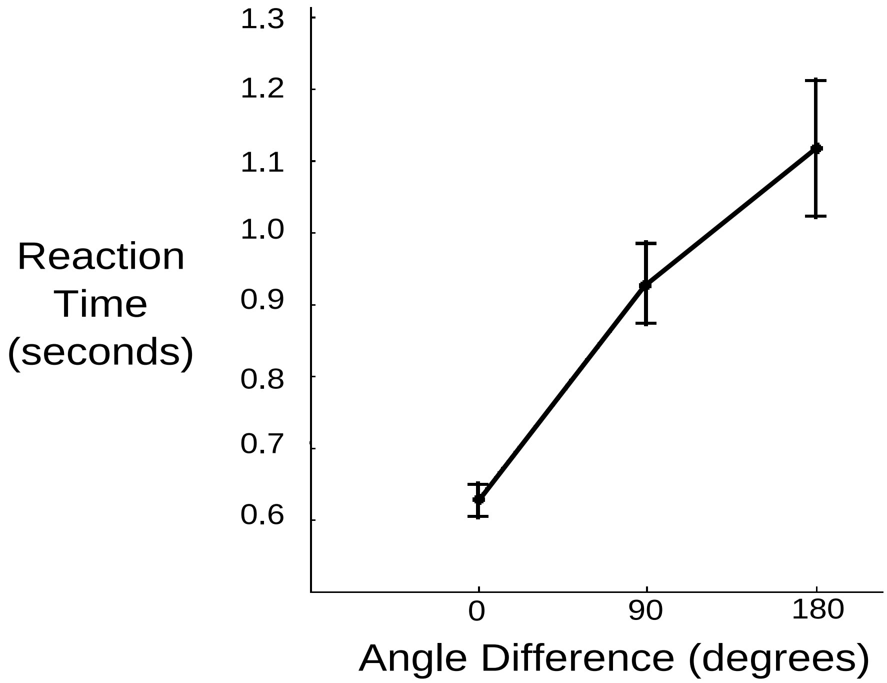

Figure 10. This graph shows the results of a pilot study on the mental rotation of Tetris shapes by players of differing skill levels. Reaction time (in seconds) is plotted against difference in orientation of two displayed L-shaped zoids (only differences from 0° to 180° are plotted). Only correct "same zoid" answers are included; i.e., conditions in which both zoids were either of type or of type . A linear relationship between reaction time and angle-difference is readily apparent. The error bars represent 95% confidence intervals.

once they have an orientation-independent representation of a zoid, it is not necessary to rotate the zoid further to test for matches. Nonetheless, external rotation is still epistemically useful because it is helpful in constructing orientation-independent representations in the first place.

What does it mean to have an orientation-independent representation?
From an experimental perspective, it means that it should take no mor
time to judge whether two shapes are the same, however many degrees apar
the two have been rotated. Players' reaction times on mental rotation
should be plotted as a horizontal line, rather than the upward
line we see in Figure 10. Total reaction time should be t
time needed to abstractly encode the first shape (present
abstractly encode the second shape (presentation), a
the abstract encodings. Moreover, we would expec
abstractly encode different presentations, and t
encodings should be constant across all tri

We have not observed flat line per
of very experienced players, so we m
players use abstract orientatio
studies of extremely prac
fact, the more expos
closer to flat
explanation
acquire
Te

perspective representations, external rotation could play a valuable role in speeding up the multiple-perspective encoding process. Consider what it means, from a computational perspective, to activate (or encode) a multiple-perspective representation. Presumably, the agent enters a state in which the complete set of orientation specific representations are active, or at least, strongly primed. The process by which this activation takes place is identi
to retrieval. Thus, each image of a shape serves as an index, or retr
for the multiple-perspective representation.

How might physical rotation help such a retrieval pro
ture, which is ripe for experimental testing, is that r
environmental support there is (Park & Shaw,
ulate that it takes less time to complete a
complete a retrieval using *n* indices
subject a total of 1200 ms to i
when shown a single token
to identify the type if
ms immediately fol
perspectives of
single pe
In
mac
ta

the useful function of speeding up the activation process. In this case, *two cues are better than one.* Because rotation is the means of generating the second cue, and rotation is quick enough to save time in the settling process, it can play an epistemically valuable role.

Rotating to Help Identify Zoids

It is an open question whether agents use multiple perspective representations of zoids (or chunks). It is not an open question whether there is a phase where zoids are first represented in their current perspective as particular zoid shapes (or chunks of zoids). On our account, the process by which particular zoids are encoded in working memory has three logical steps. In
the first, simple features such as lines, corners, and colors are extract
the image; in the second, orientation-specific corners and lines-
features of the image—are extracted; and in the third step, s
conjuntive features—perceptual chunks—are identified
itly in working memory. Both steps two and three re
reasonable to suppose, then, that fast perceptu
highly trained attentional system, and th
due to improvement in the attentional
recognition. Thus, we hypothesize t
chunks and zoids, it is because
features represented in the
We can recast this
say that the more expe
at searching for the
chunks. Accor
is in terms
by me

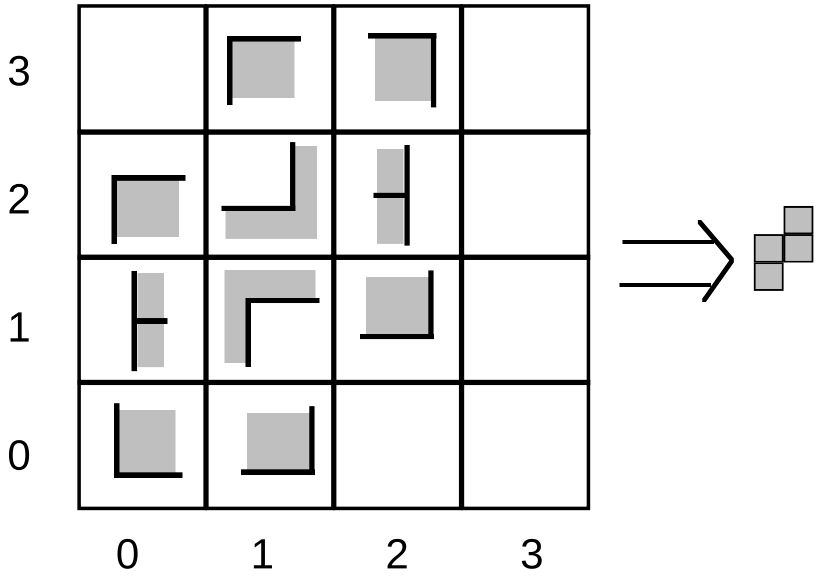

Figure 11. The iconic buffer is a 4 × 4 matrix of cells, each of which may contain a primitive feature.

minimal number of cells to reliably extrapolate to the contents of the whole matrix (see Figure 11).

Given the shape of tetrazoids, experts may sometimes rotate zoids because, if encoding operates by a mechanism at all like a decision tree, then *rotating can be an effective way of reducing the number of attentional probes needed to identify a zoid.* Compare Figures 12 and 13. The decision tree in Figure 12 assumes the expert identifies the zoid without rotating it. As can be seen, if the expert first examines cell (1,1), then, a decision will require
either one, two, or three questions directed at the matrix to identify the zoid,
depending, of course, on the zoid present and the contents of (1, 1). Th
cision tree in Figure 13, however, shows that if the agent can also
zoid between its attentional probes of the matrix, an ident
made in at most two questions. Thus, rotation can be
the program controlling attention. An expert can o
decision-tree if rotation is included in the set o

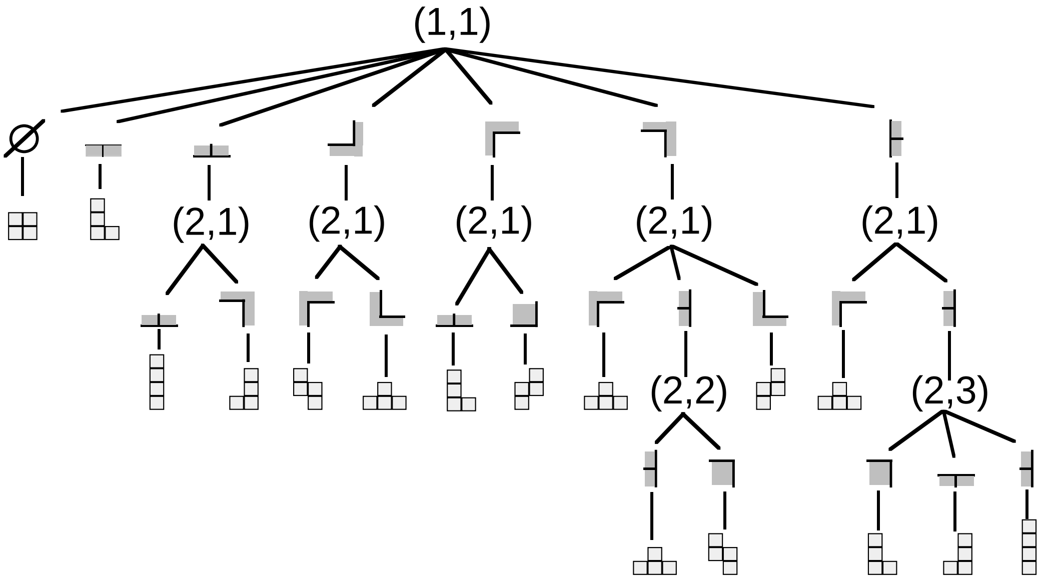

Figure 12. This decision-tree directs a series of questions at specific cells in the iconic buffer in order to identify what type of zoid is present. The tree first probes cell (1,1). If the buffer is the one in Figure 11, cell (2,1) is queried next, leading to the identification of .

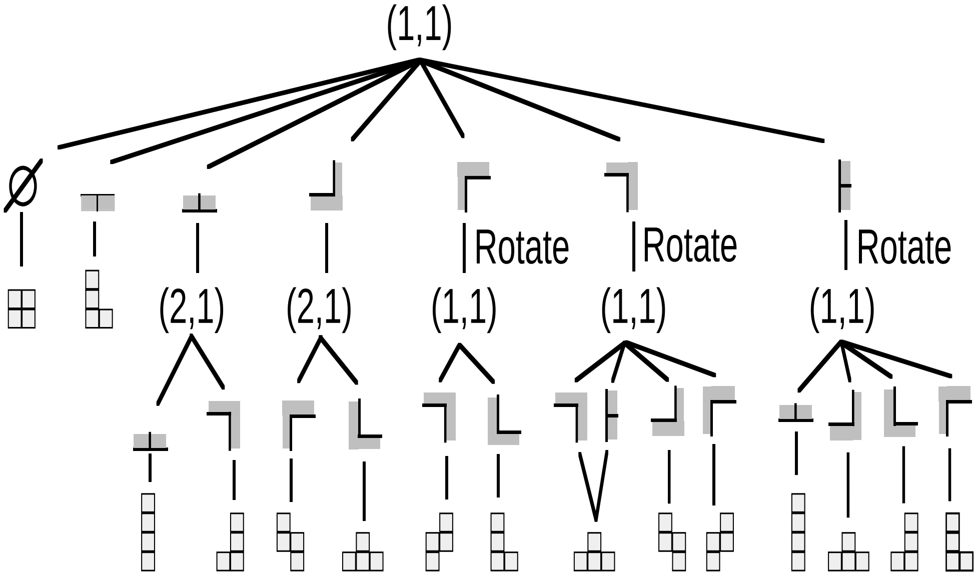

Figure 13. If the decision-tree incorporates calls to external rotation operations, its maximum depth is two. In addition, attention need not shift from cell (1,1) most of the time.

But this may be only part of the story. So far, we have argued that identification involves domain-specific control of attention, and that extra rotations may be a side effect of a streamlined program regulating this control. A second reason experts may make superfluous rotations is that, paradoxically, it is the lazy thing to do. Although we do not know if it takes less energy on the part of an attention mechanism to consult the same cell twice, it is possible that a lazy attention mechanism might prefer to re-ask for the value of a cell, rather than focus on a new cell. This is an obvious strate
when new data has just arrived because change is automatically inter
to the nervous system. This idea of finding a strategy that min
number of cells probed makes sense in a decision-tree account o
long as it costs less to consult the same cell on successi
case, the decision-tree in Figure 13 would be pref
in Figure 12 because probing the same cell
would put less strain on the attentiona

The implication of both arguments, we believe, is that it is adaptive to build attentional mechanisms that are closely coupled with actions such as rotation. The close coupling between attention and saccades is already accepted, why not extend this coupling to include more molar actions such as rotation?

Rotating to Facilitate Matching

So far we have assumed that matching is a primitive process in working memory: zoid chunk and contour chunk can be compared and matched only if they are explicitly represented in working memory. To make certain that enough chunks of different sizes are tested to guarantee finding the largest matching chunks, a player can rely on either externally rotating a zoid, mentally rotating a zoid, or mentally accessing a multiple-perspective representation of a zoid to generate as many candidate chunks as time will allow.

Are we justified in assuming that matching occurs in working memory?
And that symbolic matching, primitive or not, is really the fastest way
determining a fit between a zoid fragment and a contour fragment?

An alternative possibility is that matching is a perceptual proc
general idea is simple enough. Matching requires noting the c
two structures. If the structures are simple—such as lines
in the same orientation—it may be possible to note th
ing some attention-directed process such as a vis
applied directly to the early bitmap-like rep
ing might actually be an element of Pha
features of the situation are extract
of Phase Three—the phase in which
working memory.

External rotation plays a r
we have to explain how new
Because we are considerin
there muct also be a m
only certain way to
buffer is through
tions may
whethe
rese

for instance. Second, if mental rotation does modify the pre-attentive iconic buffer—where the bitmaps reside—players would probably prefer to create the relevant bitmaps by external rotation rather than by mental rotation because, as mentioned earlier, external rotation is faster. And third, it is likely that physical rotation is less cognitively demanding than mental rotation. Iconic memory needs to be refreshed every 200 ms (Reeves & Sperling, 1986). Thus, if a player uses mental imagery to flood the iconic buffer, he or
she will have to refresh the buffer every 200 ms. It is much easier to gene
tokens by bringing them in through the visual system than by interna
ating them. Therefore, even if matching operates by perceptua
correspondence, we have another reason for preferring exter
to mental rotation and to multiple-perspective repres

So ends our account of the epistemic uses
our discussion of the data with a brief descr
translation.

TRANSLATION AS AN EPISTEMIC ACTION

The pragmatic function of translation is to shift a zoid either r
to permit placement in an arbitrary column. Translation us
pragmatic purpose. But we have found at least one unamb
use of translation: to verify judgement of the column
of the cases when a player drops a zoid, the act
a behavioral routine of translating the zoid
See Figure 14. Because the accuracy of j
visually presented stimuli varies w
coeur, Ullman & Mackay, 1991
has a greater chance of la
from a height of three sq
to-wall routine is
zoid to the wa
column, a pl

As ep
action
in

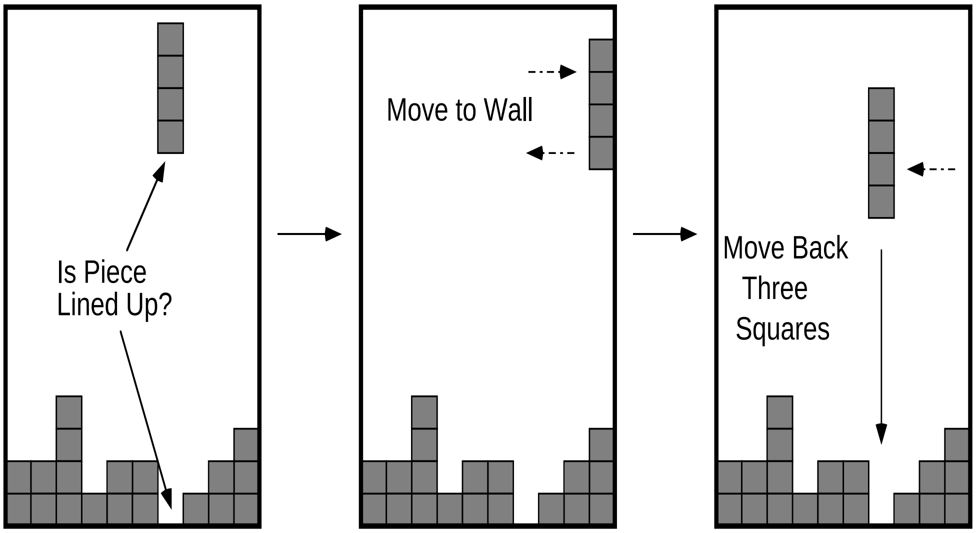

Figure 14. In a small percentage of cases players will drop certain zoids only after translating them to the nearest wall and then back again, as if to verify the column of placement. In this figure, ▯ is translated to the outer wall and back again before it is dropped. The explanation we prefer is that 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 simultaneously counting tapping out the number of squares to the intended column.

Table 1

Ordinary Drop Distance vs. Translate-to-Wall-then-Drop Distance

	Intermediate	Advanced	Expert
Mean Drop Distance	13.18	13.69	15.65
Mean Drop Distance after Translate Routine	19.04	19.33	20.05

Note. Within each skill level, the two means differ significantly as judged by a t test with $\alpha = .05$.

over, it cannot sensibly be viewed as a mistaken pragmatic action because the
procedure is more likely to occur the higher the drop. As shown in Table 1,
experts drop a zoid, on average, when it is about 13 squares from its resting
position. On those occasions when they also perform the translate-to-wall
routine, the zoid is dropped, on average, from about 19 squares above its
resting position, 6 squares higher than usual. The only reasonable account
for this regularity is that the higher the zoid, the more the player needs t
verify the column. Moreover, as shown in Figure 15, the greater the
distance, the more likely the drop is verified using the translate-t
tine. At great heights above the zoid's resting position,
of moving away from the goal column is more than offse
benefit of reducing possible error.

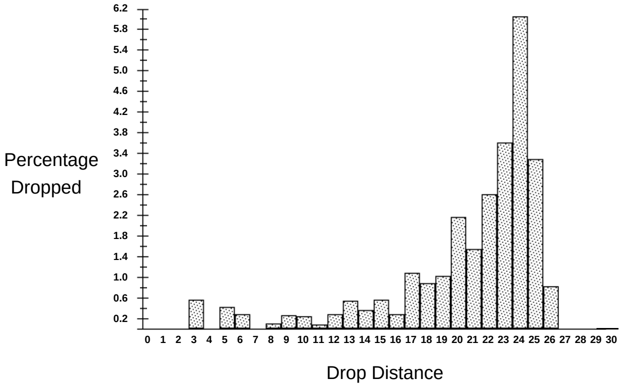

Figure 15. This graph plots the percentage of dropped zoids that followed a translate-to-wall routine against the distance they were dropped. The higher the drop, the more likely it followed a verification routine.

DISCUSSION

xplain our data on the timing and frequency of rotations and transla-
ions regularly performed by Tetris players, we have argued it is necessary to advert to a new category of action: epistemic actions. Such actions are not performed to advance a player to a better state in the external task environment, but rather to advance the player to a better state in his or her internal, cognitive environment. Epistemic actions are actions designed to change the input to an agent's information-processing system. They are ways an agent has of modifying the external environment to provide crucial bits of information just when they are needed most.

The processing model this suggests to us is a significant departure from
classical theories of action. Its chief novelty lies in allowing individual
tional units inside the agent to be in closed-loop interaction wit
world. Figure 16 graphically depicts this tighter coupli
and external processes. As in the cascade model ment
cessing starts in each phase before it is compl
in this case, the output of Phase Two can
Four, activating a motor response directly
of Phase Three can bypass Phase Four.

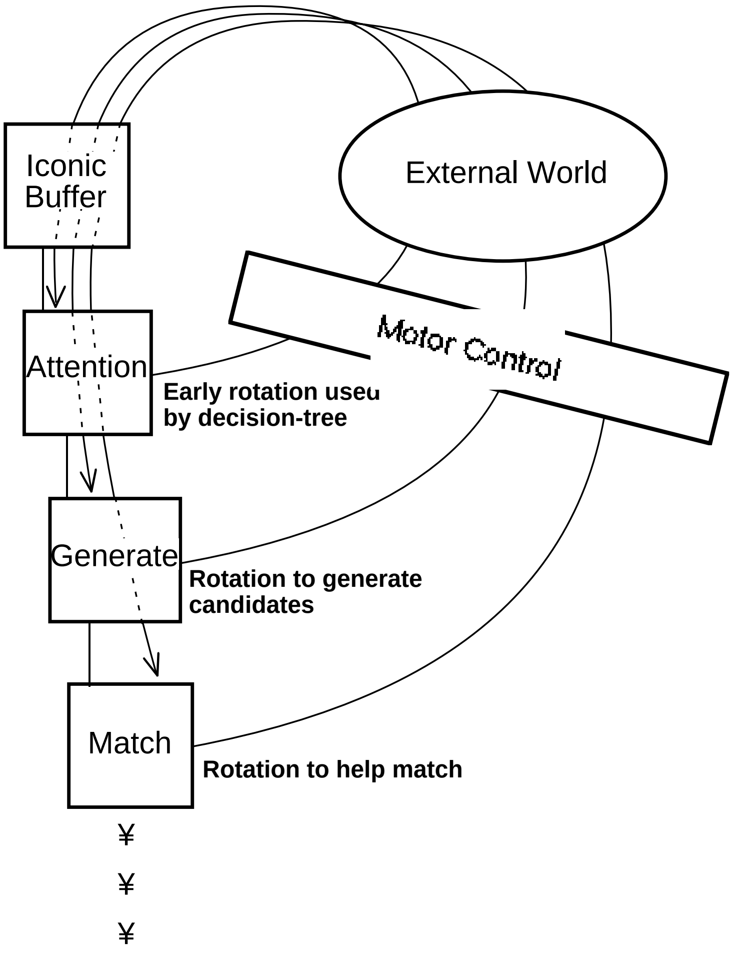

Figure 16. In this model, calls for rotation from attentional processes, or from candidate generation processes, cause changes in the world which feed back into those very processes. Because of the tight coupling between action and what is perceived, the fastest way to modify the informational state of an internal process may be to modify its next input.

To return to an example already discussed, suppose attention operates
as if driven by a decision-tree. The attentional system may request rotations
in the same way that it requests directing attention to cell (i, j) in the iconic
buffer. These requests are not sent to the Phase Three processes operating on
working memory, as if to be approved by a higher court. They are temporary,
time-critical requests which have no bearing on the pragmatic choice of where
to ultimately move. The point of the request is very specific: to cash in on
the speed at which input can be changed. If a change of input will hel
complete the computations that constitute selective attention faste
the attention system can compute on its own, it would be adaptiv
attention directly to certain simple motor actions.

The property of Tetris that makes such a strategy pa
local effects of an action are totally determinate. There
exogenous influences, or other agents to change
rotate key. There is a dependable and simple
the change in stimulus. Consequently, a wel
might incorporate simple *calls to* the wor

A similar story can be told
generated and matched or teste
can provide just the input r
a match. Again, becaus
the agent can cou
the rotate key
on to hel

One
proce
ce

tem responsible for saccades. Perhaps there is a similar connection between attention and highly trained key pressing responses.

Second, we can create a more complicated picture of the interrelations among processes involved in Tetris-playing than the one presented in Figure 2. Consider Figure 17, which displays a highly interconnected network of processes for attention, candidate generation, matching, and rotation. Obviously, this does not represent a strictly feedforward system: there are bac
ward links from **generate** candidates and **match** to **attention**, as well as from all three to **motor arbitrate**. We have already discussed how **match** and **rotate** can benefit from sending requests back to **attention**. In the same way, candidate generation can benefit from sending requests back to attention because the process of generating new candidate placements requires trying out new zoid chunks and new contour chunks, and an easy way to create such chunks is by looking at zoid and contour anew. The one complication this connection scheme adds to the process is that requests for motor actions must be arbitrated, hence the addition of the **motor arbitrate** process. This kind of model follows the distributed framework proposed by Minsky (1986).

If this way of thinking has merit, it suggests that we begin asking additional questions when studying behavior. For instance, we should
confront a task and ask not only, “How does an agent think about thi
e.g., categorize elements in it, construct a problem space repre
it?” but also, “What actions can an agent perform that wi
more *manageable*, easier to compute?”

This represents a shift from orthodox cognitivist
theme in cognitive psychology has been to discover
agents use to structure their environments. C
properties of the stimuli agents find in
the effects of these changes on such
recognize, complete, and so on
elements of the stimulus are
better, faster, more ofte
recall and recogni
subject, in impo
experiment
subje

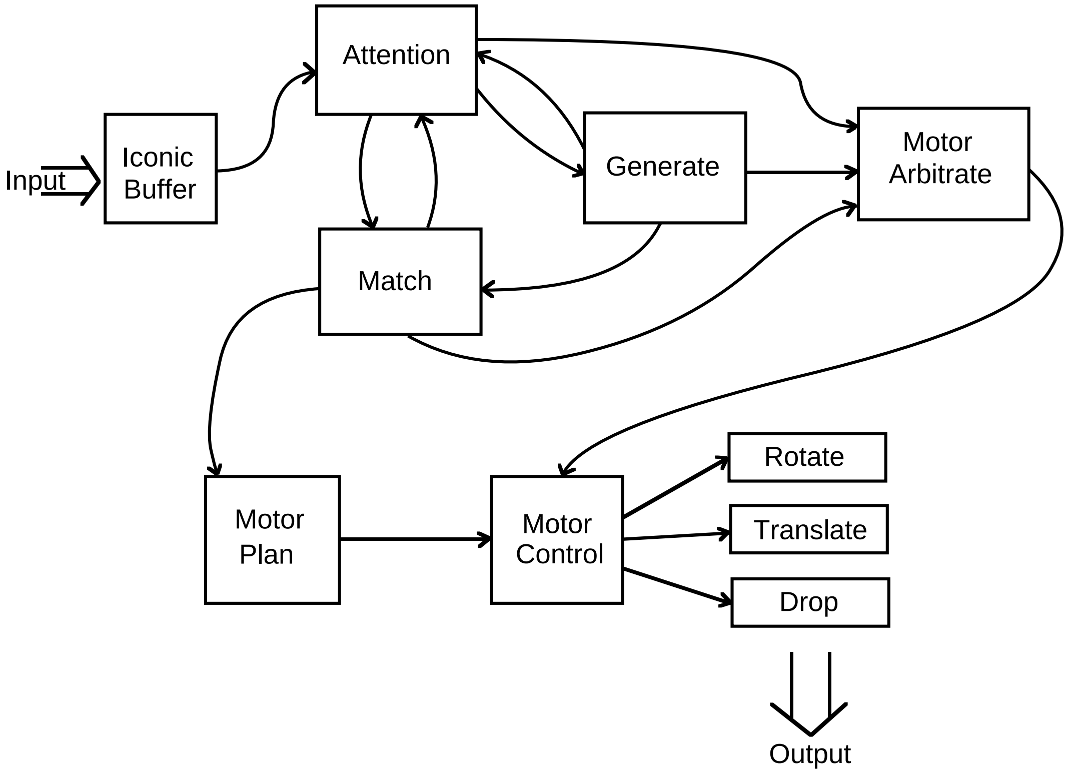

Figure 17. A more complicated model of the processes occuring in Tetris-
cognition would represent particular functional parts as a directed network of
mental processes able to pass messages between each other. The only significant
deviation from the sketch in Figure 16, is that two way links between attention,
andidate generation, and matching are shown, and a new process, called an ar-
or, is introduced to intervene between the possible calls to translate, rotate,
op.

There is, of course, nothing wrong with this approach. It permits controlled study. But it reflects a bias that the type of environmental structuring relevant to problem solving, planning, and choice, as well as to recall and recognition, *occurs primarily inside the agent.* That is, the environmental structure that matters to cognition is the structure the agent *represents* (or at least, presupposes in the way it manipulates its representations). No al-
lowance is made for *offloading* structure to the world, or for arranging th
so that the world pre-empts the need for certain representations, or
the need for making certain inferences. This leaves the perfo
pre-emptive and offloading actions mysterious.

To take a simple example, a novice chess player usu
to physically move a chess piece when thinking about
Why is this? From a problem space perspective,
superfluous. It cannot materially alter the cu
Yet, as we know, by physically alteri
imagining moving a piece, novices f
replies, and positions. In l
finds it helpful to cha
but moving to a
leap into f
way of
up o

they have on the agent.

This way of thinking treats the agent as having a more *cooperative* and interactional relation with the world: the agent both adapts to the world as found, and changes the world, not just pragmatically, which is a first order change, but epistemically, so that the world becomes a place that is easier to adapt to. Consequently, we expect that a well-adapted agent ought to know how to strike a balance between internal and external computation. It ought to achieve an appropriate level of cooperation between internal organizing processes and external organizing processes so that, in the long run, le
work is performed.

We conclude with a brief explanation of how accepting the catego
epistemic action affects traditional AI planning.

Epistemic Actions and Theories of Planning

In the introduction, we suggested that AI planners might accomodate epistemic activity by operating in a state space whose nodes were pairs encoding
both physical state and informational state. In that case, the payoffs a pl
receives from an action have two dimensions: a physical payoff, and an
mational or epistemic payoff. The clearest examples of epistemic
those which deliver epistemic payoffs rather than pragmatic o
nale, presumably, is that in each such case, after we have
of time lost performing the action, the expected ep
benefits still outweigh the expected net benefit
action.

The cost-benefit model that seems to ap
used to characterize the tradeoff betwee
since Stigler's seminal paper “The e
pointed out that for consumers
price of a camera, market inf
how much one could hope to
assume that prices fit a no
for a lower price decr
gain of one more
In c
mo

tions are most informative when what is seen is ambiguous in both shape and position. The model also fits the translate-to-wall routine. Thus, we explain why the probability of translating to the wall and back before a drop varies with drop distance by pointing out that the greater the drop height, the more informative the verification and the less risky (costly) the action. It also explains why players physically rotate to save mental rotation: they can attain the same knowledge faster and with less effort than by mentally computing
the image transformation. Rotating to facilitate matching has a favora
cost-benefit spread because matching via perception is fast, reliabl
less resources than matching in working memory.

The virtue of such a cost-benefit account is twofold. Fir
to continue modeling the decision about what to do next
among accessible actions. Without a notion of epis
justify why expert players sometimes choose prag
actions within a rational-agent calculus.

The second virtue of a cost-benefit
the superior decision-making of expe
more expert a player is, the mor
the costs of computation do
performing more epistem
But when we look
epistemic payoff of a
and benefits fail
siderable det
dertaking
Raif
I

well as on how it generates and tests candidate placements, and on how it attends to details of the contour and zoid. This requires understanding an agent's active cognitive processes to a level of detail unheard of in standard planning and rational decision accounts.

The upshot is that to incorporate epistemic actions into a planner's repetoire, we will need to cast aside the assumption that planning can proceed *without regard to* specific mechanisms of perception, attention, and reasoning. This idea is not foreign to the planning community, but to dat
has been restrictively applied. For instance, in discussions of acti
where repositioning sensors is a central concern—the decisi
to reposition a sensor is thought to depend on assumpti
sor's range, field of view, noise tolerance, and so o
inner functioning of the sensor. It is our beli
about an agent's internal machinery gener
tions, and that once more is known
selection in particular domains
prevalent than anyone woul
showing how, in a game as
actions that make it e
candidates, and i
derstand some o
we claim, h

References

Ambros-Ingerson, J. & Steel, S. (1987). Integrating planning, execution and monitoring. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 83–88.

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge, MA: Harvard University Press.

Baddeley, A. (1990). *Human memory: Theory and practice.* Boston, MA: Allyn and Bacon.

Bratman, M. (1987). *Intention, plans, and practical reason.* Cambridge, MA: Harvard University Press.

Chapman, D. (1989). Penguins can make cake. *AI Magazine*, *10*(4), 45–50.

Chernoff, H. & Moses, L. (1967). *Elementary decision theory.* New York: John Wiley.

Felleman, D. J. & Van Essen, D. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, *1*, 1–47.

Hitch, G. J. (1978). The role of short-term working memory in mental arithmetic. *Cognitive Psychology*, *10*, 302–323.

Howard, R. A. (1966). Information value theory. *IEEE Transactions on Systems Science and Cybernetics*, *2*, 22–26.

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.

Jolicoeur, P., Ullman, S., & Mackay, M. (1991). Visual curve tracing properties. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 997–1022.

Kirsh, D. (1990). When is information explicitly represented? In P. Hanson (Ed.), *Information, language, and cognition.* Vancouver, British Columbia: University of Vancouver Press.

Kosslyn, S. (1990). Mental imagery. In D. Osherson, S. Kosslyn, & J. Hollerbach (Eds.), *Visual cognition and action*, pages 73–98. Cambridge, MA: MIT Press.

Lerdahl, F. & Jackendoff, R. (1983). *A generative theory of tonal music.* Cambridge, MA: MIT Press.

Minsky, M. (1986). *The society of mind.* New York: Simon and Schuster.

Neisser, U. (1967). *Cognitive psychology.* New York, NY: Appleton-Century-Crofts.

Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

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: Lawrence Erlbaum.

Norman, D. A. (1988). *The psychology of everyday things.* New York: Basic Books.

Park, D. C. & Shaw, R. J. (1992). Effect of environmental support on implicit and explicit memory in younger and older adults. *Psychology and Aging*, *7*, 632–642.

Raiffa, H. (1968). *Decision analysis: Introductory lectures on choices under uncertainty.* Reading, MA: Addison-Wesley.

Reason, J. (1990). *Human error.* Cambridge, England: Cambridge University Press.

Reeves, A. & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review*, *93*, 180–20.

Sereno, M. I. & Allman, J. M. (1991). Cortical visual areas in mammals. In A. Levinthal (Ed.), *The neural basis of visual function*, pages 160–172. London: Macmillan.

Shepard, R. N. & Metzler, J. (1971). Mental rotation of three-dimensiona objects. *Science*, *171*, 701–703.

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

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, *74*.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, *69*, 213–285.

Tarjan, R. E. (1985). Amortized computational complexity. *SIAM Journal on Algebraic and Discrete Methods*, *6*, 306–318.

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.

Treisman, A. & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General*, *114*, 285–310.

Ullman, S. (1985). Visual routines. In S. Pinker (Ed.), *Visual cognitio* pages 67–159. Cambridge, MA: MIT Press.

Waltz, D. (1975). Understanding line drawings of scenes with shadows. P. H. Winston (Ed.), *Psychology of computer vision*. Cambridge, M MIT Press.

Author Notes

We thank Steve Haehnichen for his work on the initial implementions of Tetris and RoboTetris. Charles Elkan, Jeff Elman, Nick Flor, Jim Hendler, Ed Hutchins, and Teenie Matlock provided helpful comments and suggestions on early drafts of this paper.