

Learning on Paper: Diagrams and Discovery in Game Playing

Susan L. Epstein¹ and J.-Holger Keibel²

¹Department 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

²Germanics Department and Graduate Programme in Human and Machine Intelligence
University of Freiburg, 79098 Freiburg, Germany
keibel@uni-freiburg.de

Abstract. Diagrams play an important role in human problem solving. In response to a challenging assignment, three students produced diagrams and subsequent verbal protocols that 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, and how the diagrams both explained and drove explanation. It also considers how this process might be adapted for a computer program.

1 The Problem

This paper recounts how three people generated and explored diagrams as they learned about a simple game. Despite identical instructions, 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. Our thesis is that visual cognition interleaves with high-level reasoning in complex ways. The primary results of this paper are a preference for one-dimensional representation in thinking about a two-dimensional space, the power of the diagram to displace the original reasoning context, and a demonstration of the synergy between visual and high-level processes during abstract reasoning.

Pong hau k'i is an old Korean game, played between black and white on the board shown in Figure 1 [1]. The *mover* (the contestant whose turn it is to move) slides her piece along a line to the single empty location on the board. By game-playing stan-

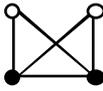
- Pong hau k'i**
- Two-player game
 - 5 locations
 - First to move plays black
 - Second to move plays white
 - Move = slide to the next empty location
 - Initially, the board is as shown
 - Mover who cannot slide loses
- 

Figure 1: As presented to the students, the game of pong hau k'i.

dards, this is a very simple game: there are only 60 possible *states* (locations of the markers plus the mover). A sequence of states is a *contest*. The state in Figure 1 is the *start state*, from which every contest begins. Contestants take turns until a mover loses because she is *trapped* (cannot move either of her pieces) or until a draw is declared because a state has repeated more than 3 times. Indeed, pong hau k'i is a *draw game*, that is, if both sides play correctly, every contest ends in a draw.

This work addresses cognition about movement in an abstract, rather than a realistic, two-dimensional space. Although there are studies of how individuals respond to a variety of representations for the same two-person game, to the best of our knowledge there are no studies of an individual's diagrammatic analysis of a game. This paper is in the tradition pioneered in [2], particularly the protocol analysis on problem reformulation from Subject 2. The first section describes the experiment that exposed a group of students to pong hau k'i. Subsequent sections describe the students' solutions and, more importantly, the processes they went through to reach them. The final sections discuss the results and the potential for automation of these processes.

2 The Experiment

Pong hau k'i was an informal assignment for students at The Cognitive Science Summer School at the New Bulgarian University in July, 2001. During a lecture they were given the rules with the slide in Figure 1, and told that it was a really simple, even boring, game. None of them had seen it before. 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 was provided.) They were encouraged to obsess over it and to report back on their experience.

That lecture included an introduction to artificial intelligence, with some discussion of the search space paradigm. One slide depicted part of the search space for another game, lose tic-tac-toe, as the *game tree* in Figure 2. In such a diagram, each *node* (picture of the board) represents a state and each *edge* (line joining two nodes) represents a move transition from one state to the next. The topmost node is called the *root*. A node that ends play is *terminal*. Later in the lecture, the students were asked 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." These directions appeared on slides.

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, (e.g. [3-6]). 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 the first 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 the empty position. Any italics are the subjects'

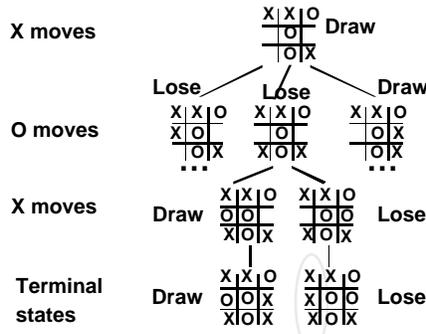


Figure 2: An example of a game tree as presented to the subjects.

own, and all names have been omitted. Prior to this paper, none of the three subjects described here had seen the protocols of the others.

3 Subject 1 — A (Relatively) Quick Response

Subject 1 is a Ph.D. candidate in cognitive psychology. Originally, he offered no verbal protocol with his analysis, only an elaborate, tree-like diagram on a single sheet of paper, with some small annotations described below. Subject 1 drew his diagram in a few hours, stopping “when I thought it was complete....” He did not use a 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.”

Subject 1’s diagram is difficult to read (and therefore not included here). It is, however, possible to deduce a good deal of his process from it. Figure 3 reproduces 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.) To the right of the start state is an encoding of it (WW-BB), from which one may deduce the 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 7 (undirected) edges along which a move may occur.

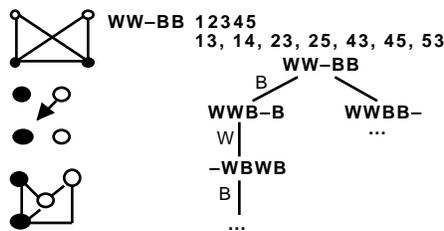


Figure 3: The top portion of Subject 1’s protocol. Ellipses are the authors’.

On the entire page there are only two other sketches of the board, both to the right. The first, Figure 4(a), appears near the two losing positions highest in the diagram, about 8 levels down in the tree. It represents the state just before a possible win for either mover. The second, Figure 4(b), appears next to a subtree placed far to the right but connected to the diagram (the *subtree-to-the-right*). It represents a state from which white, if the mover, can win; if black is the mover the contest will merely continue. These are the only sketches Subject 1 drew to check his reasoning.

Most of the page is occupied by a game tree, where each node is a 5-letter string representing the contents of the 5 locations on the board. Nodes at the same depth in the tree are horizontally aligned on the unlined paper. In most of the tree, a W or 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 are the neighbors of the start state. Subject 1 chose to follow only the move from 4 to 3, shown below and to the left of WW-BB in Figure 3. 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 its expansion, 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 undirected lines, but at some point Subject 1 noticed that some nodes repeated. When he found a repeated node (encoding and mover), he drew an unlabeled edge that was *directed* (had an arrow on at least one endpoint), as in Figure 4(c). A *backward edge* begins at one node and runs to another *higher* (closer to the root) in the tree; there are 11 backward edges. One backward edge in particular, from the left-most portion of the subtree-to-the-right, returned to the root and presumably eliminated the need to expand WWBB-. 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.

Every terminal node is either the source of a backward edge (indicating repetition in play) or is surrounded by a rectangle. 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

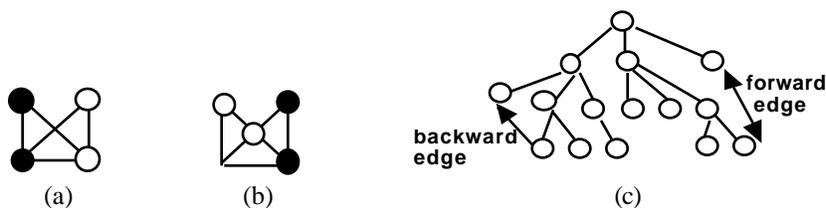


Figure 4: The other two board sketches in Subject 1's protocol: (a) the state penultimate to a loss; either contestant can win from here, (b) a state from which white, but not black, can win, and (c) orientation of forward and backward edges in Subject 1's tree.

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 1) that is fully expanded.

4 Subject 2 — An Obsession

Subject 2 is a Ph.D. student in linguistics and cognitive science with a strong mathematical background. His exploration of the game was remarkable. The day after the assignment, Subject 2 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 first five diagrams, which chronicled his path to the sixth and final one. He remembered a great deal of the process he had been through in the past 24 hours, and we numbered and annotated his early diagrams as we discussed them. From these, Subject 2 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 the diagrams, the debriefing session the day after the assignment, and the original and later versions of the written protocols. Subject 2 identifies 9 stages in his path to expertise at pong hau k'i.

Stage 1. (15 minutes) Subject 2 attempted to “get a feeling for the game and its state space.” At lunch after the lecture, several students made pieces from bread (the crusts against the bread centers) and began to play. Initially, Subject 2 played “very much at random without even knowing how a [losing] state might look....” Indeed, he hypothesized one state as a loss that 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 figures on the same (left/right) side of the game board.*” In a second contest, Subject 2 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

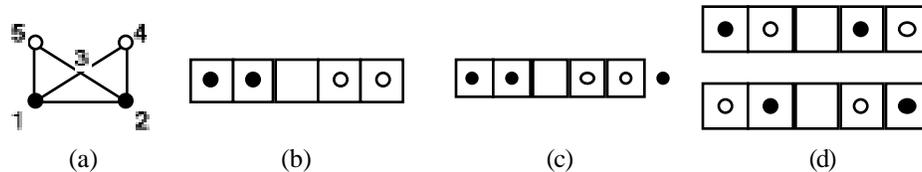


Figure 5: (a) Subject 2’s numbering of the board. (b) The starting board and (c) the start state represented in Subject 2’s rectangular coding. (d) The “seeds” for Diagram 4.

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 2 began to draw Diagram 1. 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 2 now shifted to a rectangular representation for a node. (He was drawing on graph paper, where a sequence of squares would be easy to draw clearly.) Because black moves first, Subject 2 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 5(a). Using this numbering he wrote sample nodes, such as Figure 5(b), to confirm his calculation of the possible states.

Still at lunch with his friends, Subject 2 now began to draw Diagram 2, his first full state space diagram. All nodes were in the format of Figure 5(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 2 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 2 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 2 began alone, with new determination, to draw Diagram 3 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 2 referred back to Diagram 2 regularly, checking for oversights, and frequently found them. Eventually, he extended the tree 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 representa-

tion 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 2 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 2 drew Diagram 4, beginning not from the initial state but from “another state of maximal connectivity.” He chose the states in Figure 5(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 2 then traced backward through the previous diagram to position the start state appropriately. He checked off states from Diagram 3 as he reached and reproduced them in the left half of Diagram 4, moving outward from the upper state in Figure 5(d). At the same time, he maintained elaborate tests for symmetry and edge color, for states he regarded as “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) Diagram 5 positioned the start state at the top, surrounded by empty rectangular nodes, that is, states without any playing pieces. Subject 2 was concentrating on the shape of the space now, not the moves in the game. 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 2 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.

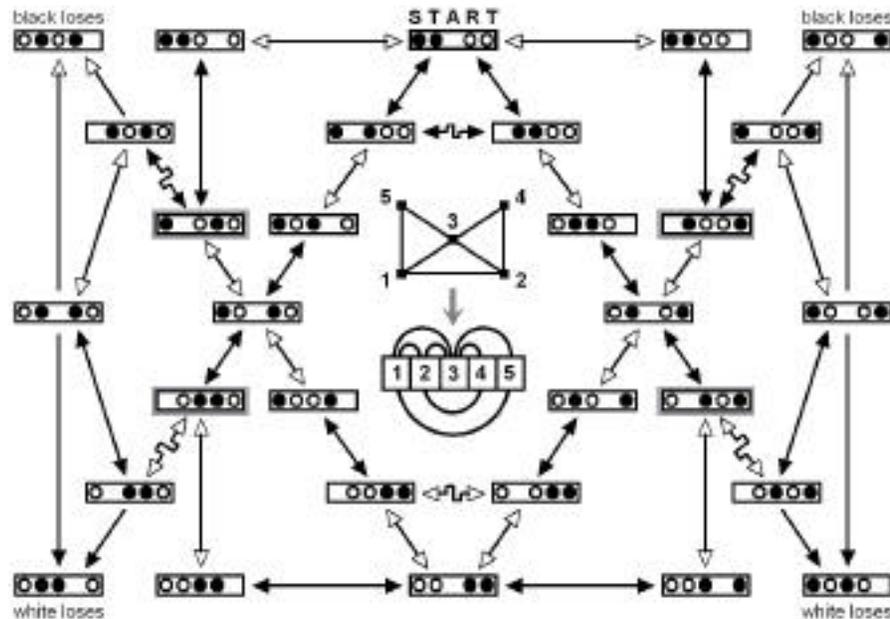


Figure 6: Subject 2’s Diagram 6, the product of 6 hours of intensive work. The four delicate states have gray borders.

Subject 2 initially drew the right side of the diagram with empty nodes and colored edges, and re-derived the nodes before confirming them with Diagram 4.

Stage 6. (duration unknown) Concerned about an asymmetry on the left in Diagram 5, Subject 2 explained the picture to his roommate, pointing out four *delicate states* 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 2 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 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 2 drew Diagram 6, shown in Figure 6. Although he also wanted to renumber the nodes, Subject 2 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 ver-

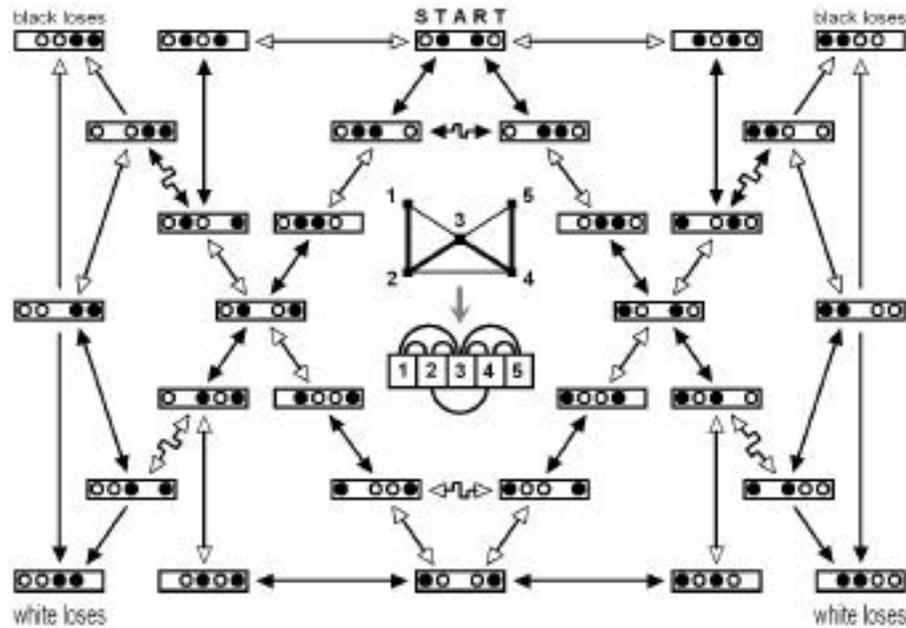


Figure 7: Subject 2’s “W-shaped encoding,” Diagram 7, with the locations renumbered as in the center of the diagram.

tically 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.”

Stages 8 and 9. During debriefing that afternoon, Subject 2 recalled his renumbering scheme from Stage 7, which ultimately removed some misunderstandings observable (but not remarked upon) during the conversation. Later that afternoon, in Stage 8, Subject 2 produced Diagram 7 with that renumbering (Figure 7), so that the patterns in nodes with the same function are symmetric to each other. Over the next few days, in Stage 9, Subject 2 revised and extended his protocol, and produced Diagram 8 (Figure 8). It classifies states and moves according to their safety, while maintaining vertical, horizontal, and through the center symmetry with respect to states and edges.

Subject 2 now felt he was an expert, and had been since Stage 5, although he had only competed once since then, as described in the next section. “...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 2 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

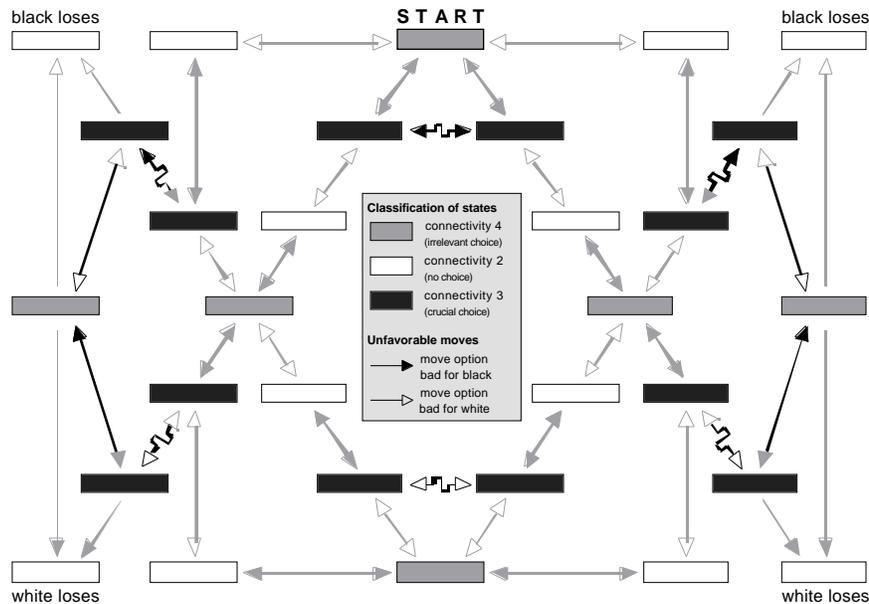


Figure 8: Subject 2's Diagram 8, showing classes of states.

light.” Nonetheless, it was a remarkable intellectual journey.

5 Subject 3 — An Automated Approach

Subject 3 is a graduate student in psychology who ostensibly chose not to draw any diagrams at all. He was among the students at lunch that day with Subject 2, but professed disinterest. Eventually the fun drew him in, and he began to play against Subject 2. “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 2 began] to draw a tree-like decision graph [Diagram 2], where every branch represented a possible move. ...The only thing I remember of his drawing was that he was thinking of a good labelling of the positions. 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 3 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. My original plan was to play with the computer” and have it count the states for him. He began to write the algorithms but then realized that the program might also perform 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 2’s] idea that a good labelling is needed. Without much thinking I came up with [Figure 9]. 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 3 had a program. The program reproduced the pong hau k’i game board, with red and blue (rather than black and white) circles for the playing pieces. The code, however, did not employ Subject 3’s numbering in Figure 9. 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 9, was 2, 1, 3, 5, 4.

Subject 3 now began to play against his program, which made random legal moves. “This was the first time, I wanted to understand the strategy.” Subject 3 soon won against his program, played “10 –30” more contests, and remained convinced that he was correct. Still later the same evening, Subject 2 appeared in Subject 3’s room.

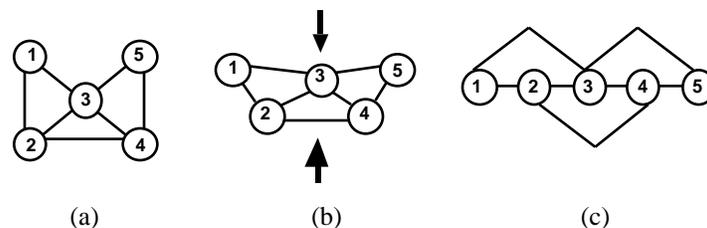


Figure 9: Subject 3’s compression of the board to create a numeric encoding, as given in his protocol. (a) The board itself, (b) the initial pressure, and (c) the collapse into a line.

“He [Subject 2] ... 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 2 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.”

6 Discussion

6.1 The Right Answers

Pong hau k’i indeed includes only 30 possible piece configurations, and 60 possible states. Subject 1 found all 30 piece configurations (several more than once), but overlooked several edges and used a backward edge to pair two different states, asserting a repetition that was not there. Subject 2’s first three diagrams included all the configurations too, but were missing edges; the others are complete. Subject 3 wrote code that was capable of drawing all possible board positions, although he only experienced the ones to which he led it to during competition.

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 2 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. When questioned weeks later, Subject 1 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 3 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 positions. [This describes both states in Figure 4.] ... 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 task was to construct a labeled diagram of the state space, a rule set, and a good heuristic. Subject 1 only constructed the space, but felt that he had completed the task. Subject 3 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 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.

6.2 The Diagrams

Although all the subjects received the same instructions, each of them approached the task differently. Subject 1 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. His directed edges were an innovation not shown during the lecture. Subject 2 began with a similar approach, but soon worked breadth first, that is, level by level. In addition, he became interested in the *shape* of the space even more than its contents. He used repetition, both of states and of piece configuration, as a way to tease out all possibilities. Finally, Subject 3, although he claimed not to have drawn at all, effectively had his program draw 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.

To describe the problem, however, all three found the physical board itself confining and soon abandoned it. Each of them began by numbering the board: Subject 1 as if he were reading text, left to right and top to bottom; Subject 2 left to right but from the bottom up, and Subject 3 with a thoughtful, visually-described process which he quickly abandoned. Subjects 2 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. Numbers quickly led to a one-dimensional, horizontal format: Subject 1 used a string of B's and W's, Subject 2 used circles in a row, and Subject 3 encoded a state as an array of numbers that denote color. The colors of the playing pieces were also clearly irrelevant. Subject 1 used letters instead, Subject 2 shifted to red and green, and Subject 3 shifted to red and blue.

Although the horizontal axis is a neutral representation [7], it fails to convey key information: 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, each subject needed a way to produce all possible moves methodically from a given state. Subject 1 extracted moves from an edge list that really contains indices into a string of B's and W's. Subject 2 memorized them as the arcs in the centers of Figures 7 and 8. Subject 3 wrote a simple routine that served the way Subject 1's edge list did.

A diagram about play must include the moves, the agent making them, and the ultimate outcome. Subject 1 labeled the edges to designate the mover, rather than labeling the states, as in Figure 2. Subject 2 colored his edges to show how pieces might move. Subject 3's program, which could have portrayed movement, simply makes the next state "appear" on the screen, without any transitional sliding. Both Subject 1 and Subject 2 highlighted winning nodes in their diagram, imputing importance to the absence of a cycle. Subject 2 went further, aligning these highlighted nodes. Subject 3 had no tangible record of winning states, only his memory of them.

The subjects' diagrams served a variety of purposes. Subject 1 used his tree diagram to check whether or not he had considered all possible states, and his occasional sketches to check his reasoning. Subject 3 used his computer-drawn diagrams to provide him with playing experience; those states were connected temporally but visible only one at a time. Subject 2 used his first diagrams the way Subject 1 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 Diagram 3 were replaced by his "semantic" considerations in Diagram 4.

Some states were particularly salient. Subject 1 sketched (i.e., translated from a string back to a diagram) only two: Figure 4(b) where the mover wins, and Figure 4(a)

where white must win or lose. Subject 3 appears to have noticed these in passing. Subject 2, however, eventually focused on four states one play removed from Figure 4(b), and did not recognize the significance of those positions until he explained his diagram verbally to a friend — producing the diagram was not as powerful as interpreting it aloud.

Each numbering of the locations on the board drives exploration of the search space differently, but there are 120 possible numberings. It is likely more than coincidental that Subject 2's second numbering, for Diagram 7, matched Subject 3's. Indeed, Subject 3'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 1.

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, and, once completed, not essential for expert performance.

6.3 Reproducing This Behavior in a Program

For a game as small as this, it is not difficult to write a program that searches exhaustively, noting potential cycles. In that way, a program could generate its own lookup table to provide a correct move from every state. From such a table, moreover, it is easy enough to extract correct (i.e., safe) moves and even delicate states — otherwise, any random move will do.

Nonetheless, pong hau k'i presents a variety of challenges for cognitive modeling. For example, a game-learning program called *Hoyle*, known for its prowess on far more difficult games, developed basic skill at pong hau k'i quickly. After observing 5 error-free contests played by others, Hoyle never lost a contest at pong hau k'i. An expert player should also win if the opposition errs. Given the opportunity, however, Hoyle never won a contest. Inspection of its knowledge base indicated that the program knew how a win would arise at the game, it just was not winning. Hoyle's standard approaches to diagrammatic reasoning were of no help here [8]. Although the program detected significant visual patterns formed by playing pieces and empty locations on the board [9], patterns that were associated with a win also occurred in contests that went on to a draw. Hoyle therefore saw no advantage to producing them.

The secret to computer expertise at pong hau k'i, for Hoyle at least, turned out to be the notion of delicate states and the appropriate way to exploit them. What Subjects 2 and 3 learned to do was to lure their opponents toward delicate states. An agent that expects flawless opposition will not do this, since pong hau k'i is, after all, a draw game. Given this understanding and a fallible opponent, Hoyle wins often and quickly. The key situations for Hoyle were described not by configurations of pieces but by paths in the search space, the same paths Subject 2 eventually saw and highlighted in his final diagrams.

7 Conclusions

AI researchers devised the game tree diagram 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 diagram has supported the development of many sophisticated search mechanisms that now underlie champion game-playing programs. Nonetheless, 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.

The subjects found 60 states far too many game boards to draw. Instead, each of them encoded the two-dimensional board in a single 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, as the centers of Figures 6 and 7, and Subject 1's need to sketch, demonstrate.) Thus they were forced to find some methodical way to generate moves: indexing into a sequence, envisioning with arrows, or writing a routine. Furthermore, Figure 9 demonstrates that the mapping from two dimensions to one is itself a diagram.

While a game tree represents cycles by endless repetition, a *game graph* links moves backwards to earlier states. Subject 1 quickly arrived at this device, using predominantly one-headed arrows to indicate repetition. This also allowed him to retain the game tree's convention that represents the mover by a node's level. In contrast, Subject 2 used two-headed arrows, but then found that in his new game graph 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 2 confirmed correctness and completeness visually, relying on the inherent symmetries of the game.

One might hypothesize some process for the representational shifts that abound in these protocols, a "think-choose-try-evaluate-revise" cycle. Certainly, each subject first selected a representation and then marched determinedly through move generation. (Subjects 2 and 3 had enough qualms about their initial choices to explore others later; they seemed most disturbed by the way the symmetries on the board failed to match the properties of these few integers.) Subject 2 became dissatisfied and abandoned several approaches while he was engaged in them, which suggests that a metalevel critic ran in parallel as he drew. He also used the properties of the current diagram (e.g., the number of edges to a node, intersecting edges, asymmetry) to structure the next one. By Diagram 4 he had abandoned any pretense at a tree; he now sought to convey not only sequences of states in play, but also the symmetries of their values. His revisions included the state representation, the mover designation, and the larger structure.

The protocols reveal a wealth of cognitive activity in which diagrams play a central role. The game tree diagram alone cannot, and did not, guarantee completeness, correctness, or well-managed repetition. Subject 1's diagram served to satisfy his curiosity, not to organize his thoughts about the problem. Subject 3's diagrams were a kind of random travel through the space, generated by his code. That travel was enough to support the development of expertise. Only with Subject 2 might we argue that dia-

grams led to deeper understanding of the game, an understanding unnecessary to expertise, but satisfying in its level of knowledge organization. His Diagrams 6 through 8 imposed order and a value system on a set of repetitive chronological sequences.

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 anonymous referees for their suggestions, and to our three anonymous, enthusiastic, and cooperative subjects, who transformed a homework assignment into an intriguing demonstration of the role of diagrams in human cognition.

References

1. Zaslavsky, C.: Tic Tac Toe and Other Three-in-a-Row Games, from Ancient Egypt to the Modern Computer. Crowell, New York (1982)
2. Anzai, Y., Simon, H., The Theory of Learning by Doing. *Psychological Review*. 36 (1979) 124-140
3. Cross, N., Christaans, H., Dorst, K. (eds.): *Analyzing Design Activity*. Series Vol. 30. John Wiley and Sons, Chichester (1996)
4. Suwa, M., Tversky, B.T., Gero, J.S., Purcell, T.: Seeing into sketches: Regrouping Parts Encourages New Interpretations. In: J.S. Gero, B.T. Tversky, and T. Purcell (eds.): *Visual and Spatial Reasoning in Design II*. Key Centre of Design Computing and Cognition, University of Sydney, Sydney, Australia (2001) 207-219
5. Pearson, D.G., Alexander, C., Webster, R.: Working Memory and Expertise Differences in Design. In: J.S. Gero, B.T. Tversky, and T. Purcell (eds.): *Visual and Spatial Reasoning in Design II*. Key Centre of Design Computing and Cognition, University of Sydney, Sydney, Australia (2001) 237-251
6. Saariluoma, P., Maartola, I.: Spatial Mental Content and Visual Design. In: J.S. Gero, B.T. Tversky, and T. Purcell (eds.): *Visual and Spatial Reasoning in Design II*. Key Centre of Design Computing and Cognition, University of Sydney, Sydney, Australia (2001) 253-268
7. Tversky, B.T.: Multiple Mental Models. In: J.S. Gero, B.T. Tversky, and T. Purcell (eds.): *Visual and Spatial Reasoning in Design II*. Key Centre of Design Computing and Cognition, University of Sydney, Sydney, Australia (2001) 3-13
8. Epstein, S.L.: Learning to Play Expertly: A Tutorial on Hoyle. In: Fürnkranz, J. Kubat, M. (eds.) *Machines That Learn to Play Games*. Nova Science (2001) 153-178
9. Epstein, S.L., Gelfand, J., Lock, E.T., Learning Game-Specific Spatially-Oriented Heuristics. *Constraints*. 3 (1998) 239-253