# Thinking through Diagrams: Discovery in Game Playing

Susan L. Epstein

Department of Computer Science, Hunter College and The Graduate School of
The City University of New York, New York, NY 10021, USA
susan.epstein@hunter.cuny.edu

**Abstract.** Diagrams often play an important role in human problem solving. A diagram can make a task more difficult, however, when it obscures important problem features, or requires repeated effort to interpret what it represents. Moreover, the nature and origin of diagrams can be as important as their exploitation during problem solving. This paper chronicles the complex interleaving of visual cognition with high-level reasoning in three subjects. Their diagrams and subsequent verbal protocols offer insight into human cognition. The diversity and richness of their response, and their ability to address the task via diagrams, provide an incisive look at the role diagrams play in the development of expertise. This paper recounts how their diagrams led, and misled, them; how their diagrams both explained and drove explanation; and how a spatial approach reported here can lead to deeper understanding of other simple games.

## 1 Introduction

Spatial information can play important roles in problem solving, learning, and discovery. The centerpiece of this paper is an experiment in complex cognition that explores this idea. Three subjects investigated the simple game described in Figure 1. Each of them invented diagrams to cope with the task's complexity and reasoned from that spatial information. The primary results of this work are the subjects' preference for one-dimensional representation in thinking about a two-dimensional space; the power of a diagram to displace the original reasoning context; a new diagrammatic convention and associated algorithm that provide insight into cyclic, two-person games; and a demonstration of the synergy between visual and high-level processes during ab-

**Pong hau k'i**

- Two-player game
- 5 locations
- First to move plays black
- Second to move plays white
- Move = slide to the next empty location
- Initially, the board is as shown
- Mover who cannot slide loses

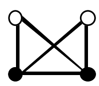

**Fig. 1.** As presented to the subjects, the game of pong hau k'i [21]. The mover slides her piece along a line to the single empty location on the board. The goal is to *trap* the other contestant, so that on her turn she cannot slide. Contestants take turns until a mover loses because she is trapped, or until a draw is declared because the mover and the location of all the pieces have been repeated more than three times.

stract reasoning. The focus here is on how spatial information supports, and interferes with, high-level reasoning.

This work addresses cognition about movement in an abstract, rather than a realistic, two-dimensional space. Although there are studies of how individuals change representation to facilitate the solution of one-person problems [1], and studies of how individuals respond to a variety of representations for the same two-person game [22], to the best of this author's knowledge there are no studies of how individuals change representations to facilitate learning a two-person game. This paper is in the tradition pioneered in [2], particularly with respect to the third subject. The first section of this paper provides terminology and an overview of related work. Subsequent sections describe the experiment, the subjects' solutions, and the processes used to reach them. Finally, the results are discussed, along with suggestions for future work.

### 1.1 Representing problem solving and search spatially

The classical AI approach to behavior in a *domain* (problem area) is to specify *states* (ways the world can appear) and *actions* that transform one state into another. Given a set of states and actions defined on them, a *problem* is a *start state* (description of the initial state of the world) and a description of one or more desired *goal states,* to which a sequence of actions should lead. *Problem solving* can thus be modeled as *search,* finding a sequence of actions that transform the start state into a goal state.

The set of all possible states in a domain is its *state space*; it represents all the states within it, and how they relate to one another A *state space diagram* represents a state space spatially, as a set of *nodes* (each of which depicts a state), some pairs of which are linked by *edges*. An edge from one node to another indicates the existence of a problem-solving action that transforms the first state into the second. In a state space diagram, problem solving can be envisioned as finding a path along edges through the diagram, from the start state to a goal state. To construct a state space diagram, one must describe both the individual states and their interrelationships.

A *game tree* is a state space diagram specialized for two-person games. An example of part of a game tree appears in Figure 2. Note the horizontal alignment of the nodes there into *levels*, and the labels to the left that note whose turn it is to move on that level. A state that ends play in a game tree is *terminal*, and is associated with an *outcome* (e.g., win, loss, or draw). *Move generation* is the enumeration of every pos-

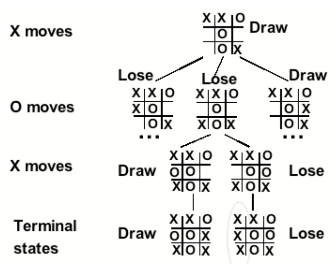


**Fig. 2.** Part of a game tree for another game, lose tic-tac-toe, where the object is to avoid three in a row, column, or diagonal.

sible action from a game state, and can itself be a complex task; some subjects in this experiment developed diagrams as aids to move generation.

For a game like that in Figure 2, where a piece, once placed, is fixed to a single location on the board, a game tree describes possible lines of play clearly: a contest moves downward from the start state along one edge at a time until it reaches a terminal state. In the game described in Figure 1, however, where a move may return the board to a previously visited state, the game tree must be adapted, either by repeating states (so the tree becomes infinite) or by introducing a *cycle* (a path of edges returning to an already-represented state).

### 1.2 Cognitive issues in state representation

A *representation* for a state describes it with some degree of detail. One might, for example, represent the state in Figure 1 as follows: "it is black's turn, the pieces are made of wood, black's pieces are on the two lower locations, white's pieces are on the two upper locations, and the center location is empty." A good representation captures the salient features of a state, without irrelevant details. In pong hau k'i, the material from which the pieces are made has no impact on problem solution. Similarly, their size and shape are irrelevant.

Once the salient features of a problem state are identified, there are still many possible ways to represent an individual state. For example, the pong hau k'i start state could be represented in any of the ways shown in Figure 3. In some games (e.g., tic-tac-toe), it is possible to determine whose turn it is to move merely by examining the board. In pong hau k'i and many other games, however, one must identify the *mover* (the contestant whose turn it is to move) as part of the state description. Figures 3(a) and 3(b) are *visualizations*, representations that convey information by the spatial arrangement of elements of the display itself [10]. Figure 3(a) is a *complex visualization* because it suggests animation (sliding pieces along the lines).

A human problem solver has at least two representations of a problem: the *external representation* provided by the environment and the *internal representation* constructed within the subject's memory. The external representation, although both explicit and implicit, is for the most part under experimental control; the internal representation can only be glimpsed through a subject's behavior. Furthermore, a problem solver may construct additional implicit representations, and, within the restrictions of the experiment, additional explicit ones.

A good state representation speeds solution. Kaplan and Simon's work on the mutilated checkerboard problem showed that cues (e.g., colors, or words that are commonly paired with each other) inherent in a representation can suggest a particular line of reasoning to people [11]. They also showed that hints to subjects who had dif-

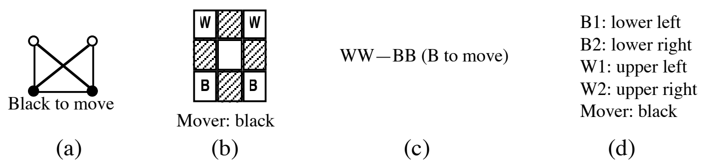


**Fig. 3.** Some representations for the start state in pong hau k'i.

ficulty solving the problem could help them change to a representation (make a *representational shift*) that incorporated such cues. Their subjects required motivation to make that shift, however, and cues proved constructive only once subjects were dissatisfied with their progress in the current representation. Their best subjects could detect the salient features of the space, and relied heavily on perceptual details. They reported that a good representation produces a smaller space with an effective generator that relies on *invariants*, things unchanged by an action.

Representational shifts sometimes increase expertise. Amarel's work on the missionaries and cannibals problem, for example, argued the increasing power of a sequence of different representations, progressing from natural language to a compact mathematical symbolism [1]. On the 3-disk tower of Hanoi problem, Anzai and Simon described how a single, untrained individual performed a sequence of representational shifts that led to better performance and to a deeper understanding of the problem [2]. Deliberately obscured representations of the 3-disk problem (involving monsters and globes), however, were far more difficult for humans than the original problem [12]. Careful analysis showed that most of those subjects' time was devoted to exploration of the possible states, followed by rapid detection of a correct sequence of actions. Fast, accurate move generation from the problem description proved essential to expertise; the ability to construct simple two- or three-move plans was a further asset. Memory load was also a consideration: the more a subject was forced to recall (e.g., imagining the monsters, the size of their globes, and changes to their globes), the more difficult it was to plan. Practice on repeated problems seemed to support planning and learning. An external representation (e.g., a diagram) also helped, but objects in more than one location, actions entailing spatial changes, or more than one imaged entity still made solution more difficult.

### 1.3 Search, spatial information, and learning

Spatial information can reduce search effort by providing problem-solving cues. For example, a program that constructed simple plane geometry proofs from an input problem and an accompanying diagram was able to construct more complex proofs when it could take metric cues from the diagram [8]. Spatial information can also encourage the substitution of (faster) perceptual judgments for logical inference. Such a judgment is a recognition that the elements in a diagram are spatially organized in a way that supports problem solving. DiBS (Display-Based Solver) used list-based problem input to construct its own set of one-dimensional diagrams [14]. These diagrams captured relationships (coded as "below") that described the goal state and represented subgoals as objects that blocked other objects from their intended final locations. The program could solve several quite different, simple, one-person problems (the 3-disk tower of Hanoi, simple algebra, and making coffee) by transformations to its diagrams alone. This approach demonstrates the wealth of information a diagram can store, the cues it can offer a solver, and the power of an automated representational shift. DiBS, however, requires careful encoding of domain knowledge and of the problem itself, and is intended for single-agent problems with a single relationship category. Moreover, DiBS requires the description of goal states in advance, although part of some tasks is to identify goal states for the first time. DiBS also assumes that the order in which actions occur is irrelevant, and that there is a single agent, the results of whose actions are certain. The rules of two-player games, however, are quite strict about sequence, and a single decision may not determine the next state in which one makes a choice.

Domain expertise can be pre-specified, or it can be learned. From an AI perspective, *learning* is a change in one's problem solving behavior based on experience. The most challenging form of learning is *discovery*, where the learner has no external teacher. In game playing, for example, learning to play merely by playing, without expert instruction, constitutes discovery.

Spatial information can support learning in a variety of ways. The construction of a state space diagram helps identify the states that actually arise during problem solving. Some spaces also lend themselves to a kind of organization that prevents repetition and supports move generation (e.g., the tic-tac-toe game tree). For game playing, the paths through a state space diagram can delineate opportunities (possibilities to win) and dangers (possibilities of loss) during play. Furthermore, a diagram can support generalizations about the paths in the state space, generalizations readily converted into heuristics for efficient, effective problem solving. Examples of this abound in the central experiment described here.

The importance of spatial information in human scientific discoveries is well-documented [9, 15, 17]. One way to validate a theory about the discovery of scientific laws is to automate the process and replicate the results. Most such AI artifacts begin with numeric data that simulate the results of hypothetical experiments, and induce rules from them [7, 13]. HUYGENS, however, first transformed the data into a set of one-dimensional diagrams, and then sought perceivable regularities within those diagrams, using heuristics to focus attention [4]. HUYGENS could extend its initial diagrams during search, so that it not only reasoned about diagrams but also constructed new ones. The program was limited, however, to domains where one-dimensional diagrams could represent salient features.

Diagrams can also make learning more difficult. People learning to play isomorphs of tic-tac-toe experienced varying degrees of difficulty with eight disguised versions of the game [22]. Each isomorph was accompanied by a diagram that offered or obscured invariants ordinarily perceivable without high-level reasoning. Zhang showed that perceptions about those invariants supported or undermined one's ability to learn to play this simple game, to explore the space, and to discover structure within it. Tic-tac-toe has a large state space: 5746 reachable distinct states, and 120 different possible paths to a terminal state. Zhang argued that, if accessible, the symmetric invariants made search in so large a space more manageable, and that the winning invariants elicited a particular human bias about the world: more is better. His work is an elegant argument, at least in this specific domain, for *representational determinism*, the idea that a representation both guides and constrains what a problem solver can achieve.

## 2 The Experiment

Pong hau k'i was an informal assignment for students at the 2001 Cognitive Science Summer School at the New Bulgarian University. During a lecture, they were given the rules for this old Korean game, with the slide in Figure 1. None of the students knew the game. They were told that it was a very simple, even boring, game, played between black and white. They were asked to "play it, take protocols on the development of your own expertise, and capture any representational shifts," defined as "a different way of seeing things." (No example of a representational shift was provided.) They were encouraged to obsess over it and to report back on their experience.

The same lecture introduced state space search, including the slide in Figure 2. Later in the lecture, the students were asked, on additional slides, to "draw out the state space for pong hau k'i, label [it] with outcomes, produce a rule set to play it, …and produce a good heuristic for it." No guidance on representation for repeating states was provided, and no representation for the pong hau k'i states was specified. The author expected that the students would use the representational conventions of Figures 1 and 2, but, as we shall see, the students quickly abandoned them.

This task began as a teaching device, not as a planned experiment. As a result, it lacks the timed photographs, videos, and recordings that lend rigor to some other work in this field, [5, 16, 18, 19]. Once several students began to respond in such great depth, however, documentation was as extensive as possible. The students of the Cognitive Science Summer School are a diverse lot, drawn from many countries and many disciplines. The Summer School is conducted in English, but the students' written materials were idiosyncratic, and included notation in both German and Polish. For uniformity, in what follows this author has made minor spelling corrections, inserted clearer or more grammatical language [in brackets like these], and substituted B for "black," W for "white," and "–" for an empty location. Any italics are the subjects' own, and all names have been omitted. During the experiment, no subject described here saw the protocols of the others.

### 2.1 Pong hau k'i

By game-playing standards, pong hau k'i is very simple: there are only 60 states in its search space. Moreover, pong hau k'i is a *draw game*, that is, if both sides play correctly, play always ends in a draw. Even if one is not familiar with the game, it ought to be relatively easy to play well.

Compared to other simple games (e.g., tic-tac-toe), however, pong hau k'i is actually rich and challenging, for several reasons. The pieces are permitted to move about the board; this encourages competitors to envision "movies" rather "snapshots" of the board, and places a high load on their spatial memory [12]. It is also impossible to tell from the board alone whose turn it is to move; it could be black's or white's from any state. Finally, there is only symmetry across the vertical axis on the board; diagonal and horizontal symmetries are invalid.

### 2.2 Subject 1 — an automated approach

Subject 1, a graduate student in psychology, ostensibly drew no diagrams at all. At lunch after the lecture, several students made pieces from bread (the crusts against the bread centers) and began to play. Although he professed disinterest, eventually the fun drew Subject 1 in, and he began to play against Subject 3. "After a while the game became dull, because neither [of us]… understood the point of it.… We started to discuss, that there must be some good strategy, and we must find it.… [Subject 3 began] to draw a tree-like decision graph, where every branch represented a possible move. …The only thing I remember of his drawing was that he was thinking of a good labeling of the [locations]. We agreed that numbers are good indicators of the locations, but we didn't talk about [how to assign numbers to locations]. He wanted to count how many [distinct states there were]. But it was really time to go back to [class]."

Subject 1 is not a trained programmer. Nonetheless, alone later that evening he wondered "why people should bother with counting [moves] so much … if computers are much more suitable…. Therefore I decided to build a program that can play this game." Although his original plan was to play against the computer and have it count the states for him, he subsequently realized that the program might also serve as an opponent. "At this point, I did not think of any strategy, only the rules [of the game]… I remembered the lunchtime games, and [Subject 3's] idea that a good labelling is needed. Without much thinking I came up with [Figure 4]. The reason I chose this labelling is because if you make a 'line' from the original shape by 'pushing' from the top and from the bottom the numbers will increase from 1 to 5."

Two hours later Subject 1 had a program that reproduced the pong hau k'i game board, but with red and blue (rather than black and white) circles for the playing pieces. The code, however, did not employ Subject 1's numbering in Figure 4. The order in which the circles were drawn on the screen on each turn, and the order in which they were stored, using his numbering from Figure 4, was 2, 1, 3, 5, 4.

Subject 1 now began to play against his program, which made random legal moves. "This was the first time, I wanted to understand the strategy." Subject 1 soon won against his program, played "10 –30" more contests, and remained convinced that he was correct. Still later the same evening, Subject 3 appeared in Subject 1's room. "He [Subject 3] … seemed very tired. He became delighted that" another person was working on pong hau k'i so intensely. "So we quickly started to play," but Subject 3 warned that he was undefeatable. "After 5 – 8 steps we gave up, because I couldn't lose [either]. Tie game. By that time, we realised that both of us [had become experts at pong hau k'i]. He did his way theoretically, I [did] mine practically."

### 2.3 Subject 2 — a (relatively) quick response

Subject 2, a Ph.D. candidate in cognitive psychology, originally offered no verbal protocol with his analysis, only an elaborate, tree-like diagram on a single sheet of

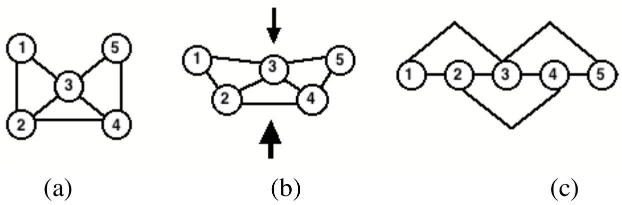


**Fig. 4.**Subject 1's compression of the board to create a numeric encoding, as given in his protocol. (a) The board itself, (b) the initial pressure , and (c) the collapse into a line.

paper, with some small annotations described below. Subject 2 drew his diagram in a few hours, stopping “when I thought it was complete….” He did not use a game board, rather, “At the beginning I tried to imagine the situation on it. Then I just manipulated the letters and numbers. I drew a sketch when I was afraid I was lost, to check it.”

It is possible to deduce a good deal of Subject 2’s process from his diagram (which is difficult to reproduce and therefore omitted). Figure 5(a) replicates the top portion. In the upper left corner is the start state from Figure 1. Below it are two sketches symbolizing how white could win. (Line segments were indeed omitted, and most sketches, as reproduced here, were also more square than the original in Figure 1.) To the right of the start state is an encoding of it (WW–BB), from which one may deduce the location numbering: 1 = upper left, 2 = upper right, 3 = center, 4 = lower left, 5 = lower right. “12345” lists the labels assigned to the vertices. Immediately below it is a list of the seven pairs of locations between which a move may occur.

There are only two other sketches of the board on the page, both to the right. The first, Figure 5(b), is near the two losing states highest in the diagram, about eight levels down in the tree. It represents the state just before a possible win for either mover. The second, Figure 5(c), appears next to a section of the game tree that is placed far to the right but connected to the diagram (the *subtree-to-the-right*). Figure 5(c) represents a state from which a white mover can win; if black is the mover the contest will merely continue. These are the only sketches Subject 2 drew to check his reasoning.

Most of Subject 2’s page is occupied by a diagram in the spirit of Figure 2. Each node is a five-letter string representing the contents of the five locations on the board. Nodes that are the same number of moves from the start state are horizontally aligned in levels on the unlined paper. In most of the tree, a W or a B to the left on each level indicates the mover in the previous state; in the subtree-to-the-right these labels were to the right. The two alternative first moves for black appear beneath the start state. Subject 2 chose to follow only the move from 4 to 3, shown below and to the left of WW-BB in Figure 5(a). At most levels there are only two or three nodes, but occasionally there are as many as six. The diagram reaches all margins of the paper, and includes 80 nodes in all. One node, on the eleventh level, points with a long arrow to the subtree-to-the-right. Otherwise, exploration appears to have had a strong downward tendency, and to have moved from left to right at every level.

Most tree edges are simply lines in his diagram, but at some point Subject 2 noticed that some states repeated. When he found a repeated state (encoding and mover),

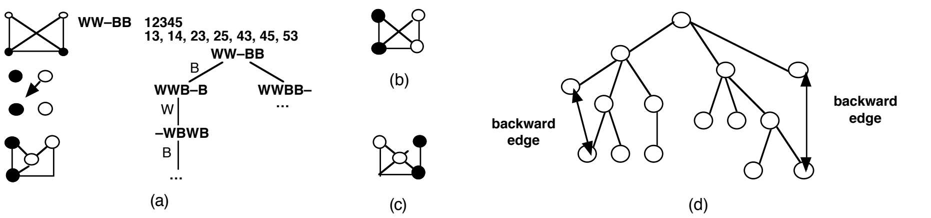


**Fig. 5.** (a) The top portion of Subject 2’s protocol (ellipses are the author’s), (b) a board configuration penultimate to a loss; either contestant can win from here, (c) a board configuration from which white, but not black, can win, and (d) orientation of forward and backward edges in Subject 2’s tree.

he drew an unlabeled edge that was *directed* (had an arrow on at least one endpoint), as in Figure 5(d). A *backward edge* begins at one node and runs to another *higher* (closer to the start state) in the tree; there are 11 backward edges. One backward edge in particular, from the leftmost portion of the subtree-to-the-right, returned to the start state and presumably eliminated the need to explore the opening move to WWBB– on the right in Figure 5(a). A *forward edge* runs between two nodes and has arrows on both ends. The diagram includes two forward edges; the only difference apparent between backward edges and forward ones is their length — the repeating pairs joined by a forward edge are those farthest apart on the page. There is also one directed edge between two nodes on the same level, and one dashed forward edge between identical nodes associated with different movers. In the latter case, Subject 1 presumably observed the similarity, noted it with an edge, and judged it not worth pursuing.

When a node has no others beneath it in Subject 2's diagram, it is either the source of a backward edge (indicating repetition in play) or is surrounded by a rectangle (indicating the end of a contest). There are eight such nodes, each with a W or B label indicating the winner in that state. Only four of these are distinct: two wins for white and two for black. There is one significant error: the forward edge from the unexplored first move WWBB– is not to a later copy of itself. There is, however, another copy of WWBB– (one unremarked upon by Subject 2) that is correctly considered.

### 2.4 Subject 3 — an obsession

Subject 3, a Ph.D. student in linguistics and cognitive science, has a strong mathematical background. He had been among the students at lunch that day with Subject 1. The day after the lecture, Subject 3 proudly displayed a diagram he had drawn of the search space. He seemed exhausted, and questioning revealed that he had devoted a good deal of time and energy to the task. Although he thought little of them, he had not yet discarded his earlier diagrams, which chronicled his path. He remembered a great deal of the process he had been through in the past 24 hours, so together the author and Subject 3 numbered, discussed, and annotated his early diagrams. From these, Subject 3 began to write what eventually became a 13,000 word protocol, covering about 6 hours of intensive exploration and many more of retrospective reconstruction. The material in this section is based on observation of his diagrams, the debriefing session the day after the assignment, and the original and later versions of the written protocols. Subject 3 identifies 9 stages in his path to expertise at pong hau k'i.

**Stage 1.** (15 minutes) Subject 3 attempted to "get a feeling for the game and its state space. Initially, he played "very much at random without even knowing how a [losing] state might look…." Indeed, he hypothesized one state as a loss which was not a loss at all. Nonetheless, "After a few more (random) moves my opponent pointed out that his [pieces] were stuck and that he had lost. I looked and realized that indeed, by accident, I had won this first game. Both his [pieces] were on the [right side] of the board and I was blocking him from moving…. From seeing this I figured that *one could lose if (and only if) one ended up having both [pieces] on the same (left/right) side of the game board*." In a second contest, Subject 3 lost. Then, as the others watched and commented, he began to experiment alone, moving "… according to the rules, without consciously following any strategy. … I got the feeling that the dynamics of the game contained several symmetries.… by a few straightforward moves I could invert the initial configuration (interchange the black and white pieces). And I also got the feeling that I would always end up in these same two [states] (the initial one and its inverse).… I had the sensation of running into 'attractor states.' Then I tried to … [force either player to lose]. It only occurred occasionally and rather by accident than by controlled force. I could not really find out which moves one had to do/avoid in order to win."

**Stage 2**. (20 minutes) Subject 3 began to draw his first diagram. He computed 30 possible *configurations*, ways to place the pieces on the board; with the label for the mover, he arrived at 60 possible states. "However, I was not yet sure that all of these states could actually be reached by legal moves from the initial configuration." For compactness, Subject 3 now shifted to a rectangular representation for a node. (He was drawing on graph paper, where a sequence of squares is easy to draw quickly and clearly.) Because black moves first, Subject 3 labeled the 2 lower locations 1 and 2 (left to right, respectively). He labeled the center 3 because he felt it retained "some of the symmetry of the original game board." Originally the top locations were (left to right) 4 and 5, but he interchanged them (5 and then 4) so that 2 to 4 and 1 to 5 were moves. This produced the numbering in Figure 6(a). Using this numbering he wrote sample nodes, such as Figure 6(b), to confirm his calculation of the possible states.

Still at lunch with his friends, Subject 3 now began to draw his second diagram, his first attempt at the full state space. All nodes were in the format of Figure 6(c), where the additional circle to the right denotes the mover. Although he initially checked his diagram against the pong hau k'i board, Subject 3 eventually relied solely on his new representation. He set up the two opening moves at the top, beneath the start state, just as Subject 1 had. "Then I followed only one of the two now opened paths …depicting white's and black's possible moves in alternating order. At a certain point I would interrupt myself and work a bit on the right path hoping for some cycle to close. The whole diagram became very confusing: with almost every new state that I drew I would get the feeling of having already encountered this particular one." As he searched his diagram looking for cycles, Subject 3 became "rather annoyed by the fact that I had to distinguish the mover with the state." He decided not to extend "paths starting from those states that occurred for the second time, but with the opposite mover. I believed I could just refer to some sort of symmetry with the paths starting at the corresponding state that I had visited first." At some point he realized that he had forgotten about the edge from 1 to 2: "… indeed, I found two new 'twigs' to be grown in my graph." Returning to it hours later, however, "Everything just became more and more messy and confusing…." He abandoned the task.

**Stage 3**. (45 minutes) Somewhat later, Subject 3 began alone, with new determina-

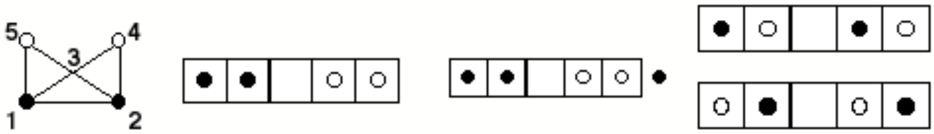


(a) (b) (c) (d)

**Fig. 6.** (a) Subject 3's numbering of the board. (b) The starting board and (c) the start state represented in Subject 3's rectangular coding. (d) The "seeds" for his fourth diagram.

tion, to draw a third diagram, this time in color. The playing pieces became green and red, edges were colored green or red to denote the mover transforming the state, and the external circle for the mover was eliminated. Edges now had arrows on both ends, indicating that moves were invertible. Subject 3 referred back to his previous diagram regularly, checking for oversights, and frequently found them. Eventually, he extended the space methodically, keeping both sides at about the same depth, and soon reached four winning nodes, two for red and two for green. These were highlighted in orange to denote termination of search. Then he expanded the non-terminal states. "I started to count the states. There were 30 of them. I took this correspondence with my earlier calculation as a strong confirmation for finally having covered the entire state space. … But my representation was just unbearably entangled. (And in fact, I was surprised by the apparent complexity of the state space.) I wanted to shape it up towards a comprehensible new diagram." Subject 3 considered the diagram crowded, and it did not support the symmetry he expected; in particular, he "had an expectation that the opposite of the start state would be on the very bottom…naturally."

**Stage 4**. (50 minutes) Subject 3 drew his fourth diagram, beginning not from the initial state but from "another state of maximal connectivity." He chose the states in Figure 6(d), and positioned them first, horizontally, in the center of his diagram, with the upper one on the left, and the lower one on the right. Subject 3 then traced backward through the previous diagram to position the start state appropriately. He checked off states from his previous diagram as he reached and reproduced them in the left half of the new one, moving outward from the upper state in Figure 6(d). At the same time, he maintained elaborate tests for symmetry and edge color, for states he regarded as

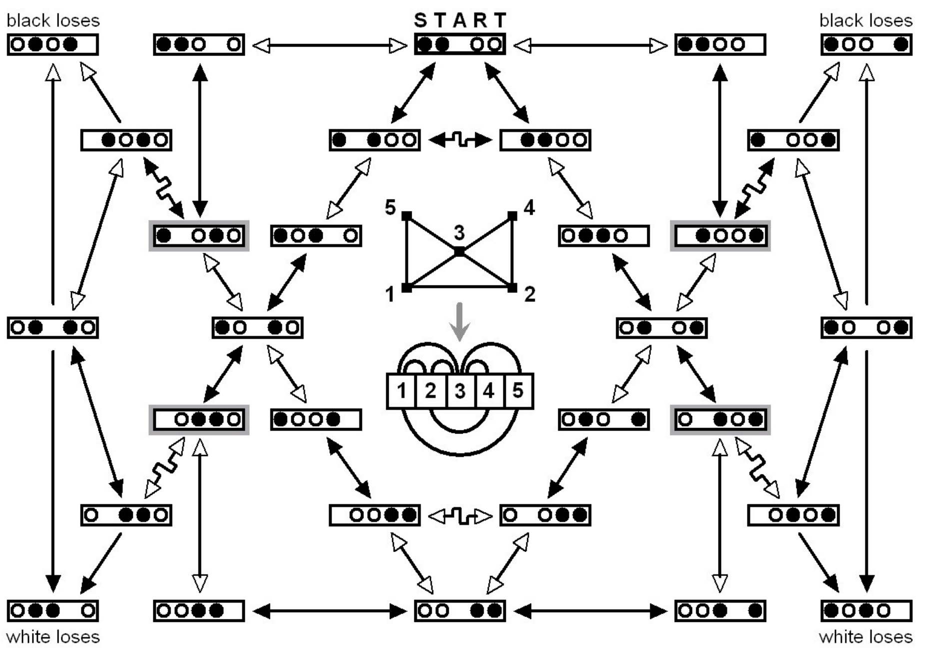


**Fig. 7.** Subject 3's sixth diagram, the product of 6 hours of intensive work. The four delicate states have gray borders.

“opposites,” and for edges that were incorrect but somehow “missing. Although he “was expecting some asymmetries,” intersecting edges drove him to redraw it. After the 50 minutes of this stage, he writes that he now considered himself “obsessed.”

**Stage 5**. (50 minutes) His fifth diagram positioned the start state at the top, surrounded by empty rectangular nodes, that is, states without any playing pieces. Subject 3 concentrated on the shape of the space now, not the moves. At some point he filled in the nodes, and then focused on introducing the losing states so that the arrangement was spatially pleasing. To his surprise, Subject 3 found that his desire for symmetry now led to the discovery of edges that he had completely overlooked in earlier diagrams. Moreover, the angle of the edges now became significant to him. Subject 3 initially drew the right side of the diagram with empty nodes and colored edges, and re-derived the nodes before confirming them with his previous diagram.

**Stage 6**. (duration unknown) Concerned about an asymmetry on the left in his fifth diagram, Subject 3 explained the picture to his roommate. He pointed out four *delicate states* (ones where it is possible to make a fatal error) located between “the secure central area” and the areas on either side where one can lose. “The discovery of the *delicate states* and all my following considerations about strategies were entirely based on the arrow pattern and the state frames – which in my mind appeared almost deprived of their contents.” Subject 3 now correctly recognized that pong hau k’i depends entirely on play in the delicate states: “…to prevent losing, I … simply have to avoid the two bad moves… I only need to memorize two specific configurations… whenever my opponent allows me to win, I … simply have to make the one good move…. For this I only need to memorize two further specific configurations….” He

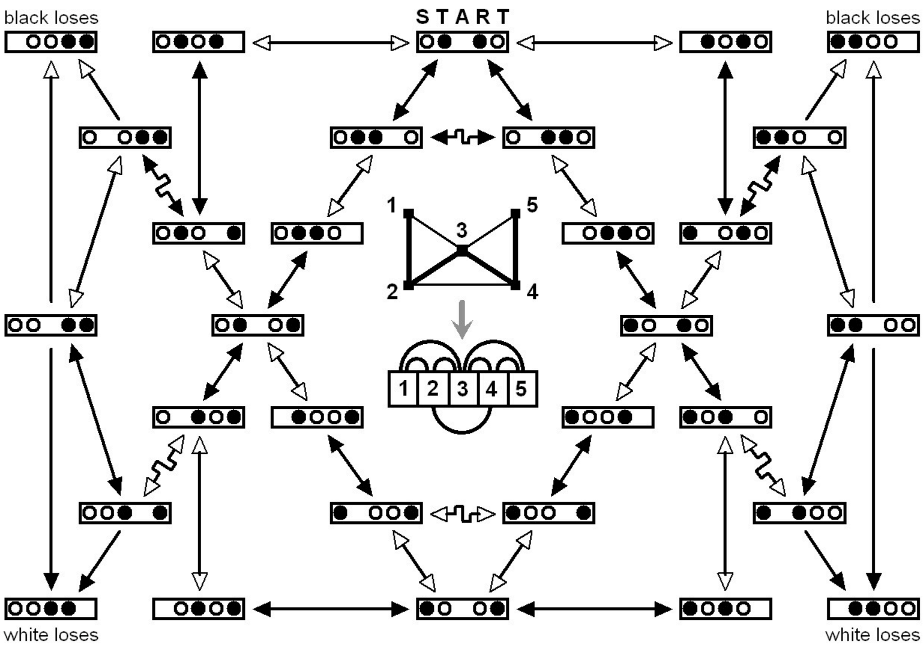


**Fig. 8.** Subject 3’s “W-shaped encoding,” seventh diagram, with the board locations renumbered as in the center.

was startled, however, by how irrelevant his full diagram was to actual play: “I am surprised that … one would not have to know anything about the global pattern of the state space. … My two strategies … are local rules – local in terms of time (which translates to space on my diagram). … I also noticed that the … strategy …for making the winning move …also serves for …preventing my defeat.…”

**Stage 7**. (20 minutes) The next morning, driven by “global symmetry,” Subject 3 drew his sixth diagram, shown in Figure 7. Although he also wanted to renumber the nodes, Subject 3 was pressed for time, and “wanted to focus on the intended rearrangements on the left-hand side of the preceding draft.” Therefore he “… drew all 30 empty state frames starting on the right half where I did not have to change much. Then I added all the arrows until finally I copied the states themselves.” He carefully aligned states vertically and horizontally, highlighting terminal states and delicate ones. The result was satisfying — he felt like an expert player and remarked on the symmetry of his diagram. “The stunning thing that was new here is the ‘reflection through the center’ symmetry and the fact that it operates not only on the state frames and arrow pattern but as well on the state configurations!! I liked this.” This was the diagram he initially presented to the author.

**Stages 8 and 9**. During debriefing that afternoon, Subject 3 recalled his renumbering scheme from Stage 7, which ultimately removed some misunderstandings observed (but not remarked upon) during the conversation. Later that afternoon, in Stage 8, Subject 3 produced his seventh diagram, using that renumbering (Figure 8), so that the patterns in nodes with the same function are symmetric to each other. Over the next few days, in Stage 9, Subject 3 revised and extended his protocol, and produced his final diagram (Figure 9). It classifies states and moves according to their safety

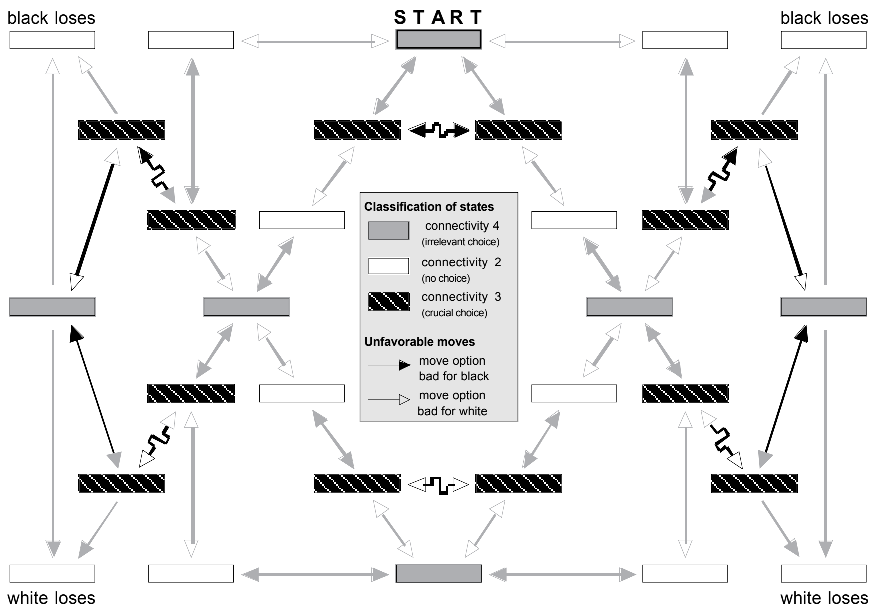


**Fig. 9.** Subject 3’s eighth diagram, showing classes of states.

and the number of moves they afford; it also maintains vertical, horizontal, and through-the-center symmetry with respect to states and edges. Note that there are no pieces whatsoever in the nodes.

Subject 3 now felt he was an expert, and had been since Stage 5, although he had only competed once (described Section 2.2) since then. "…in a way I had reformulated the game in my own terms and then detached it to some extent from its original shape. My (final) game consisted of boxes … and arrows and I had mastered these …." Why then had he continued to draw? His additional work, he believed, increased his subtle "…knowledge of the structure underlying the game. Knowledge that might not directly influence performance." Even as he dissected his experience, Subject 3 remained puzzled by his own lack of rigor. "… I probably could have proceeded in a much more intelligent/elegant way and in less time. I am not ashamed of this but rather surprised. For before [writing this protocol] … I [thought I had] approached the task in a more controlled and directed way than this documentation has brought to light." Nonetheless, it was a remarkable intellectual journey.

## 3 Discussion

The instructions to the subjects in this experiment facilitated solution [12]. Memory overload was prevented by asking them to draw a state space diagram. Enough experience to seed planning, learning, and generalization was provided by asking them to compete. Despite identical instructions and similar educational levels, however, their verbal and graphic protocols reveal a surprising diversity of process and of final perspectives, albeit a similar level of skill at the game itself.

### 3.1 The right answers

Pong hau k'i includes 30 possible piece configurations and 60 possible states. Subject 1 wrote code that could draw all possible board states, although he only experienced the ones to which he led it to during competition. Subject 2 found all 60 states (some more than once), but overlooked several edges and in one backward edge paired two different states, asserting an incorrect repetition. Subject 3's first four diagrams included all the configurations, but were missing edges; his others were complete.

Pong hau k'i is a draw game. One cannot win unless one's opponent errs; from most states, any move is safe. Only in the four states Subject 3 called "delicate" is it possible to make a fatal error, and only one of the two moves in each of those states will cause a loss. Even after the wrong move from a delicate state, one's opponent has two possible moves, only one of which wins. Questioned weeks later, Subject 2 wrote that after "a few games with one person" he realized that "it is quite difficult to lose the game, and I had an idea about situations I should avoid." Subject 1 made a similar discovery: "After my first [non-draw] I realised that the only way [to win] is to put the opponent into the side [locations. This describes both states in Figure 5.] … I tested this idea, and tried to avoid [having my pieces there] …." Thus, one would expect all three subjects to play expert pong hau k'i; drawing against each other, and winning if an opportunity arises.

The students were asked to construct a labeled diagram of the state space, a rule set, and a good heuristic. Subject 1 constructed a program that "knew" the space, without ever drawing it out. He learned rules and heuristics from experience that moved him through the space as he competed against his program, a kind of virtual representation. Subject 2 only constructed the space, but felt that he had completed the task. Subject 3 developed game graphs that proved the heuristics and the rule set that the other subjects arrived at. His final product is a powerful, rigorous, visual argument.

### 3.2 Invention of state representations

Although all these subjects received the same instructions, each of them approached the task differently. To represent an individual state, all three found the physical board itself confining and soon abandoned it. Subject 1's program was in part motivated by the burden of repeatedly drawing the board. Subject 2 switched because he wanted to fit the entire state space on a single page. Each subject began by numbering the board: Subject 1 with a thoughtful, visually-described process which he quickly abandoned; Subject 2 as if he were reading text, left to right and top to bottom; and Subject 3 left to right but from the bottom up. Subjects 1 and 3 considered more than one numbering. In every numbering, however, the central position was always labeled 3, presumably because it is also the median of the numbers from 1 to 5.

Clearly, perception and symmetry are powerful influences on state representation. Each numbering of the locations on the board drives exploration of the search space differently, and there are 120 possible numberings. It is likely more than coincidental that Subject 3's second numbering, in Figure 4, matched Subject 1's. Indeed, Subject 1's original idea was to explore several such numberings as if they were preference strategies, playing one against the other. Remarkably, the other numbering he planned to explore was the numbering used by Subject 2. The preferred numberings either move through the board as if reading English (Subject 2), or trace an "M" (Subject 1's program) or a "W" (Figures 4 and 9) along its lines.

Numbering quickly led to a one-dimensional, horizontal format: Subject 1 encoded a state as an array of numbers that denote color, Subject 2 used a string of B's and W's, and Subject 3 used circles in a row. The colors of the playing pieces were also clearly irrelevant. Subject 1 shifted to red and green, Subject 2 used letters instead, and Subject 1 shifted to red and blue. As observed in Section 1, operators that generate moves quickly and correctly are essential to expertise. Although the horizontal axis is ordinarily a neutral representation [20], it fails to convey key information here: the legal moves as defined by the edges. The edges on the game board serve as a move generator. Once they abandoned the two-dimensional board, the subjects needed a way to produce all possible moves methodically from a given state. Subject 2 extracted moves from an edge list that really contains indices into a string of B's and W's. Subject 1 wrote a simple routine that served the way Subject 2's edge list did. Subject 3 memorized them as the arcs in the centers of Figures 8 and 9.

Each subject established a state syntax with little deliberation. Subject 1 appeared relatively unscathed by his. Subject 2 labored under his arbitrary syntax and, as predicted by [3], struggled to generate moves correctly. Subject 3 began with one syntax and then progressed to another that gave him a more "satisfying" diagram. In every case, the linearity of the representation appeared to have made part of problem solving (move generation) more difficult because the subjects lost an important invariant, connectedness [22]. This accounts for missing nodes, missing edges, and incorrect edges in the diagrams.

There are some indications about these subjects' internal representations of a game state. Subject 1 used the computer as a sketch pad; once his code was complete, he appears to have relied completely on the board from Figure 1 for his own move generation. Subject 2 manipulated his linear state representation comfortably, but periodically converted back to the Figure 1 representation to consider what was really happening at certain points in the space. At crucial moments in its construction, he returned to the board diagram for confirmation, trusting his visual perception more than his mechanical process. Only Subject 3 appears to have made a complete shift to a one-dimensional state representation: "At the beginning I tried to imagine the situation on [the board]. Then I just manipulated the letters and numbers."

### 3.3 Reliance on a state space representation

There is no evidence that any of these subjects was able to retain or represent an entire state space diagram in memory. Two relied upon a diagram for memory, and their state space diagrams are more complex than their state diagrams [10]. This complexity derives from the number of states in the space, the introduction of color, the orientation of states with respect to each other (particularly in the later diagrams of Subject 3), the implicit animation (the arrows in Subject 3's later diagrams), the presence of two competitors, and the uncertainty about the other competitor's response.

Subject 1, although he claimed not to have drawn at all, effectively had his program draw state diagrams for him. While he competed against it, he certainly saw on the screen far many more pictures of the board itself than either of the other two students, but any representation of the state space he had was purely internal. As a result, one can only guess at the operators available to him.

In contrast, Subject 2 was strongly influenced by Figure 2. He explored the game by expanding the tree depth-first, generating moves with his list of edges. He noted repetition in the diagram, but did not attempt to winnow it out — rather, he used it to curtail further search until all nodes had been explored. Rather than imagine the result of a move, Subject 2 represented each result on paper. As a result, he was able to generate all the states correctly, but he also produced duplicates. His challenge then became detecting repeated states. Subject 2's directed edges were an innovation not shown during the lecture; they were in direct response to that need.

Subject 3 began in the style of Figure 2, but soon worked breadth first, that is, level by level. In addition, he became interested in the shape of the space more than its contents. He used repetition, both of states and of piece configuration, as a way to tease out all possibilities. His color-coded operations were player-dependent moves.

A diagram about play must include the moves, the player making them, and the ultimate outcome. Subject 1's program, which could have portrayed movement, simply makes the next state "appear" on the screen, without any transitional sliding. Subject

2 labeled the edges to designate the mover, eschewing the state labels of Figure 2. Subject 3 colored his edges to show how pieces might move. Both Subject 2 and Subject 3 highlighted winning nodes in their diagram, imputing importance to the absence of a cycle. Subject 3 went further, aligning these highlighted nodes. Subject 1 had no tangible record of winning states, only his memory of them.

Subjects 2 and 3 also established a state space syntax. While a game tree represents cycles by endless repetition, a *game graph* links moves backwards to earlier states. Initially, Subject 2 was seriously hindered by his tree representation; it made the detection of repetition difficult, and unexpanded states were not obvious. He soon incorporated elements of a game graph, using predominantly one-headed arrows to indicate repetition. This also allowed him to retain the game tree convention that represents the mover by a node's level. Subject 3, in contrast, soon moved from a game tree to a game graph, but with two-headed arrows. He found, however, that in his new syntax he could no longer "read" the board to identify the mover. This gave rise to incorrect associations among states with the same visual properties but different movers. In response, Subject 3 confirmed correctness and completeness visually, relying on the inherent symmetries of the game. Eventually he enhanced his edge representation with the identity of the mover. Furthermore, Subject 3 repeatedly rearranged the elements of his diagrams to highlight their role in play. By his fourth diagram he had abandoned any pretense at a tree; he then sought to convey not only sequences of states in play, but also the symmetries of their values. His final diagram emphasized the number of moves available from a state, which moves are dangerous, and which states are losses. This variety of perspective is predicted in [3]. Figure 9 includes the perspective of each mover, the perspective of "safe" moves, and the perspective of winning, along with an elegant statement of symmetry.

Within the pong hau k'i state space, some states are particularly salient. Subject 2 sketched (i.e., translated from a string back to a diagram) only two: Figure 5(b) where the mover wins, and Figure 5(c) where white must win or lose. Subject 1 appears to have noticed these in passing. Although Subject 3 eventually focused on four states one ply removed from Figure 5(b), he did not recognize the significance of those delicate states until he explained his diagram verbally to a friend — producing the diagram was not as powerful as interpreting it aloud.

The subjects' state space diagrams served a variety of purposes. Subject 1 used his computer-drawn states to provide him with playing experience; those states were connected temporally but visible only one at a time. Subject 2 used his tree to check whether or not he had considered all possible states, and his occasional sketches to check his reasoning. Subject 3 used his first diagrams the way Subject 2 used his tree, and subsequently used one diagram to check another, and portions of one diagram to generate other portions. He suspects that he began to think about strategy around the same time; his "structural" concerns in his third diagram were replaced by his "semantic" considerations in his fourth.

Kaplan and Simon's result that a good representation produces a smaller space [11] is borne out by Subject 3's switch to mover-less states. There is, however, no evidence here of an effective generator that relies on invariants. Finally, none of these subjects used their diagrams to compete. Each of them believed he had internalized the most important facets of his diagram(s), and could play perfectly without reference to them. The diagrams were thus an exploratory device. Once completed, they were not essential for expert performance.

### 3.4 Beyond the assignment

As one becomes immersed in these protocols, it is difficult to remember precisely what the assignment was — these subjects turned it into individual discovery odysseys. The first part of the assignment was to "play it, take protocols on the development of your own expertise, and capture any representational shifts." All three played pong hau k'i, and were able to report on the development of their own experience to varying degrees. Subject 1 animated his numbering decision as Figure 4. Subject 2 drew his representational shift for the state diagrams at the top of the page. Subject 3 detailed the evolution of both his state diagrams and his state space diagrams.

The second part of the assignment was to "draw out the state space for pong hau k'i, label [it] with outcomes, produce a rule set to play it, and produce a good heuristic for it." Subject 1 never drew the state space, but he did realize "that the only way [to win] is to put the opponent into [two locations on the same] side," which is necessary and sufficient for expertise. Subject 2 first played a few contests, during which he realized that "it is quite difficult to lose the game." This gave him "an idea about situations [he] should avoid." His state space diagram contained a few errors, but he too had the right idea. Nonetheless, this experiment indicates that a traditional state space representation obscures the structure that gives rise to heuristics of the kind people readily develop.

Subject 3, in his search for structure, produced an increasingly symmetrical sequence of state space diagrams, eventually *planar* (without crossed edges) ones. These representations were driven by his determination to understand the relationships within the space, long after he could play expertly. His final diagram, Figure 9, classifies states by their strategic importance and the number of choices they offer, rather than by the location of their pieces. It also makes clear where important and irrelevant decisions lie. Thus his final concepts were about vulnerability and freedom, not about states at all. His success is in part attributable to his HUYGENS-like search through diagrams for regularity, and in part to the insight stemming from his search through representation space [11]. Once Subject 3 was able to abandon the mover as part of the state representation, he could attend to the relevant features.

### 3.5 Discovery and planning

This task was deliberately posed as a discovery problem. The exploratory behavior predicted by [12] was visible in the subjects' early play: "…we were placing the

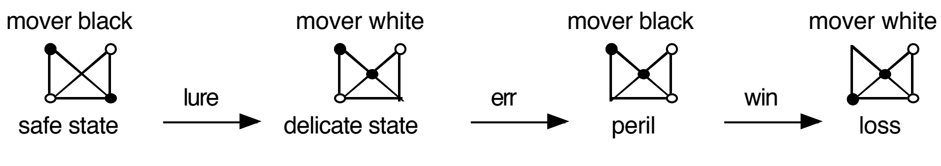


**Fig. 10.** An abstract 3-step plan as a sequence of actions, with an example from pong hau k'i.

bread without any strategy, [we] just kept [following] the rules." Discovery in a game is made more difficult by one's own inability to predict the next state where one will be forced to act.

The only way to win at pong hau k'i is to lure one's opponent into a delicate state and hope that she errs. This three-step plan (lure, err, win) relies on the fallibility of the opposition. Figure 10 descriptively labels these actions and the states they transform, each from the perspective of the mover, with an example. The only other move from the delicate state restores the contest to a neutral state. All three subjects were aware of the dangers of the delicate state, but the protocols of Subjects 1 and 2 indicate only defensive use of it. Subject 1 wrote, "I played a few games … during [which] I realized that it is quite difficult to lose the game, and I had an idea about situations I should avoid." Subject 2 wrote, "…we gave up, because I couldn't lose [either]." Only Subject 3 saw the plan's offensive and defensive aspects, once he discovered delicate states.

Subject 3's planning knowledge is closely linked to his state space diagrams. He writes of "the secure central area" and "the … terminal columns … [where] the game would either end in a terminal state or return to the secure central area … via a delicate state. He also recognized that "it is not even guaranteed that any of the four delicate states will ever be reached." Thus a contest could circle only the innermost loop.

What supported discovery here? For Subject 3, it was the elimination of the mover label on the state, and the way the states oriented themselves on the paper. For Subject 1, the program that made random moves in such a small state space likely led to every state several times. Although this is insufficient to support the development of expertise in a larger space [6], it was adequate for pong hau k'i. For Subject 2, the move generator was essential, as predicted by [12].

There is much empirical evidence that diagrammatic representations support problem solving. Cheng has identified four processes in conceptual learning that diagrams can support: observation of the space, modeling to link and generalize observations, acquisition of new concepts, and integration of those concepts [3]. He argues that a diagram to support learning should encode the laws that govern the relevant phenomena in the problem domain, and highlight differences among concepts. Subject 3's final diagrams meet these criteria.

### 3.6 Testing the representation

One hallmark of a good representation is its applicability to other, similar problems [3]. As a test of Subject 3's techniques, the author created a thought experiment: ap-

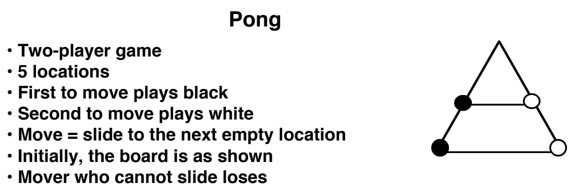


**Fig. 11.** The (invented) game of pong.

ply these methods to another simple game, called merely "pong." The rules of pong are identical to those of pong hau k'i, and its board also has five locations. Only its layout differs, as shown in Figure 11. At its conception, the author was unaware of the properties of pong's state space, but anticipated that one player could win by trapping the other on the bottom of the triangle.

Subjects in the central experiment and in [12] make clear that it is important to number the locations on the board in a way that supports move generation. Thus the "natural numbering" for pong with 1 at the top was immediately discarded for one that begins at the bottom left, goes up to the top, and down the right side. With this numbering a one-dimensional representation for the start state is BB–WW. Once again there are 30 possible board configurations and 60 possible states. The next step is to lay out the states on paper, and attempt to find a planar representation. Surprisingly, that proved more difficult; checking for repeated states requires careful bookkeeping. A list of all possible configurations proved a helpful check, as did the expectation of symmetry in the state space diagram.

When a diagram clarifies things, it becomes a tool. The pong state space diagram of Figure 12 adapts Subject 3's representation of the mover by edge color. That reduces the number of states to 30, represented in Figure 12 only as numbers (the order in which they were generated by hand). Legal play thus involves a sequence of edges that alternate in color. An algorithm to identify all reachable states from such a diagram, using spreading activation, appears in Table 1. It begins from the start state (labeled "1"), and alternately collects states reachable by edges of one color, then states reachable from those just-added states by edges of the opposite color. When the algorithm executes on Figure 12 with 1 as the start state, 10 of the theoretically possible pong states are never reached (those numbered above 20). As a result, the anticipated loss states (two pieces of the same color at the bottom of the triangle) never arise. Instead, the only realizable losses (states 15 and 20) are when a player occupies the top vertex and one of the mid-side locations. (As a matter of fact, only one such loss for each side is accessible, depending upon the first player to move. If white were to go first, 24 and 28 would be accessible instead.)

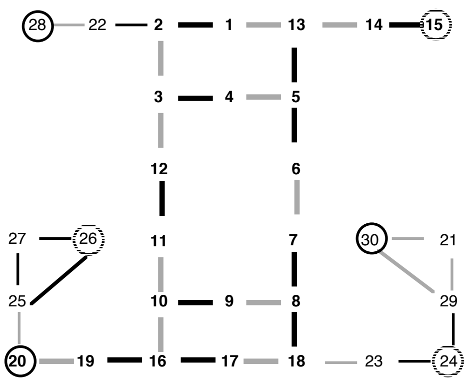


**Fig. 12.** An abstraction of pong's state space, with states numbered in the order in which they were generated. The start state is labeled 1. States reachable from 1 are in bold. Black lines denote a move for black; gray lines a move for white. Black circles denote a loss for black; gray circles a loss for white.

**Table 1.** A high-level algorithm to identify all reachable states in an edge-colored state space like that in Figure 11. The function `other` toggles between two colors. When the algorithm halts, all nodes accessible from the start state are displayed.

```
reachable ← {start-state}
fringe ← {start-state}
new-fringe ← null
hue ← first mover's color
until fringe is empty
      select node from fringe
            for each edge with color = hue from node to new-node
               unless new-node is in reachable, add new-node to reachable and to new-fringe
      fringe ← new-fringe
      new-fringe ← null
      hue ← other(hue)
display reachable
```

Could a better start state have been chosen for pong? The diagram in Figure 12 makes clear that no state and designated first mover in pong provide access to the full space. Ideally a diagram teaches the solver about the nature of the problem, that is, the person learns from the diagram she has drawn. The dangerous states in pong are 13 and 16. Since play begins at state 1 and black moves first, state 13 is entered only through a black move, where white has two choices, one of which will lead to its defeat. Similarly, the wrong decision in 16 will cost black the contest. Thus a player can make any legal move, except in the single state of concern to her. Unlike pong hau k'i, once your opponent has made an error, a win in pong requires no skill at all; it is inevitable.

This analysis is a demonstration of the power of a good diagram to highlight the key features of a problem. The author has also applied the same technique to a variety of other 5-location boards, and derived similar insights. The diagram moves one beyond the game itself, to concepts about what is possible within the game's framework.

Figure 12 also facilitates planning. As black, one wants to reach 15 which is only accessible through 13. In turn, 13 is only reachable from 5, which is itself inevitable from 1. So one must plan to move from 5 to 13 (and not to 6), hoping white will move to 14. This is another example of the (lure, err, win) plan structure. If black moves from 5 to 6 instead, white will have an opportunity at 10 to execute a similar plan.

## 4 Conclusions

Perceptual information from a diagram can provide guidance during problem solving. This experiment has shown how diagrams can be used during problem solving to develop representations, to inventory possible actions, to confirm representational shifts, to identify opportunities and dangers, and to plan. Ideally, the function of a state space diagram is to support display-based problem solving, done entirely within the context of the diagram without additional logical formalisms. Ultimately, display-based problem solving should support the extraction of simple but cogent rules for expertise, as it did here. Thus it is a tool to develop intelligent behavior, one that should eventually be exchanged for the clarity and speed of the heuristics it engenders.

In particular, AI researchers devised the game tree to facilitate their analysis of two-person games. A game tree captures turn-taking and the impact of sequences of decisions, and serves as a repository for states and their relationships. Indeed, this representation has supported the development of many sophisticated search mechanisms that now underlie champion game-playing programs. Nonetheless, in human hands, a game tree does not facilitate either the correct and complete generation of moves, or the detection of duplicates in cyclic games, both of which are crucial to expert play. Furthermore, both the "natural" extensions to game trees for duplicate states make them visually more complex. The mover-free state representation developed here is a step toward more comprehensible diagrams for games where states repeat.

Diagrams can also mislead. Each subject conserved effort by encoding the two-dimensional board in one dimension. Ironically, the linear representations they devised to simplify their task made move generation considerably more difficult. (It is, we suspect, far easier to imagine sliding along visible lines than jumping in a set of apparently disconnected leaps.) Thus they were forced to find some methodical way to generate moves: indexing into a sequence, envisioning with arrows, or writing a routine.

Diagrams are dispensable. In a relatively simple problem, like those addressed here, fortuitous induction may serve equally well. Although the state space diagram is an essential memory device, this experiment demonstrated that it does not necessarily reveal the deeper structure of its associated task, in part because of its size and complexity. To deduce deeper structure, however, some representation of state space is essential, the more compact the better, as a variant on the traditional game tree proved here. A good diagram not only captures the salient portions of a problem, it focuses attention on them. The best of these diagrams did this so well that they themselves became obsolete.

The "more is better" bias previously identified for learning tic-tac-toe and its isomorphs is valid here, but "more of what?" is the issue. There are five (non-disjoint) position categories in pong hau k'i: left, right, top, bottom, and center. The winning invariant is the number of distinct locations adjacent to those you hold. (An obvious error is to prefer more edges through the occupied locations instead. That heuristic prefers a top and a bottom location on the same side to two top locations, and leads to a loss.) Although all three subjects learned to play well, no representation specifically incorporated the winning invariant. Indeed, in some sense, the description of movement in the rules emphasizes the lines on the board, rather than the locations to which they lead. It thereby masks the winning invariant and makes the task more difficult. It would be interesting, therefore, to study isomorphs of this problem, including deformations of the essentially-square board.

One might hypothesize some process for the representational shifts that abound in these protocols, a "think-choose-try-evaluate-revise" cycle. This is supported by all the subjects' selection of a state representation, and by the third subject's intensive development of a search space diagram. In the latter, each new approach is repeatedly criticized and abandoned. This suggests that a meta-level critic ran in parallel as he drew, mediating between the visual representation and some more abstract expectations.

These protocols reveal a wealth of cognitive activity in which diagrams play a central role. The paradigmatic game tree diagram alone cannot, and did not, guarantee completeness, correctness, or well-managed repetition. Subject 1's diagrams were a kind of random travel through the space, but travel adequate to develop expertise. Subject 2's diagram served to satisfy his curiosity, not to organize his thoughts about the problem. Only with Subject 3 might one argue that diagrams led to deeper understanding of the game, an understanding unnecessary to expertise but satisfying in its level of knowledge organization. His final diagrams imposed order and a value system on a set of repetitive chronological sequences. Furthermore, the technique he developed has proved applicable to other simple cyclic games that could be applied for game-analysis and game authoring. It is a model of how diagrams support thought.

## Acknowledgements

Thanks to Boicho Kokinov for his invitation to teach at The Cognitive Science Summer School, and for attracting such thoughtful students. Thanks too to the three anonymous, enthusiastic, and cooperative subjects, who transformed a homework assignment into an intriguing demonstration of the role of spatial cognition in human expertise. This work was supported in part by NSF IIS-0328743. This paper develops and expands upon material first presented at the Diagrams 2002 conference.

## References

1. Amarel, S.: On Representations of Problems of Reasoning about Actions. In: Michie, D. (ed.): Machine Intelligence 3. Edinburgh University Press, Edinburgh (1968) 131-171
2. Anzai, Y. and Simon, H.: The Theory of Learning by Doing. Psychological Review 36 (1979) 124-140
3. Cheng, P. C.-H.: Unlocking conceptual learning in mathematics and science with effective representational systems. Computers in Education 33 (1999) 109-130
4. Cheng, P. C.-H. and Simon, H. A.: Scientific discovery and creative reasoning with diagrams. In: Smith, S., Ward, T., and Finke, R. (eds.): The Creative Cognition Approach. MIT Press, Cambridge, MA (1995). 205-228
5. Cross, N., Christaans, H., and Dorst, K. (eds.). Analyzing Design Activity. John Wiley and Sons. Chichester (1996)
6. Epstein, S. L.: Toward an Ideal Trainer. Machine Learning 15 (1994) 251-277
7. Falkenhainer, B. C. and Michalski, R. S.: Integrating Quantitative and Qualitative Discovery. Machine Learning 1 (1986) 367-401
8. Gelernter, H.: Realization of a Geometry-Theorem Proving Machine. In: Feigenbaum, E. A. and Feldman, J. (eds.): Computers and Thought. McGraw-Hill, New York (1963) 134-152
9. Hadamard, J.: The Psychology of Invention in the Mathematical Field. Dover Publications. New York (1945)
10. Hegarty, M.: Mental Visualizations and External Visualizations. In: Gray, W. and Schunn, C. (eds.): Proceedings of the 24th Annual Meeting of the Cognitive Sci-

ence Society, Lawrence Earlbaum, CA (2002) 40-45
11. Kaplan, C. A. and Simon, H. A.: In Search of Insight. Cognitive Psychology 22 (1990) 374-419
12. Kotovsky, K., Hayes, J. R., and Simon, H. A.: Why are some problems hard? Cognitive Psychology 17 (1985) 248-294
13. Langley, P., J. M. Zytkow, Simon, H. A., and Bradshaw, G. L.: The Search for Regularity: Four Aspects of Scientific Discovery. In: Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (eds.): Machine Learning 2: An Artificial Intelligence Approach. Tioga Publishing, Palo Alto (1986) 425-469
14. Larkin, J. H.: Display-based Problem Solving. In: Klahr, D. and Kotovsky, K. (eds.) Complex Information Processing: The Impact of Herbert A. Simon. Lawrence Erlbaum Associates: Hillsdale, New Jersey (1989) 319-341
15. Pascal, B.: Pensées de Pascal. Éditions Garnier Frères, Paris (1964)
16. Pearson, D. Alexander, G., C., and Webster, R.: Working Memory and Expertise Differences in Design. In: Gero, J. S., Tversky, B. T. and Purcell, T. (eds.) Visual and Spatial Reasoning in Design II. Key Centre of Design Computing and Cognition, University of Sydney, Sydney, Australia, (2001) 237-251
17. Poincaré, H.: La Valeur de la Sciènce. Flammarion, France (1970)
18. Saariluoma, P. and Maarttola, I.: Spatial Mental Content and Visual Design In: Gero, J. S., Tversky, B. T. and Purcell, T. (eds.) Visual and Spatial Reasoning in Design II. Key Centre of Design Computing and Cognition, University of Sydney, Sydney, Australia, (2001) 253-268
19. Suwa, M., B. T. Tversky, Gero, J. S., and Purcell, T.: Seeing into sketches: Regrouping Parts Encourages New Interpretations. In: Gero, J. S., Tversky, B. T. and Purcell, T. (eds.) Visual and Spatial Reasoning in Design II. Key Centre of Design Computing and Cognition, University of Sydney, Sydney, Australia, (2001) 207-219
20. Tversky, B. T.: Multiple Mental Models. In: Gero, J. S., Tversky, B. T. and Purcell, T. (eds.) Visual and Spatial Reasoning in Design II. Key Centre of Design Computing and Cognition, University of Sydney, Sydney, Australia, (2001) 3-13
21. Zaslavsky, C.: Tic Tac Toe and Other Three-in-a-Row Games, from Ancient Egypt to the Modern Computer. Crowell, New York (1982)
22. Zhang, J.: The nature of external representations in problem solving. Cognitive Science 21 (1997) 179-217