Projection, Problem Space and Anchoring

David Kirsh (kirsh@ucsd.edu)
Dept of Cognitive Science, UCSD
La Jolla, CA 92093-0515

Abstract

When people make sense of situations, illustrations, instructions and problems they do more than just think with their heads. They gesture, talk, point, annotate, make notes and so on. What extra do they get from interacting with their environment in this way? To study this fundamental problem, I looked at how people *project* structure onto geometric drawings, visual proofs, and games like tic tac toe. Two experiments were run to learn more about projection. Projection is a special capacity, similar to perception, but less tied to what is in the environment. Projection, unlike pure imagery, requires external structure to anchor it, but it adds 'mental' structure to the external scene much like an augmented reality system adds structure to an outside scene. A person projects when they look at a chessboard and can see where a knight may be moved. Because of the cognitive costs of sustaining and extending projection, humans make some of their projections real. They create structure externally. They move the piece, they talk, point, notate, represent. Much of our interactivity during sense making and problem solving involves a cycle of projecting then creating structure.

Keywords: Projection, interactivity, imagination, sense making, cost structure, externalization, visual thinking, situated cognition.

Introduction

Why do people typically perform better by staring at a chess board, a tic tac toe board or a geometric proof and *project* what they might do, rather than memorize the board or proof as it is initially, then close their eyes while they think of possibilities? When subjects consider possible moves in a chess game, one popular account is that they are searching a problem space; they are exploring a *purely mental representation* of the game's states, entertaining possible actions and evaluating consequences. This way of speaking leaves unexplained the relation between the physical board that is perceived and the mental process of searching an internal representation. The two might be uncoupled. And in fact, masters rarely need the cognitive support provided by a physical chessboard. They can do all the work in their heads. So a purely mental representation seems apt for them. But less expert players do benefit from a board's presence. They interactively coordinate their projections – their simulation of what if's – with the board as they see it outside. Why does a board help them project? How?

My real concern here is with interactivity: how, when and why do people interact with their environment when making sense of situations, solving problems and so on. I present a truncated account of what I believe is a key, perhaps *the* key interactive process in reasoning and sense making: the project-create-project cycle. I believe this cycle lies at the heart of much sense making, especially problem oriented sense making. It lies, as well, at the heart of most planning and tangible reasoning. A complete analysis of these phenomena would require the simultaneous study of behavior and brain. My analysis here is confined to the fine grain of behavior, involving scrutiny of the details of what people do when they make sense and reason.

In videographic studies of people understanding such things as illustrations, instructions, models and diagrams we found that subjects typically find ways of interacting with at-hand tools and resources – often in creative ways – to help them make sense of those targets. Sometimes these sense-making actions are as simple as gesturing or pointing with hand, body or instrument, muttering while looking, marking or note taking, or shifting the orientation of the target. Sometimes they involve talking with others. When tools are placed near subjects – manipulable things such as rulers, pencils, and physical parts of models – we found subjects regularly use these as 'things to think with'. They use them to create or supplement local structure to facilitate projection and mental experimentation. This is the heart of the project-create-project cycle: use what is perceived to help you do what you can in your head – namely, try to understand things by projecting possibilities, by somehow *augmenting* what you see – then externalize part of that mentally projected augmentation so that you free up cognitive resources. This process of externalization simultaneously changes the stimulus and makes it easier to project even deeper. If tools make it easier to externalize what you are thinking, then tools are used. This cycle of projecting, externalizing, then projecting again continues as long as subjects stay focused – though as with any exploratory or epistemic process a subject may soon loop, get stuck, or run out of novel projections. Let me define some terms and properties.

Projection: The basic idea

Projection is a way of 'seeing' something extra in the thing present. It is a way of *augmenting* the observed thing, of projecting onto it. In contrast to perception, which is concerned with seeing what is present, projection is concerned with seeing what is not present but might be. It is sensitive to what is present yet sufficiently controlled by a subject to go beyond what is perceived.

In figure 1 two rather different illustrations are displayed. The first – a cartoon – requires subjects to interpret the symbolic meaning of the key elements. The image must be recast as a 'keyframe' in a narrative invented by the reader, in this case, a narrative of retirees watching helplessly as

Kirsh, D. (2009). Projection, Problem Space and Anchoring. In N. A. Taatgen & H. van Rijn (Eds.), Proceedings of the 31st Annual Conference of the Cognitive Science Society (pp. 2310-2315). Austin, TX: Cognitive Science Society.

their pension money is lost forever to inflation. If you ask subjects to tell you what precedes this image and what is likely to follow they usually offer a brief story, as if running a movie forward or backward. The account typically involves a few cartoon frames describing events leading up to the current situation, but more importantly it usually has a gloss about the meaning of the loss, which is not highly visual. In observing narrative projection we found that people rarely have an urge to mark up the picture.

In figure 1b we see a geometric illustration. It is meant to show the givens for the question: will the line extending from A through D bisect BC? To solve the problem subjects scan the figure to interpret the labeling and invariably interpret the claim by imaginatively projecting a line from A through D and into BC. They may estimate magnitudes, such as whether the projected line cuts BC into equal lengths, or whether BD splits angle ABC in two. At some point, though, they are likely to reach for their pencil to add constructions or annotations to make their mental projections physical. They mark angles as equal, segments as bisected, lines as parallel or perpendicular, triangles as congruent. This interactive process of projecting and marking continues until they solve the problem or they clutter the figure so badly that they cannot keep in mind what goes with what. At that point they erase dead ends and backtrack to an earlier point in their search process.

Both illustrations 1a and 1b involve projection of meaning, but in 1b subjects also seem to be augmenting what they see. I am concerned in this paper with this second type of sense making, where subjects project quasi-perceptual structure – imagery of sorts – onto the target as if marking it up. The first example, framing and embedding the image in a narrative, is a worthy sense-making topic; but my focus is on projection that is more perception-like. It is a form of projection that typically leads to physically creating new structure.

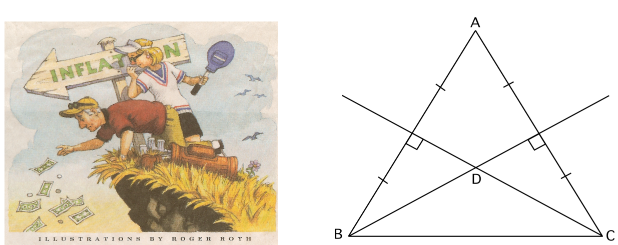

1a 1b.

Figure 1. Two illustrations: 1a is a narrative illustration requiring the viewer to make a sensible story out of the image. It involves identifying narrative worthy elements and interpreting them. For instance, the birds are not part of the narrative content but the lost money is. 1b is a figure showing a few geometric constructions. In one type of math problem, subjects are given a linguistic statement of a problem; they convert the key premises into visible shapes, and then using the figure as an aid they prove certain truths, such as that a line through AD will bisect BC. The diagram is used to clarify the givens and support inferences and allowable augmentations (new constructions). Both phases of diagramming – the conversion of linguistic to visual form, and the construction of additional lines and property labels – involve projection. In the first phase, before a subject inks the figure in the first place s(he) usually formulates a partial plan concerning where to draw the lines and how they will look. In the next phase, conjectures are often tested by projection, as if 'seen' through augmented reality, before pen is once again taken to paper.

Relation to Perception and Imagination

Projection, perception, and imagination lie on a continuum of stimulus dependence, with perception being the most dependent and imagination the least.

Perception is strongly dependent on the physical stimulus it is about. We cannot see what is not there. Even on those occasions where we have a perceptual experience of something that is, in fact, not there, such as the illusory edges shown in the first portion of figure 2, the experience is justified by the stimulus. Sometimes, real objects do produce that very effect. For instance, a solid white triangle occluding a black edged triangle would create the illusory edges shown. So the presence of perceptual mistakes and illusions is consistent with perception being stimulus dependent. Our perceptual system has been designed to recover real structure. It is tightly coupled to the outside.

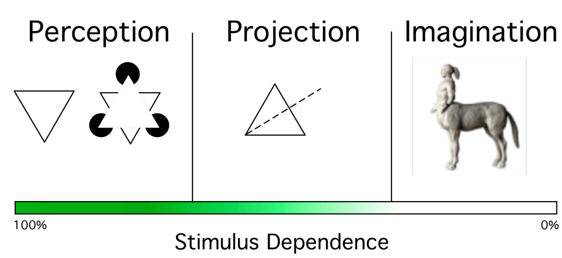

Figure 2. Perception, projection and imagination differ in their dependence on external stimuli. Perception is meant to be stimulus dependent – our perceptual systems were designed to perceive what is there. Projection is anchored to stimuli but not 100% dependent on them. We project onto external structure but what we project is not yet there. Imagination is not anchored, good imagers are able to produce vivid images unconnected to what is present and manipulate them voluntarily.

Projection is also dependent on present stimuli but much less so than perception. The coupling is looser because projection offers a peak into the possible, into what could be there, or what might be useful if it were there, but is not. It is like wearing augmented reality glasses. But with one difference. In staring at a chessboard and seeing how a knight might move, a subject must mentally remove the knight from its current spot. The layout must be changed as well as augmented.

It is an open empirical question how much external structure must be present and for how long it must persist for it to serve as the substrate of projection. But there has to be something present to project onto, something to *anchor* projection. (cf. Hutchins 2005). To grant the mind the flexibility it constantly reveals, the anchoring structure need not be persistent – we can anchor a thought in a person's gesture or in the direction a person points. But whether persistent or ephemeral there must be some external structure present, else there is nothing to distinguish projection from pure imagery. Whether that structure is enduring enough to provide a stable understructure to support repeated projection – as a chessboard does for our projection of possible moves – or whether it has a fleeting presence, triggering a single projection but then is gone, the stimulus enabling the projection in both cases has a reality outside the agent.

What can be said about the coupling between projection and substrate, between projection and anchor? First, projections are most often momentary. Some may persist in mind for minutes but usually they do not. Visual markers nicely demonstrate this temporal distribution. A visual marker or FINST, as Pylyshyn (1994) calls them, is the internal counterpart of a physical indicator laid down to mark an object or location. FINST's help us keep track of multiple objects in the visual scene. If two bees, similar enough to be perceptually indistinguishable, were buzzing around our den, and we wanted to keep an eye on each, our visual system would conjure up two FINSTs to mark the bees, thus providing the extra structure needed to keep the two distinct. FINST's of this sort have enough mental persistence to provide the stability needed for referential thought. "I wonder if that bee is going to land beside the other?" But once the tracking task ends the FINST's are released. And more often than not the need for simultaneously tracking two objects in one's nearby visual environment is not long lived. So if FINST's are representative projections they do not endure long.

In chess, the duration of projection too may vary, though the structure projected is more complex than a FINST. Most often, if a novice relies on a board to facilitate projection the projected move is considered briefly, then rejected. The ones that are not summarily dismissed form the first step in a chain of moves, so they must persist at least as long as the time it takes to create the chain. In both cases, despite their differences, projection relies on external structure being present. This contrasts with chess masters who do not need the board to think about moves. They operate more in a 'virtual reality' of their own making rather than the 'augmented reality' I am introducing as the mark of projection. And their board is conceptualized and chunked to a much greater degree.

Imagination better describes the chess master's mental activity. Their representation of the board and current situation is completely sustained internally and they have control over what they imagine next. Imagination is often defined as "a mental representation of a nonpresent object or event" Solso (1991 p.267). In psychological accounts of imagery (e.g. see Denis, 1991; Kosslyn, 2005) mental images have two primary dimensions: vividness and controllability. Vividness refers to the clarity, "sharpness" or sensory richness of an image (Richardson, 1999). Controllability refers to the ease and accuracy with which an image can be mentally transformed or manipulated (Kosslyn, 1990). We may assume that masters have both vivid and well-controlled images of chess situations.

Externalization: part of the project-create cycle

Externalization is a way of taking information or mental structure generated by an agent and transforming it into epistemically useful structure in the environment. It is a way of materializing structure that first was mental – it is the create part of the project-create-project cycle.

Externalizations are everywhere: annotations, notes, constructions in geometry, gestures, utterances, encoding order in layout, (Kirsh 1995, 2008) etc. Often the action of externalizing alters the information or projection in useful ways. This is a key factor in thinking with things, in knowing what you are thinking by seeing what you are saying, and so on. Externalizations may leave persistent traces, as in annotations or rearrangements, or they may be present only during the externalization process, as when someone gestures or talks while thinking.

Externalizations always serve an epistemic function. But they also may have pragmatic consequences too. A chess move is at once an externalization of an inner projection and a move in the game. And of course there are other actions that change the environment in epistemically useful ways that are *not* externalizations: registration of maps, turning on the news channel, etc. These are actions that alter the epistemic landscape of activity but they do not bear the right relation to internal activity to qualify as externalization, and they are not part of the project-create-project cycle.

When does external structure help performance?

Externalizing a mental projection allows a person to release at least some of their working memory, replacing it with perceived structure. So it serves as an effective interactive strategy for increasing mental power. The value of this interactive method is easy to appreciate when the structure being created is something like a construction in a geometric proof. A construction typically starts out first as a mental projection and then, if it seems fruitful, is materialized by marking the illustration. But when is structure *necessary* to improve performance? Some chess masters can play equally well with eyes closed. When does externalizing help?

For example, would staring at a blank tic tac toe board while calling out moves help performance? A blank board does not carry any state information. How could it help?

To answer that question we ran a few simple experiments, video'ed and analyzed performance to see how behavior and cognitive strategy differs when a board is present from when it is not. We used a 3 by 3 tic tac toe board first, then we scaled the game to a 4 by 4 board to see if the complexity of the game affected the value of external structure. Our conjecture was that having an empty tic tac toe table would help in both cases.

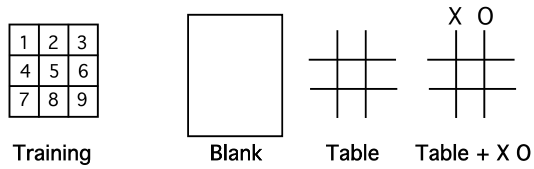

Figure 4. Training image and the three conditions.

Experiment One. **Procedure**: In figure 4 the training stimulus and all experimental conditions are shown. Experimenter and subject took turns calling out numbers 1 to 9 corresponding to the cells in a 3 by 3 board. Subjects could *not mark* their paper but could gesture if they wished. The goal was to get three in a row defined in the classical tic tac toe manner. Subjects were given an initial training period during which they mastered the translation of number to position on the table. A within subject design was used. Each subject played in each of the three conditions: blank, table, and table + XO. There were three games to a condition, three conditions to a block and two blocks to an experiment. All conditions were counterbalanced within and between subjects to control for order effects and microgenetic learning. In the table condition, subjects were given a sheet with a blank tic tac toe table to view if they so wished. In table + XO a similar tic tac toe table was given to subjects but with the letters X and O above it. In the blank condition subjects were given a blank piece of paper to look at. This was meant to serve as the imagination case – the unanchored problem space case. During our pilot study with 7 subjects we found that performance varied considerably among subjects. In particular, there were a few subjects who regularly did best on the blank condition. During debriefing it was apparent they had good imagery abilities. So all pilot subjects were called back and given a standard imagery test: vividness of visual imagery questionnaire (VVIQ-2) by D. Marks (1995). In the end 27 subjects were run and all tested with VVIQ-2.

Results. As shown in Table 1 the mean time to make a move was relatively close in all conditions and statistical tests showed no significant differences between conditions. Apparently, seeing a table does not help in 3 by 3 tic tac toe. When we divided the subject pool into strong visualizers – the upper 33% of our VVIQ scores - and weak visualizers – the bottom 33%, there were differences in means but none that were statistically significant. We also checked for order effects, to see if subjects showed significant learning during the experiment. None was noted. Nor were their significant order effects (microgenetic learning) among strong and weak visualizers. Other individual differences were more suggestive, however. Fully half of our subjects actually did better on the blank condition than the table condition. This was significant (p=.002). This difference does not correlate with visualization ability.

Discussion. Coming into the 3 by 3 case, and on the basis of our pilot data, we assumed that staring at an empty tic tac toe table would help subjects - at least weak imagers - because we thought an empty table would function as an aid to memory. Without a table an agent must remember the

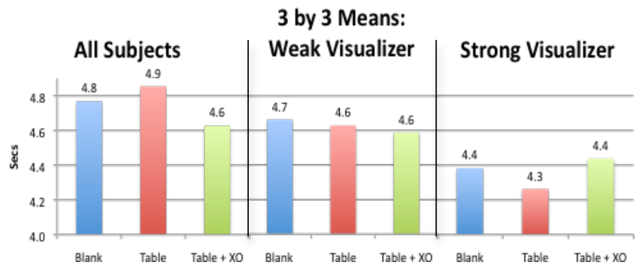

Table 1. Mean performance 3 by 3. Shorter is better. Differences are not significant.

structure of the table as well as the values in all its cells. So having a table to observe ought to reduce memory load.

Apparently, our conjecture is wrong in the case of 3 by 3 tic tac toe. Overall, nothing is to be gained by projection. Imagination is just as good. Either the memory task is not challenging enough to warrant offloading memory, as it is in chess where the board and piece configuration contains a huge amount of information, or subjects are already at ceiling.

There is, however, another possibility. Projection is an expensive process. It requires anchoring imagined elements – mental X's and O's – with physical locations. There is no a priori reason why mental tic tac toe elements should easily fit the physical table subjects look at. Some might like a large table others a small one. Indeed, several subjects reported a disconnect between their imagery, or their mental imagery strategy, and the table they were asked to use. One reason there is no general effect in 3 by 3 tic tac toe, then, is that, for many subjects, the benefits of projecting may not overcome the costs. For those subjects, projection is not a good strategy.

Some support for this interpretation can be found from the surprising finding that a full 50% of our subjects actually did worse in the table condition than in the imagery condition. What might explain that other than posting a cost to projecting – a cost to anchoring? See table 2.

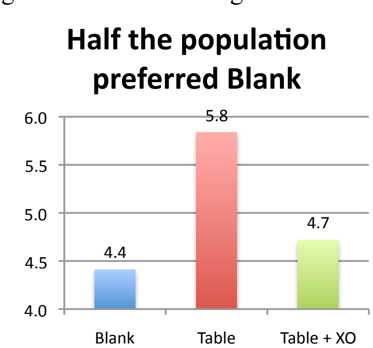

Table 2. Many subjects found it easier to play the 3 by 3 game in their imagination.

To test whether there is threshold where the urge to use external structure to support projection becomes overwhelming we ran a second experiment in which we scaled up tic tac toe to a 4 by 4 table. Intuitively everyone has a visualization limit. Once that limit is reached the cost of projection is more than paid back by the reward in memory saving, or visualization reliability, or reduced mental effort.

Experiment Two. Procedure: cells are identified with numbers 1-16 and learned beforehand. The goal of this enlarged game of tic tac toe is to be the first to get four in a row.

1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	16

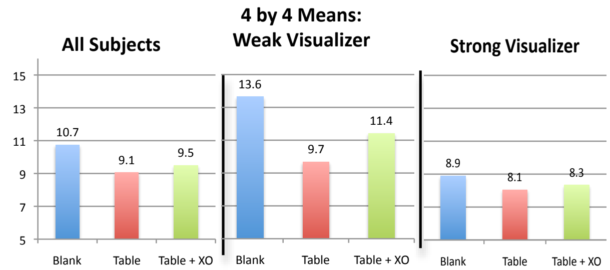

Table 3. On a 4 by 4 game the importance of having a table to work with becomes more valuable, especially for weak visualizers.

Results: Unlike the 3 by 3, all subjects in the 4 by 4 reported that external supports were helpful. But their empirical performance did not always confirm this claim. As shown in Table 3, overall performance on the table condition was significantly better than in others (Blank is slower than Table p =.03, Blank slower than XO p=.01). But we observed considerable variance in table performance among some subjects. Most significantly, strong visualizers were significantly faster than slow (p=.05). And weak visualizers showed a much larger facilitation when using tables. In fact, although strong visualizers (n=8) trended toward a preference for the table condition their improved performance was not significantly different than in the pure imagery condition.

Discussion: These findings are consistent with the hypothesis that projection has a cost that is offset once the complexity of a problem passes a threshold. The differences we observed between weak and strong visualizers suggests, further, that this threshold varies considerably with imagery capacity. Of course this does not prove that projection happens whenever the going gets rough. But it suggests that how necessary it is to project depends on both the information size of a problem, and the 'effort' an individual must expend in anchoring imaginary elements.

Several qualitative observations add to this picture. During debriefing interviews several subjects reported that they used different strategies in the table and blank conditions. When no table was present they felt overwhelmed and played defensively, using a strategy of blocking the opponent as quickly as possible rather than trying to win. Imagery alone is hard. In the table condition, however, strong 3 by 3 players – those who typically are better visualizers – initially believed they could project enough state to play offensively. They felt that with the support given them by a table they could compete with an experimenter who played with paper and pencil. Invariably subjects made errors and soon shifted their table strategy to a defensive one, and their table performance improved, soon becoming their best condition. What is interesting is that they believed in their projective ability and that they could endure the mental effort of following a harder strategy than they would consider in the pure imagery condition. When this strategy proved unreliable they fell back on using the table with a defensive strategy for greater reliability and speed. (Fewer errors were made in the table condition but not significantly so).

A second qualitative observation we made concerns the number and type of gestures made in 4 by 4 versus 3 by 3 games. It soon became apparent that the more difficulty a subject had with the tic tac toe task, even with the help of a table, the more likely they were to externalize state information to help them out, in this case with hands and fingers. Humans are ingenious at finding ways of overcoming internal state limitations. They invent methods of reducing the overall cost of performing a task, especially when the alternative is failure. They project then create structure.

For example, subject M, found a clever way of placing his fingers on the cells and the lines between the cells in the 4 by 4 table to encode more than 10 cells worth of information. Obviously he would have had far more difficulty encoding this information without the table there to 'lean' on since he would have had to project a visual structure with lines and cells 'under' his fingers.

There is much more to be said here concerning the nature of coding with hands and gesture and the timing of these interactions. But it may be more worthwhile tying this study back to the question of how people use projection to make sense of diagrams such as visual proofs and illustrations of mechanical systems.

The idea I am exploring, is that projection is related to perception, perhaps continuous with it, though it cannot be identical with perception because it is directed at augmenting the world. You cannot see what is not there. Yet in some theories of perception, most notably enactive theories, O'Regan and Noe (2001), Noe (2004) perception already contains a component of 'seeing the future'. For instance, when we see an object we do not literally see its back, but knowing it has a back is part of our perceptual experience. A more mechanistic or computational way of putting this is to say that when we perceive an object we simultaneously activate or prime a constellation of sensations we would experience if we were to move to the right or left, manipulate the object, saccade to the top, and so on. Projection is like perception, understood in this special way, because it is a process of increasing the priming level of some of the things we would see if we were to act in certain ways. The relevance to tic tac toe and to geometry is that because we are able to label a tic tac toe table by writing on it, we have a weakly primed version of the labeled table already in mind. The stronger our disposition to add labels the stronger is the priming for seeing a table augmented with those labels. Projection is a way of intentionally increasing the level of primed states of the world. It lets us entertain what the world would be like if we did act to make it so.

Now consider the visual proof shown in figure 6a that the sum of $\frac{1}{2^n}$ converges to 1. As you immerse yourself in the proof do you feel you are recreating a progression of cuts? Do you see that the operation of halving a square whose sides are 1 by 1, and then halving the remainder (whether it be a square or a rectangle) is a recursive process that will never yield a structure larger than the original square? I contend that this quasi-simulation of cutting is a form of projection and lies at the heart of making sense of visual proofs.

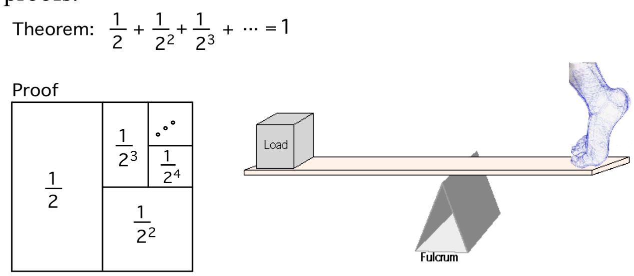

6a. 6b.

Figure 6. Two different types of visual reasoning are at play in 6a and 6b. 6a shows a geometric proof and requires understanding the recursiveness implicit in cutting regions in half. 6b requires physical understanding of the effect of pushing down on a lever.

Now look at figure 6b. Is more force or less force required to lift the load when the fulcrum is moved closer to the foot? How did you find the answer? By mentally moving the fulcrum and then simulating the consequences for the foot? This projected animation cannot be perception because, presumably, perception requires that what you see is, in fact, there to be seen. You cannot see the future. Yet there is something perception-like in this projected imagery even if it is not nearly as vivid as perception. Again projection seems to lie at the heart of our sense making visual thinking here. As we found with tic tac toe, our proxy for thinking in abstract problem spaces, projection lets us probe problems by tying our thinking to external structure. It lets us anticipate how the world might be, when we act on it.

Conclusion

I have been arguing that projection is a basic cognitive capacity involved in visual thinking, in much problem solving and in making sense of illustrations, diagrams, and many types of planning and reasoning situations. Projection differs from imagination and imagery in being anchored to visible structure. When we project it is like wearing augmented reality glasses: we lay structure over existing structure. There are no doubt other modalities of projection beside vision, but I have not considered these here. An experiment was presented in which subjects were tested to see if they performed better when there was more structure of the right sort to anchor their projections. As predicted, when useful anchoring structure is present subjects score more highly. The usefulness of such anchors depends on a subject's imagery ability and the complexity of the problem. For tasks that are simple relative to a subject's imagery ability, external anchors are of no value. Projection is replaced by imagination. But as a task increases in complexity, projection and anchoring becomes important even for good imagers. At some point everyone benefits from external structure. The costs associated with projecting and anchoring are offset by the returns derived from mental ease, memory saving and reliability.

The relevance of this to problem solving and our opening question concerning chess players should be obvious: masters who can play chess in their imagination do not need to look at a board when playing. They have so overlearned the chess board and possible configurations that they can play equally well with or without a board. But less practiced players need a chessboard and pieces. They cannot sustain a meaningful problem space for chess without the help of perception to provide anchors to a real board. Their problem space is more a projection than autonomous mental space. If anchors do not exist they must create them.

References

Denis, M. (1991). *Image and Cognition*. New York: Harvester Wheatsheaf.

Hutchins, E. (2005). Material anchors for conceptual blends. Journal of Pragmatics, Vol. 37, No. 10. pp. 1555-1577

Kirsh, D. *The Intelligent Use of Space.* Artificial Intelligence, Vol. 73, Number 1-2, pp. 31-68, (1995).

Kirsh, D. (2008). Problem Solving and Situated Cognition. In Robbins, P. and Aydede, M. (eds.) *The Cambridge Handbook of Situated Cognition.* New York: Cambridge University Press. 264-306.

Kosslyn, S.M. (2005). *Mental Images and the Brain.* Cognitive Neuropsychology (22) 333-347.

Kosslyn, S.M. Flynn RA, Amsterdam JB, Wang G. (1990). *Components of high-level vision: a cognitive neuroscience analysis and accounts of neurological syndromes*. Cognition. 34(3):203-77.

Marks, D.F. (1995). "New directions for mental imagery research". Journal of Mental Imagery, 19, 153-167.

O'Regan, J. Kevin and Noe, Alva (2001) *A sensorimotor account of vision and visual consciousness*. Behavioral and Brain Sciences. (24) 939-973

Noe, Alva (2004). *Action in Perception.* MIT Press.

Pylyshyn, Z.W. (1994). *Some primitive mechanisms of spatial attention* .Cognition, 50, 363-384.

Richardson, J.T.E. (1999). *Mental Imagery*. Psychology Press: Hove, U.K.

Solso, R. L. (1991) Cognitive Psychology. 3rd ed. Allyn & Balon