Goel, V. (1992). Ill-structured Representations for Ill-structured Problems. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.

# "Ill-Structured Representations" for Ill-Structured Problems

**Vinod Goel**
Institute of Cognitive Science
University of California, Berkeley
goel@cogsci.berkeley.edu

## Abstract

While the distinction between well-structured and ill-structured problems is widely recognized in cognitive science, it has not generally been noted that there are often significant differences in the external representations which accompany the two problem types. It is here suggested that there is a corresponding distinction to be made between "well-structured" and "ill-structured" representations. Such a distinction is used to further differentiate diagrams into finer-grained types, loosely corresponding to sketches and drafting-type diagrams, and it is argued that ill-structured, open-ended problems, like the preliminary phases of design problem solving, *need* "ill-structured" diagrammatic representations. Data from protocol studies of expert designers are used to support the thesis.

## Introduction

Cognitive science has made considerable progress in understanding how certain well-structured problems[1] are solved and the role external representations play in such solutions (Newell & Simon, 1972). A typical example of a well-structured problem is the game of chess. In chess, the start state is specified, as is the goal state and the set of legal transformations (though generating or selecting the "best" transformation at any given point is a non-trivial task). The representation of the task (whether it be in internal or external memory) is also "well structured" in that it is clear what state is being instantiated, by virtue of what it is that state, what states of affairs are being referred to, and what the set of legal transformations from one state to another state are.

While it has been frequently noted that many of the problems that we confront in life are not well structured, it has generally not been appreciated that the representations which often accompany the solutions to such problems are also not "well structured" in the above sense. In fact, the predicates 'ill structured' and 'well structured' have not to my knowledge been applied to representations. The major differentiating criterion for representations has been, and continues to be, the diagrammatic (or pictorial or imagistic) and propositional (or linguistic) dimension (Kosslyn, 1981; Larkin & Simon, 1987; Pylyshyn, 1981; Simon, 1972).

While I do not deny the importance of the diagrammatic and propositional distinction, I will here focus on and argue for an "ill-structured" and "well-structured" distinction. In fact, I will use the "ill-structured" and "well-structured" distinction to further differentiate diagrammatic representations into finer-grained types. Informally, and as a first pass, one might understand the distinction between "ill-structured" and "well-structured" diagrams in terms of the difference between fast freehand sketches and formal, box-like drafting diagrams, where sketches correspond to the former, and drafting diagrams to the latter.[2] A more formal statement follows.

The goal of this paper is to differentiate "well-structured" diagrammatic representations from "ill-structured" diagrammatic representations and to show that some ill-structured problems require "ill-structured" representations to prevent premature crystallization of ideas and to facilitate the generation and exploration of alternate solutions. This is a very brief summary of work reported in full elsewhere (Goel, 1992a; Goel, 1992b).

## Differentiating "Well-Structured" & "Ill-Structured" Diagrams

Consider the diagrammatic representation in Figure 1. It depicts two states and a transformation in a game of chess. The representation and the symbol system it belongs to have the following seven properties:

(p1) *Syntactic Disjointness:* Each token belongs to at most one symbol type. Thus for example, no tokens of the type 'rook' belong to the type 'queen'.

(p2) *Syntactic Differentiation:* It is possible to tell which symbol type a token belongs to. So given the types 'queen' and 'rook' and a token of the type 'rook', it is possible to tell which type it does and does not belong to.

[1] A well-structured problem is one in which the information necessary to construct a problem space is specified.

[2] Strictly speaking, this is not true. But it is a useful starting point.

(p3) *Unambiguity:* Every symbol type has the same referent in each and every context in which it appears. Thus no 'bishop' refers to a knight regardless of context.

(p4) *Semantic Disjointness:* The classes of referents are disjoint; i.e., each object referred to belongs to at most one reference-class. So, for example, no pawn belongs to the class of rooks.[3]

(p5) *Semantic Differentiation:* It is possible to tell which class a particular object belongs to. Thus, given a king and two classes of objects, one could determine which class, if any, the king belongs to.

(p6) The rules of transformation of the system are well specified. Thus, for example, there is no question as to what does or does not constitute a legal move for a bishop.

(p7) The legal transformations of the system are such that these properties are preserved at each and every state.

The first five of these properties (p1-p5) are adopted from Goodman (1976). The reader is referred to Goodman (1976), Elgin (1983) and Goel (1992b) for a more complete discussion.

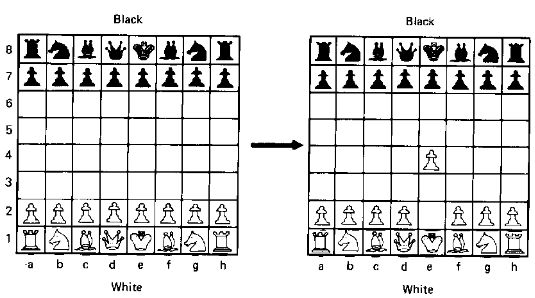


Figure 1: States and transformation from a game of chess (reproduced from Rich (1983, p.65))

Six of these seven properties of symbol systems are actually presupposed by the notion of a computational problem space (Goel, 1991b). The satisfaction of properties p1 and p2 is necessary for there to be a discernable fact of the matter as to what state is being instantiated. Satisfaction of properties p3 and p5 is necessary for there to be a discernible fact of the matter as to what state of affairs is being referred to.[4] Property p6 is necessary to constrain the class of possible transformations, while property p7 is necessary to maintain the above properties during the transformation of one state to the next.

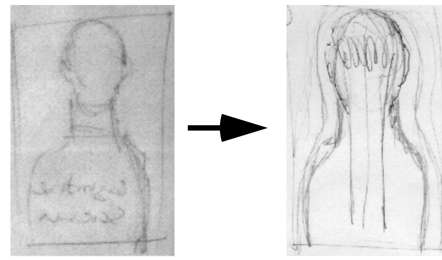

a b

Figure 2: States and transformation from early part of a graphic design session

In contrast, consider the diagrammatic representations in Figure 2 extracted from the early part of a graphic design problem-solving session. They belong to the symbol system of sketching and differ from the representations in Figure 1 with respect to each of the above seven properties. In fact, they fall on the opposite extreme with respect to each of the seven properties (p1-p7):

(p1') *Failure of Syntactic Disjointness:* Each token may belong to many symbol types at the same time. That is, in the absence of any agreement as to the constitutive versus contingent properties of tokens, there may be no fact of the matter as to which equivalence-class they belong to. Thus for example, what equivalence-class does token *a* in Figure 2 belong to? Do tokens *a* and *b* belong to the same equivalence-class? There may be no agreed-upon answers to these questions.

(p2') *Failure of Syntactic Differentiation (through density):* Because the symbol system of sketching allows for a dense ordering of symbol types (i.e., between any two types there is a third), it is not always possible to tell which type a token belongs to. So, for example, even if we agree that the token *a* in Figure 2 belongs to only one equivalence-class, it may not be possible to tell which of several classes it does or does not belong to.

(p3') *Ambiguity:* Symbol types do not have the same referent in each and every context in which they appear. For example, the token *b* in Figure 2 was interpreted as a human head and later reinterpreted as a light bulb.

(p4') *Failure of Semantic Disjointness:* The classes of referents are not disjoint; i.e., each object referred to may belong to many reference-classes. So, for example, the human figure referred to by the symbol-type *a* may belong to the class of humans and the class of students.

(p5') *Failure of Semantic Differentiation:* The system of sketching allows for a dense ordering of reference-classes. When this is the case, it is not possible to tell which class a particular

---

[3] This is, of course, true only in the vocabulary of chess, narrowly defined. In the larger context, a pawn also belongs to the class of chess pieces, the class of material objects, etc. But this is consistent with the point being made here.

[4] Property p4 is necessary to go from the referent, to the referring state. But it is not clear whether this is *necessary* for the notion of a problem space.

object belongs to. For example, in a perspective drawing of a human figure, every height of the figure would correspond to a different class of heights of humans in the world, and these classes of heights are of course densely ordered. In such a case it would not be possible to tell which class a particular human height belongs to.

(p6') The system of sketching has no well-specified rules for transforming one state into another. There is no transformation of the token *b* which would be "incorrect" or "illegal."

(p7') As the properties p1-p6 are not present to begin with, they are not preserved in the transformation of the system from one state to the next.

Having defined "ill-structured" and "well-structured" representations as such, I will henceforth dispense with the scare quotes.

It makes sense that the representations which underlie well-structured problems (e.g. cryptarithmetic, Moore-Anderson task, Tower of Hanoi, 8-Puzzle problem, checkers, etc.) should be well structured (by virtue of having properties p1-p7). After all, if these properties were absent, then the states, operators, and evaluation functions could not be specified and the problem would not be a well-structured problem.

There is, however, no similar reason for ill-structured problem spaces to be accompanied by representations belonging to well-structured symbol systems. The fact that the problem spaces are ill structured means that states, operators and evaluation functions are not defined thus, there is little need for the information which specifies them to be actually present. In fact, not only is there no compelling reason for representations accompanying ill-structured problems to be well structured, there actually seems to be a case to be made to the effect that they *need* to be ill structured to facilitate certain cognitive processes. This point is argued for in the next sections with some results from design problem solving.

## The Role of Ill-Structured Diagrams in Design Problem Solving

Two empirical studies of design problem solving were conducted. The first examined some of the cognitive processes involved in design problem solving while the second focused on the impact on these cognitive processes when the symbol systems the designers were allowed to use were manipulated along the well-structured and ill-structured dimensions.

One result of the first study (Goel, 1991b; Goel & Pirolli, in press) was that the development of design solutions has several distinct phases. Four of these phases are problem structuring, preliminary design, refinement, and detailing.[5] These phases differ with respect to the type of information dealt with, the degree of commitment to generated ideas, the level of detail attended to, and the number and types of transformations engaged in.

What is of interest to us here is the contrast between the preliminary design phase and the refinement and detailing phases. Preliminary design is a classical case of creative, ill-structured problem solving. It is a phase where alternatives are generated and explored. This generation and exploration of alternatives is facilitated by the abstract nature of information being considered, a low degree of commitment to generated ideas, the coarseness of detail, and a large number of lateral transformations. A lateral transformation is one where movement is from one idea to a slightly different idea rather than a more detailed version of the same idea. Lateral transformations are necessary for the widening of the problem space and the exploration and development of kernel ideas (Goel, 1991b).

The refinement and detailing phases are more constrained and structured. They are phases where commitments are made to a particular solution and propagated through the problem space. They are characterized by the concrete nature of information being considered, a high degree of commitment to generated ideas, attention to detail, and a large number of vertical transformations. A vertical transformation is one where movement is from one idea to a more detailed version of the same idea. It results in a deepening of the problem space (Goel, 1991b).

It was also noted that the preliminary design phases were accompanied by ill-structured representations (belonging to the symbol system of sketching), while the refinement and detail phases were accompanied by more well-structured representations, belonging to the system of drafting (Goel, 1991b). A second study was conducted to investigate the role played by ill-structured representations in the preliminary design phase.

In the second protocol study the following four of the seven properties of ill-structured representations were examined and manipulated: (p1') failure of syntactic disjointness, (p2') failure of syntactic differentiation, (p3') ambiguity, and (p5') failure of semantic differentiation. It was predicted that when these properties are absent (i.e., when the symbol system is well structured) the number of lateral transformations is likely to be hampered. The underlying rationale was that properties p1' and p3' facilitate lateral movement by allowing multiple interpretation of symbol-types while properties p2' and p5' facilitate lateral movement by allowing for overlapping (or closely ordered) symbol-

[5] While these categories may seem trivial, they do constitute a significant claim about the design problem space because they are not found in at least some nondesign problem spaces (Goel, 1991b; Goel & Pirolli, in press).

types and ideas (Goel, 1991b; Goel, 1992b). The balance of the paper describes this study.

***Subjects & Design:*** Nine expert designers from the disciplines of industrial design and graphic design were engaged in two (one-hour) problem solving sessions while the symbol systems they were allowed to use were manipulated along the dimensions of ill-structured and well-structured representations. In the one case subjects were allowed to use an ill-structured symbol system with properties p1', p2', p3', and p5'. In the other case they were requested to use a well-structured symbol system with properties p1, p2, p3, and p5.

***Manipulation of Symbol Systems:*** The manipulation of symbol systems was through the manipulation of drawing tools and media. In one session each designer was allowed to use the tools, media, and symbol systems of his/her choice. They invariably chose to use paper and pencils and did a lot of sketching. In the second session they were requested to use a computational interface. Specifically, they were asked to use a subset of the drawing package MacDraw[6] (version 1.9.5; with the freehand tool turned off and the grid turned on) running on a Mac II[7] with a large two-page monitor. MacDraw is not a sketching tool; it is a restrictive subset of a drawing or drafting tool. It only allows one to make precise lines, boxes, and circles. The subjects all used sophisticated computational drawing tools as part of their jobs and so were proficient with MacDraw.

The expectations were that freehand sketching would be used to generate substantially ill-structured representations while the representations generated with MacDraw would be substantially well structured. It was also expected (as noted above) that ill-structured representations would result in more lateral transformations.

***Task Descriptions:*** There were three graphic design tasks and two industrial design tasks. The graphic design tasks were to design (i) a poster for the new cognitive science program at UC-Berkeley, (ii) a poster for the Shakespeare Festival at Stratford-on-Avon, and (iii) a poster promoting the city of San Francisco. The industrial design tasks required the design of (i) a desk time piece to commemorate Earth Day, and (ii) a toy to amuse and educate a 15-month-old toddler.

***Informal Overview of Data:*** Informally, the difference between the two cases seems to be the following:In freehand, when a new idea is generated, a number of variations quickly follow. The variations expand the problem space and are necessary for the reasons noted earlier. One actually gets the sense that the exploration and transformation of ideas is happening on the paper in front of one's eyes as the subject moves from sketch to sketch. Indeed, designers have very strong intuitions to this effect.

When a new idea is generated in MacDraw, its external representation (in MacDraw) serves to fixate and stifle further exploration. Most subsequent effort after the initial generation is devoted to detailing and refining the same idea. One gets the feeling that all the work is being done internally and recorded after the fact, presumably because the external symbol system cannot support such operations.

***Hypotheses:*** It is necessary to measure two things: (1) How are the two symbol systems being used with respect to the ill-structured/well-structured properties? (2) How does this impact the number of lateral transformations and reinterpretations? The hypotheses with respect to (1) are the following:

(H1) Free-hand sketching is syntactically more dense than MacDraw.

(H2) Free-hand sketching is semantically more dense than MacDraw.

(H3) Free-hand sketching is more ambiguous and/or nondisjoint than MacDraw.[8]

The specific hypothesis with respect to (2) is the following:

(H4) Well-structured representations will hamper the exploration and development of alternative solutions (i.e., lateral transformations) and force early crystallization of the design.

***Coding Scheme:*** A coding scheme was developed to measure syntactic and semantic density, ambiguity and/or nondisjointness, and lateral transformations. A few aspects of the scheme are presented here. A full discussion, complete with examples, appears in Goel (1992a).

The protocols were segmented into episodes along the lines of alternative solutions (which correlated with drawings on a one-to-one basis) and analyzed at this level. Syntactic and semantic density were measured in terms of a *variation* relationship between episodes (and the accompanying drawings).

A *variation* rating means that the current drawing is recognizably similar to earlier drawings. At the syntactic level this means that the equivalence-class of marks (i.e., syntactic types) constituting the drawing are closely related to, but distinct from, the equivalence-class of marks constituting one or more previous drawings. A *variation* rating at the semantic level means that the idea or content of the drawing is similar (but not identical) to the ideas or contents of one or more previous drawings.[9]

---

[6]'MacDraw' is a registered trademark of Apple Computer, Inc.

[7]'Mac II' is a registered trademark of Apple Computer, Inc.

[8]Notice here the collapse of the logically distinct notions of unambiguity and disjointness. It was not possible to distinguish between them with the given methodology.

[9] The connection between the variation rating and density can be seen with the aid of the following example. Consider two symbol systems, *SS1* and *SS2*. In *SS1* characters consist of equivalence-classes of line lengths which, when measured in feet, correspond to

Ambiguity and/or nondisjointness was measured in terms of *reinterpretations* of drawings. Reinterpretations occurred whenever subjects returned to earlier drawings and gave them a different interpretation.

In addition to the relationship between drawings/episodes and the interpretation of drawings, the types of operations which transformed one drawing into another were also coded for. A *lateral transformation* was one which modified a drawing into another related but distinctly different drawing (as opposed to a more detailed version of the same drawing, a totally unrelated drawing, or an identical drawing).

***Results:*** Sequences of episodes which received a *variation* rating were considered to be more densely ordered than those which received some other rating. Measured as such, the ordering of episodes (or alternative solutions) is significantly denser in freehand sketching than in MacDraw. The first row of Table 1 (Syntactic Density) shows the number of densely ordered drawings in freehand sketching versus MacDraw per session. The second row of Table 1 (Semantic Density) indicates that the number of densely ordered ideas per session is also much greater in freehand sketching than in MacDraw.

Table 1: Mean Numbers of Densely Ordered Episodes and Reinterpreted Episodes per Session.

| | Free-hand | MacDraw |
|---|---|---|
| Syntactic Density | 11.2 | 3.0** |
| Semantic Density | 10.4 | 4.1** |
| Reinterpretations | 2.4 | 0.67* |

**p<.005, one-tail; *p<.05, one-tail

There was also a significantly greater number of reinterpretations in freehand sketching than in MacDraw (see Table 1, row 3). Thus as predicted, the freehand sketches displayed greater ambiguity and/or lack of syntactic disjointness.

On the basis of these results, and converging verbal evidence, it is concluded that the two symbol systems are indeed being used in the way predicted. That is, the freehand sketches belong to a substantially ill-structured symbol system while MacDraw drawings belong to a substantially well-structured system. Finally, we want to know whether this has the predicted impact on the number and types of transformations.

Table 2: Mean Numbers of Lateral Transformations per Session

| | Free-hand | MacDraw |
|---|---|---|
| Syntactic Lateral Transformations | 8.9 | 3.2* |
| Semantic Lateral Transformations | 8.0 | 3.9* |

*p<.05, one-tail

It turns out that there is a statistically significant difference in the number and types of transformations (see Table 2). As predicted we get significantly more lateral transformations, at both the syntactic and semantic levels, with the ill-structured representations (freehand sketching) than with the well-structured representations (MacDraw).

## Discussion & Conclusion

Before rejecting the null hypothesis associated with H4, and concluding that well-structured representations hamper the exploration and development of alternative solutions (i.e., lateral transformations) and force early crystallization of the design, it is necessary to examine some alternative interpretations of the data. A rather obvious alternative interpretation is that the behavioral differences have nothing to with the theoretical differences underlying the manipulation, it is simply that freehand sketching is easier to use than MacDraw. One would get a similar hampering of exploration and development if, instead of MacDraw, subjects were forced to draw with a twelve-foot pencil.

Table 3: Mean Numbers of Sessions & Episodes in Minutes, & Mean Number of Episodes & New Solutions per Session.

| | Free-hand | MacDraw |
|---|---|---|
| Duration of sessions (min) | 57.7 | 53.2 |
| Duration of episodes (min) | 2.5 | 2.8 |
| Number of episodes | 16.4 | 14.4 |
| Number of new solutions (syntactic level) | 5.2 | 4.0 |
| Number of new solutions (semantic level) | 5.6 | 3.9 |

This interpretation is not, however, supported by the data. The effects of the manipulation are selective, and as predicted. There are no significant differences (F<1) along a number of other important dimensions, including the duration of the sessions, number of episodes per session, the duration of episodes, and the number of new solutions generated per session (Table 3). If the difference was just one of ease vs. difficulty

the integers. So we have lengths of 1', 2', 3', etc. In *SS2* characters consist of equivalence-classes of line lengths which, when measured in feet, correspond to the rational numbers. So we have lengths of 1', 2', 3'...; but also lengths of 1.5', 2.5', 3.5'... and 1.25', 2.25', 3.25'... and 1.125', 2.125', 3.125'... and so on. Lines of lengths 1.125' and 1.25' are no more identical than lines of length 1' and 2'; neither of these pairs belongs to the same equivalence-class. However, line lengths of 1.125' and 1.25' are much more "similar" or "closer to each other" -- with respect to length -- than lines of 1' and 2'. Thus the notions of "similarity" or "closeness" seem to be an integral (necessary?) part of density.

of use, then one would expect significant differences along each of these dimensions.

There is, however, a second alternative interpretation which needs to be taken more seriously. On this account, there are no behavioral differences across the two treatment conditions. What seems like a difference (the hampering of lateral transformations) is just an artifact of the methodology. It is a well-accepted assumption of protocol analysis that a more complete record of internal activity will occur when there is a good match between internal and external symbol systems (Ericsson & Simon, 1984). If, this is the case, and one also assumes that the system of internal representation is ill structured then it follows that the freehand sketching record is more complete than the MacDraw record. So the appearance of behavioral differences is caused by a different degree of completeness in the records.

However, this interpretation -- by assuming that the system of internal representation is ill structured – violates some very important metatheoretical constraints on the system of internal representation (Goel, 1991a; Goel, 1991b), and leads to a much stronger conclusion. I am postulating that different symbol systems are correlated with different cognitive functions. The alternative interpretation requires one to make an assumption about the structure of internal representations which is very strong, and contrary to much of the literature. It seems more prudent to accept the original interpretation.

On the basis of these results, the failure of alternative interpretations, and the assumption that lateral transformations are desirable,[10] it can be concluded that, at least some ill-structured problems -- like design -- require (or at least benefit from) ill-structured diagrammatic representations during the early, explorative and generative phases of problem solving.

On one reading, this is a rather unremarkable conclusion. Any designer can tell us that sketching is important for preliminary design. Why does this obvious fact need to be established by experiment? What makes the conclusion interesting is that the analysis of symbol systems employed, and the design of the study, allow us to tie the results to certain specific properties of symbol systems, namely density and ambiguity. So the study not only confirms the obvious, it offers an explanation of it in terms of ambiguity and density of symbol systems.

[10] This, I take it, is an unproblematic assumption. It amounts to little more than the claim that better solutions will result if one is allowed to explore the space of solutions and to customize any preexisting solutions to the present context.


## Acknowledgements

The author is indebted to Peter Pirolli, Mimi Recker, and Susan Monaghan for helpful discussions and comments. This work has been supported by a Gale Fellowship, a Canada Mortgage and Housing Corporation Fellowship, a research internship at System Sciences Lab at Xerox PARC and CSLI at Stanford University, and an Office of Naval Research Cognitive Science Program grant (# N00014-88-K-0233 to Peter Pirolli).


## References


Elgin, C. Z. (1983). *With Reference to Reference*. Indianapolis, IN: Hackett Publishing.

Ericsson, K. A., & Simon, H. A. (1984). *Protocol Analysis: Verbal Reports as Data*. Cambridge, Massachusetts: The MIT Press.

Goel, V. (1991a). Notationality and the Information Processing Mind. *Minds and Machines*, *1* (2), 129-165.

Goel, V. (1991b) *Sketches of Thought: A Study of the Role of Sketching in Design Problem Solving and its Implications for the Computational Theory of Mind*. Ph.D. Dissertation, University of California, Berkeley.

Goel, V. (1992a). The Cognitive Role of Sketching in Problem Solving. Manuscript submitted for publication.

Goel, V. (1992b). Specifying and Classifying Representational Systems: A Critique and Proposal for Cognitive Science. In *Proceedings of Conference on Cognition and Representation*. SUNY, Buffalo.

Goel, V., & Pirolli, P. (in press). The Structure of Design Problem Spaces. *Cognitive Science*.

Goodman, N. (1976). *Languages of Art: An Approach to a Theory of Symbols (second edition)*. Indianapolis, IN: Hackett Publishing.

Kosslyn, S. M. (1981). The Medium and the Message in Mental Imagery: A Theory. *Psychological Review*, *88* (1), 46-66.

Larkin, J. H., & Simon, H. A. (1987). Why a Diagram is (Sometimes) Worth Ten Thousand Words. *Cognitive Science*, *11*, 65-99.

Newell, A., & Simon, H. A. (1972). *Human Problem Solving*. Englewood Cliffs, N.J.: Prentice-Hall.

Pylyshyn, Z. W. (1981). The Imagery Debate: Analogue Media Versus Tacit Knowledge. *Psychological Review*, *88* (1), 16-45.

Rich, E. (1983). *Artificial Intelligence*. N.Y.: McGraw-Hill Book Company.

Simon, H. A. (1972). What is Visual Imagery?: An Information Processing Interpretation. In L. W. Gregg (Eds.), *Cognition in Learning and Memory*. NY: John Wiley & Sons, Inc.