ARMY RESEARCH LABORATORY

Situation, Domain, and Coherence: Toward a Pragmatic Psychology of Understanding

John Warner
Manhattanville College

David D. Burnstein
U.S. Army Research Laboratory

ARL-CR-306

JULY 1996

prepared by

Department of Psychology
Manhattanville College
125 Purchase St
Purchase, NY 10577

under contract

DAAL03-91-C-0034

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	July 1996	Final

4. TITLE AND SUBTITLE

Situation, Domain, and Coherence: Toward a Pragmatic Psychology of Understanding

5. FUNDING NUMBERS

PR: 62716AH70
PE: 6.2
DAAL03-91-C-0034

6. AUTHOR(S)

Warner, J.; Burnstein, D. D.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Department of Psychology
Manhattanville College
125 Purchase St
Purchase, NY 10577

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

U.S. Army Research Laboratory
Human Engineering & Research Directorate
Aberdeen Proving Ground, MD 21005-5425

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

ARL-CR-306

11. SUPPLEMENTARY NOTES

The point of contact is Beverly Knapp, U.S. Army Research Laboratory, ATTN: AMSRL-HR-MY, Ft Huachuca, AZ 85613-7000 (telephone 520-538-4704).

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

This report reviews theoretical perspectives about how people use and understand information. This is particularly relevant to an information production domain such as military intelligence. We focused on approaches in cognitive psychology and on problem solving in particular. These approaches proved limited in several important ways but specifically because the emphasis in cognitive psychology is information processing rather than information content and on "toy" problems rather than the pragmatic use of information to meet a domain-relative objective. Drawing upon social cognition and semiotics, we propose a conceptual framework that emphasizes coherence, or how people make sense of information concerning a domain objective, and pragmatics, or how people actually use information to meet that objective. Finally, we suggest that while understanding mental processing that underlies the use of information may be valuable, our pragmatic coherence framework is better suited to address questions about how people use information. It will provide a more direct link between theory and application.

14. SUBJECT TERMS

cognitive processes	information processing	military intelligence	semiotics
coherence	information production	narrative thought	social cognition
comprehension	information technologies	pragmatics	

15. NUMBER OF PAGES

45

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	

AMS Code 622716.H700011

Technical Note ARL-CR-306

SITUATION, DOMAIN, AND COHERENCE: TOWARD A PRAGMATIC PSYCHOLOGY OF UNDERSTANDING

John Warner
David D. Burnstein

July 1996

APPROVED: Robin L. Keesee
ROBIN L. KEESEE
Director, Human Research &
Engineering Directorate

Approved for public release; distribution is unlimited.

U.S. ARMY RESEARCH LABORATORY

Aberdeen Proving Ground, Maryland

ACKNOWLEDGMENTS

This work was supported by the Human Research and Engineering Directorate of the U.S. Army Research Laboratory (Dr. Julie Hopson). This support included a contract awarded to John Warner under the auspices of the U.S. Army Research Office Scientific Services Program administered by Battelle (Contract No. DAAL03-91-C-0034).

CONTENTS

TABLES

SITUATION, DOMAIN, AND COHERENCE: TOWARD A PRAGMATIC PSYCHOLOGY OF UNDERSTANDING

INTRODUCTION

This report reviews theoretical perspectives about how people use and understand information and how that understanding serves as the basis for decisions and actions. This is particularly relevant to an information production domain such as military intelligence (MI), where information goes into the system and information returns. Digitization across the battlefield is expected to increase the information explosion, making it imperative that analysts identify critical information and enhance their ability to assign meaning in a useful and timely way. How does the intelligence analyst assign meaning that serves some particular purpose or goal? What accounts for individual differences in how this meaning is assigned?

Although this report is motivated by the needs of the Army intelligence community, these questions are far reaching and very pertinent to the information age. More and more decision making is operating within information production domains where people have access to greater volumes and sources of information. To act efficiently and effectively, people must be able to determine what information is important and what it means. Designing technology to assist people will require understanding the way in which people need to use and interact with the content of information.

MAINSTREAM COGNITIVE PSYCHOLOGY

Within cognitive psychology, influenced by a computational information-processing metaphor, there are two approaches one can take in examining how people interact with information and respond. One approach is to look at the interaction from the perspective of thinking, that is, reasoning, problem solving, and decision making; the other is to see it from the perspective of language comprehension.

Work in <u>language comprehension</u> provides useful insight into how people interpret messages, particularly within the field of discourse comprehension. However, most research concerns verbal and textual materials and particularly within psycholinguistics, is often concerned with information processing. Also, it is rarely concerned with the objective people have in comprehending nor with the specific domains in which the comprehension is situated.

Work in thinking is more concerned with people's objectives since it is usually assumed that one has a purpose that is directing the thinking. It is also more likely to take some account of specific domains, particularly in studies of expertise. The lines between the sub-disciplines of reasoning, decision making, and problem solving are blurred. We will use "problem solving" as the broadest and most inclusive term. Information production can be viewed as a type of problem domain. The problem is to determine, given the information you have and your knowledge, what information needs to be produced to meet some goal.

We begin by looking at how cognitive psychology has defined and explored problem solving to see whether this is a useful way to account for information production. Although this review invokes historical roots, it is not exhaustive. The focus is on the current state of the field and the usefulness of current theoretical perspectives for the issues raised. We follow this by raising one of the major shortcomings of mainstream cognitive psychology, a lack of situated pragmatics. We are concerned with the impact of this shortcoming (and some others) on cognitive approaches in two areas: modeling how people understand information content and designing technological aids for understanding information. We address these shortcomings by reviewing related work within social cognition and some related applied fields. These are fields where cognitive psychology has been influential but where there is a greater concern with pragmatics. Finally, we will propose a new framework for thinking about the information production that we have developed for MI and show how it attempts to satisfy the concerns raised.

COGNITIVE PERSPECTIVES ON PROBLEM SOLVING

Complex problem solving and language are viewed as highly specialized human uses of general cognitive abilities and resources such as perception, attention, working memory, and knowledge retrieval. Cognitive psychology, like the older gestalt tradition, takes a top-down approach to problem solving. This is in contrast to the more bottom-up approach of the behaviorist-associationist traditions. The key assumption for cognitive psychology is that perception and memory involve the construction and retrieval of representations. By this view, thinking is then the active restructuring of these representations or knowledge to create new knowledge, usually to achieve a particular goal (Glass & Holyoak, 1986).

A problem is usually defined as a situation in which the current state of affairs fails to match one's current goal (Glass & Holyoak, 1986; Holyoak, 1990; Martindale, 1991; Matlin, 1994). This definition of problem can be interpreted quite broadly and sometimes has been. For example, Van Lehn (1989) suggested that "virtually any human activity can be viewed as the solving of a problem." Martindale (1991) commented, "Technically, even getting out of bed involves problem solving. Because you know perfectly well how to solve this problem, the difficulty has to do with motivation." However, most researchers limit the definition to situations that involve both focused mental activity (Lesgold, 1988) and a goal that is unattainable without changing the current state.

Heuristic Search

Most cognitive approaches to studying problem solving are based on the general framework developed by Newell and Simon (1972). This model characterizes problem solving as a search through a space of possible solutions (i.e., possible and legal transformations of the representation of the current situation into the desired goal situation) (Holyoak, 1990). Problem solving involves several steps. First, the problem representation must be adequately constructed, including representations of the goal, objects that can be transformed or manipulated, possible operations or actions, obstacles, and constraints on possible actions. Next, these represented elements must be combined and related so that the result produces a solution. This representation is a plan or simulation of action in the world. Finally, the plan must be executed. The plan is the procedural path through the solution space that results from the search. A solution is the sequence of operations that transforms the current state into the goal state.

Difficulties may occur at any of the three broadly defined steps above. First, if one misrepresents the elements of the problem, one may fail to solve the problem. This can be attributable to failure to use a needed object, failure to account for a serious obstacle, or even failure to represent an appropriate goal. Second, one can represent the problem well but then fail to identify the key relationships that allow you to know how to form the plan. Third, of course, the plan can simply fail when you try it.

One major difficulty is that many problems have a potential solution space that is much too large for any human to do an exhaustive search. In chess, for example, a problem for which all components can be clearly represented, you know what the goal is, what different objects are used, what operations you can use, and what your constraints are. The problem is that the solution space is very large with very many paths that may or may not lead to the goal. Clearly, skillful chess-playing humans do not play chess by a brute force search of this solution space (De Groot, 1965), though some computer programs do. What is it, then, that people do?

Within Newell and Simon's framework, it is assumed that what the human problem solver does is engage in heuristic search. That is, a rule is generated to guide a hopefully smart but not exhaustive search of the solution space. These are sometimes called "weak" methods because they do not guarantee that you will reach a solution. Much of the research in problem solving within cognitive psychology has focused on these weak methods that people use in attacking unfamiliar problems. Four such methods have been systematically identified from human problem solving protocols and have received the greatest attention (Newell & Simon, 1972): generate and test, working forward, working backward, and means-end analysis.

"Generate and test" is the least focused strategy. The person simply proceeds in a systematic fashion trying each possible action. Then they reevaluate the problem to see if they are closer to the goal. Such a crude strategy is little different from brute force search, except when there are few options and there is some systematic way to ensure that each option is tried only once. Thus, generate and test is not a useful strategy for writing a novel, but it can be adequate if you are trying to decide which desk drawer your paper clips are in. Lacking adequate knowledge, it may be the only strategy you can use.

Working forward and working backward are related, but different strategies. When working forward, one considers the actions possible, selects the best one, and observes what happens. The cycle is then repeated, accounting for any changes that have occurred. Working backward involves decomposing the problem from the goal state to determine what solution path leads to that goal. That is, you are tracing back from the goal to find the logical sequence that leads to your current state. Which strategy you use will be determined by whether the initial situation or the end state presents the fewest options. For example, consider a complex wooden puzzle. It is probably easier to determine how to put it together by starting with the completed puzzle and taking it apart, rather than trying to decide how the individual pieces should be put together.

One of the most important strategies is means-end analysis (Greeno & Simon, 1988; Holyoak, 1990; Newell & Simon, 1972). It is a mixture of working forward and backward, but the emphasis is different. In means-end analysis, you compare the current state to the goal state and identify differences between the two. You then identify an operator or action that will reduce one of those differences. If an operator cannot be applied, you create a sub-goal of changing the situation so the operator can be applied and apply means-end analysis to this new sub-goal. The key to means-end analysis is that it allows you to decompose any problem into smaller, hierarchically nested sub-problems (many of which you may know how to solve). You then reduce differences between the current state and each sub-goal. Because the sub-goals are nested, solving a sub-goal may make it much easier to solve the more complex higher level goal. Also, by breaking the problem into linked sub-goals, less demand is made upon the human's working memory limitations.

Another area of study concerns blocks to problem solving. A "block" may occur from having too large a search space. More often, they result from the problem solver being too fixed or inflexible in his or her representation of the problem. The two most famous examples are <u>mental set</u> and <u>functional fixedness</u>. Mental set involves trying to apply a strategy from a previous problem to new problems when other solutions might be more effective or when the strategy is no longer appropriate. Functional fixedness involves failing to consider a potential solution because of a failure to realize that an object can be used in a non-conventional way. This would be like someone who intends to bend a nail, spending a long time looking for a hammer when he or she has a lug wrench in hand. In the problem representation terms we used before, the person has failed to represent the wrench as "heavy tool that can bend nails."

Attempts to develop computer programs to solve problems have led to a conscious, serial, symbolic rule-based architecture, called a production system (Anderson, 1983). Production systems fit the search framework of Newell and Simon (1972) and have been used to model such heuristic problem solving (Greeno & Simon, 1988; Anderson, 1983). A production system involves knowledge being encoded as a series of pattern-action rules called productions. Whenever the appropriate pattern or condition or antecedent is satisfied, the rule applies a particular procedure (either a mental or a physical operation). Thus, the knowledge pertaining to stopping at a red light could be represented by the simple production:

IF light is red, THEN press foot on brake.

More complicated tasks might be represented by a series of productions as

IF goal is to eat AND you are hungry,
THEN search kitchen for food.
IF no food in kitchen,
THEN make "get food" current goal.
IF current goal is "get food,"
THEN drive to store.

Notice that even in this small set, a nesting is already apparent. You can easily represent how a problem is broken into sub-goals. Presumably, when you begin a task, all rules that partially match a pattern or condition fire but are strictly sequenced based on a hierarchical scheme in which the goal takes precedence. The orders of precedence will dictate that, if the top-level goal cannot be met, a lower precedence rule can operate to achieve some sub-goal that will further satisfy the goal condition(s). If you need to drive to the store, but you do not have a car, you would fire a rule that can accomplish the sub-goal of obtaining a car.

Production systems as knowledge representation are attractive for several reasons. First, they resemble the stimulus-response framework that psychologists are already familiar with. It is a small leap from physical state-physical operation contingencies to perceptual state-mental operation contingencies. We also know that contingencies can be learned, and we understand how such learning looks. Such learning functions can further lead to adaptiveness in novel situations. Reorganization, nesting, and reprioritization can be used to model the development of expertise.

Production systems are useful for programming a procedural search model of thinking such as Newell and Simon's General Problem Solver (GPS). However, it suffers from the limitations of the problem solving approach that it mimics. It is easy to identify productions for a clearly defined problem and show how that knowledge could be acquired, but what about problems with less clear representations. Can a production system negotiate for peace?

Domain-Situated Problems and Expertise

Most of the work just cited was generated around knowledge-lean tasks (Van Lehn, 1989) such as difficult word problems and puzzles. These problems require very little knowledge. Usually, the knowledge presented in the instructions is sufficient to reach (or stumble upon) a satisfactory solution. Another way to characterize these types of problems is to call them "domain free"; they are not situated in a specific domain where domain knowledge might come into play. Simon and others were well aware of the over-simplification involved here and later rejected a model based on this approach as an adequately generalizable model of human problem solving (Matlin, 1994). Nonetheless, useful information has been gleaned from this work, some of it relevant to more realistic, ecological problems. The main focus of research into knowledge-rich or domain-situated problems has been the study of expertise. What are the factors that distinguish novices from experts? The amount of experience they have with the domain is different, but is that it? The old joke is that the way you get to Carnegie Hall is "practice, practice, practice." What is it that the expert acquires from all this practice or experience? In fact, a large variety of important differences have been identified. Three areas are most important: cognition, the interaction of perception and knowledge, and metacognition.

The cognition of experts is more effective than that of novices, but they are not more intelligent in any generic sense (no strong positive correlation between expertise in a domain and IQ). If you compare performance in domain-situated problems between experts in that domain and experts from another domain, the other domain experts behave just like novices (Voss & Post, 1988). General cognitive abilities in experts are not better developed than in novices, other factors being equal. Rather, it is the case that their cognitive skills within their own domain are more efficient. The classic example of this comes from De Groot's (1965) study of expert versus novice chess players. De Groot showed that chess experts were superior at retrieving actual middle game chess configurations but were no better than novices at retrieving random (i.e., non-meaningful) chess configurations.

Experts also show differences in the way that they perceive and represent problems and the way that their knowledge interacts with their representations (Matlin, 1994). One obvious aspect of this is that experts have more knowledge about the problem domain. This allows them to draw upon richer relations between concepts in the domain. In addition, experts have much better developed domain-knowledge schemas to structure the problem in a more useful way (Holyoak, 1990). These schemas allow them to reduce memory requirements by allowing information to be meaningfully chunked (as in the De Groot chess studies). Since these schemas are based on abstracting deep structure characteristics, experts can more quickly and efficiently recognize what is important about a problem or find a meaningful analog. However, on the flip side, experts are more flexible in their use of schemas (Lesgold, 1988). They are much better at tolerating and accommodating exceptions and deviations from their schemas. In addition, experts simply spend much more time elaborating the initial problem representation.

Finally, experts are much better at monitoring the progress of their problem solving, of accurately estimating the difficulty of the problem, estimating how long the problem may take, and at knowing when they have made an error.

Representation and Restructuring

Much of the discussion so far assumes that the problem (and its solution) can, in fact, be well represented. However, the kinds of problems humans encounter in real life and thus the most interesting, are those that are ill defined. These are problems in which one or more of the elements needed for a solution are not clear and cannot be represented. Consider writing a novel or determining enemy course of action (ENCOA). Although you know the goal in some abstract sense, it is not clear what that end state will be like. In this case, your solution space is being constructed as you solve the problem.

Many times, the difficulty with ill-defined problems is not that the solution space is too big, although in the abstract it may be. Rather, it is that we fail to find any solution at all until we look at the problem in a whole new way. The notion of insight problems was developed by the early Gestalt psychologists. These are problems in which the parts of the problem seem unrelated. While people may attempt to search for solutions, these searches are rarely even close to successful. What seems to happen when people solve insight problems is that the solution just suddenly appears via a sudden shift in perception from the unrelated parts to a whole where it is clear how the parts are related (for some empirical evidence for this intuition, see Metcalfe & Wiebe, 1987). This can be considered as a type of problem solving block where you are viewing the problem in a fixed and non-useful way. It is usually a shift in your perception of the problem that frees you from the block. What seems to be clear about restructuring and insight is that a problem that requires restructuring is one for which the solver cannot predict the solution, even while solving it, and the solver is usually surprised by the solution (Metcalfe, 1986). Further, trying to verbalize or think out loud about the problem seems to actually interfere with reaching a solution (Schooler, Ohlsson, & Brooks, 1993).

One area where restructuring has been extensively studied is in the use of analogies in solving problems. If you can identify similarities between a problem you are trying to solve and a previous problem, you can restructure the new problem to highlight these similarities. You can then map the previous solution to the new problem. Two important findings have emerged from examining the use of analogies in problem solving. First, analogical transfer is difficult to reliably demonstrate, with subjects sometimes failing to notice the usefulness of an analogue even if the two similar problems are presented in quick succession (Cummins, 1992; Holyoak, 1990). Second, when it occurs, what has transferred is an abstraction of deep structural similarities in the problems, and the new problem must provide explicit cues for that abstract structure.

What is most important in insight problems and other ill-defined problems is that they are solved by restructuring the problem representation, not by brute force search and stumbling upon an answer. This restructuring involves implicit, possibly parallel processing below the level of awareness. In addition, people (e.g., experts) who spend much on elaborating their problem representations and are flexible in their use of those representations seem less susceptible to problem blocks (Holyoak, 1990).

Attempts to model ill-defined problem solving in computer programs has led to implicit, parallel, sub-symbolic activation architectures sometimes called connectionist networks (Holyoak, 1990; Martindale, 1991; Rumelhart, 1989; Smolensky, 1988). Connectionist approaches model thinking as changing weights in an activation network that converges, in parallel, on a particular representation. The analogy is to parallel activation in perception (Holyoak, 1990). As you view an apple, for instance, parallel analyzers in the brain are responding to basic features such as roundness, redness, hardness, size, and so forth. Certain other concepts may become activated by the features they share with apple (e.g., cherries). When activation settles, that combination of features should converge on a representation of apple. While the leap from perception (in which we have a richer grasp of brain mechanisms) to thinking is a large one, it is reasonable. However, you have to assume constraints on the network since the number and degree of overlap on possible interpretations is so much greater (Holyoak & Thagard, 1989). This architecture reflects problem solving, not as serial search in which sub-goals are solved sequentially but as parallel cycles of activation and inhibition. Complex interactions between sub-goals are weighed to find the best fit solution, given conflicting constraints, represented as a stable global pattern of activation over a network. This is sometimes called "soft constraint satisfaction" since each constraint influences the behavior of the network, but they are not inviolate as a rule would be (Holyoak & Spellman, 1993).

There are two main limitations to representing human thinking with connectionist architectures. First, purely bottom-up approaches to higher cognition create a combinatorial explosion that defies human cognitive capacity limitations. Second, connectionist architectures have difficulty representing any relation richer than single place predicates and Boolean logic (Holyoak, 1990). Nonetheless, there are distinct advantages to soft constraint satisfaction for ill-defined, perceptually noisy, or ambiguous problems. In addition, because both representations and control are distributed over the simple units (like abstract neurons) (Rumelhart, 1989), there is no need for a central executive function for decision and control (e.g., deciding which production rule has priority). Such functions seem uncomfortably close to postulations of a homunculus in the head whose cognition also needs to be explained.

More recent approaches have used architectures combining both symbolic and connectionist assumptions. Soft constraint satisfaction is performed in a network but over units that are more local and symbolic, thus having relational structure. One example is the Analogical Constraint Mapping Engine (ACME) (Holyoak & Thagard, 1989; Holyoak & Spellman, 1993) used to perform analogical mapping (i.e., to map a novel situation to features of a previous situation in memory in order to create a similar structure for understanding the present situation).

ACME applies a set of abstract constraints on a network of symbolic, predicate-calculus-style representations. It then performs parallel constraint satisfaction to settle the network to reflect confidence in possible mappings. Another example is Kintsch's (1988) construction-integration model of text comprehension. The model was motivated in part by difficulties presented to modeling text comprehension by ambiguities. Top-down symbolic approaches were too inflexible and given to errors that were not representative of human performance, but bottom-up approaches created the above-mentioned combinatorial explosion. Kintsch's approach was to model the process as parallel, bottom-up activation of memory concepts, and propositions by words and phrasal structure with limited spreading activation to close associates. Additional propositions are then activated via inference rules. Finally, the model applies soft constraint satisfaction to the network with an interpretation of some small portion of text represented in the stable and settled network. The process then is reiterated for the next portion of text. Certain units that are still active after the network has settled are used in the next portion (for a conceptual theory of comprehension that is quite compatible with this architecture, see Gernsbacher, 1990).

Whereas production systems correspond to symbolic search models, connectionist models correspond to the associationist learning theory that the cognitive revolution sought to supplant. These behaviorist frameworks view problem solving as learning--the acquisition, through experiential contingencies with the environment, of a set of behaviors that produce an adaptive response that changes the environment. Certainly, an important step in solving a problem is to "learn" how to solve that problem. This learning may also generalize to other sufficiently similar problems. In addition, different problems that share common sub-goals (the behaviorist would say sub-components) can take advantage of behavioral chaining to benefit from previous learning (the common behavioral explanation for purported "insight learning"). Even symbolic approaches talk about processes such as "proceduralization," which can be paraphrased as "practice increasing the speed and reliability with which a particular stimulus patten elicits a particular set of responses," although for the cognitivists, these "responses" can be mental operations. Clearly, one of the key differences between experts and novices is that experts have learned more about solving the problems in their domain.

However, this approach fails to handle some important aspects of skilled problem solving well, in particular, the use of knowledge and the structural aspects inherent in the organization of that knowledge. Because the behaviorist tradition views learning as the acquisition of S-R contingencies, knowledge becomes trivialized as simply the storage and triggering of the associations. Just as true connectionist architectures have difficulty representing relations, this framework has trouble accounting for the expert's flexibility in his or her use of knowledge to restructure a problem as presented.

Another, older framework that has an affinity with connectionist approaches is the Gestalt approach. This approach views problem solving, such as perception, to involve parallel, unconscious restructuring of knowledge. This can result in a particular stimulus configuration appearing one way and then very suddenly, without awareness of the steps that caused the change (hence, "insight"), appearing in a very different configuration. As cited, there is empirical support for such unconscious operations in problem solving, and they are clearly important. The framework has generally been considered too vague to give an account of the conscious aspects of planning and problem solving or the strategic use of knowledge. On the other hand, symbolic approaches have been used to accommodate many of the findings associated with the Gestalt notion of "insight" (Kaplan & Simon, 1990).

There has been a movement toward hybrid architectures (combining symbolic and connectionist approaches) as well as movement toward explaining cognition as involving both symbolic and connectionist computations. These approaches combine both conscious, serial, symbolic mental operations with more Gestalt-like parallel, unconscious, activation processes that can influence ongoing conscious processing. One example is Kintsch's (1988) comprehension model, already mentioned. While still in their early stages, these hybrid theoretical frameworks offer a greater possibility of capturing the richness and flexibility of human cognition abilities such as problem solving.

MILITARY INTELLIGENCE AS PROBLEM SOLVING

Military intelligence can be considered an information production domain in the sense that information (signals, reports, commander's intent) serves as input, and information (estimates, recommendations, requests) are output. That is, information content is both the raw material and the product. Further, since military operations are strongly goal driven, intelligence production in an operational context is a good example of problem solving (or reasoning, or decision making) in a domain. Is it useful to think of MI this way? Are the current theoretical frameworks previously outlined, adequate for describing what is important about what the MI officer or MI staff needs to do? If not, what is missing?

What is the Problem in MI?

In MI operational contexts, the overall operational goal and even tactical goal may be clear, but important sub-goals for the MI staff may not be. The MI problem is to receive incoming information, sort and interpret it in light of operational goals and commander's intent. This could be represented, for instance, as a complex means-end analysis (this is only one possibility). That is, suppose the goal is to produce conclusions concerning enemy activity. The operational goal is the top-level goal. The enemy activity must be represented in order to understand its impact on reaching the top-level goal. Now the intelligence staff makes some conclusions to move the present situation closer to the top-level goal, which may involve a conclusion about the enemy activity. The problem may be easy to solve. For example, the expert recognizes the problem that the enemy activity presents and may retrieve an appropriate schema to guide their response. The problem may also be difficult and the MI staff may have to restructure the information before they can resolve what the enemy objective is. The problem-solving frameworks we have discussed are general enough to fit into any domain, this one included.

While the MI staff has a set of goals leading to satisfying requests for information, would it be appropriate to think of them as searching a space of possible plans to satisfy that request? How would the elements of this solution space look? It is more likely they are "searching" for a consistent representation or map of the various sources of information upon which their satisfaction of this request depends. This can be viewed as a type of problem solving, but it is a sub-type that may be unique to information production domains. This type of "problem" is called comprehension.

A comprehension task, such as reading, involves encoding information for more than ensuring retrieval or understanding the meaning of the symbols (words). Rather, the purpose is to embellish and enrich the text by constructing relations to what you already know. This results in understanding not only what the text states but what it implies about the world that text represents. This could be described as a "search" for understanding, but it would be a very different sense of search. Someone does not search a space of alternate interpretations, but rather, information and knowledge mutually constrain interpretation until a particular interpretation is reached, similar to Kintsch's (1988) comprehension model.

Thus, in many respects, what MI does can be better termed "situation understanding" rather than problem solving, an area that has not been given as much attention within cognitive psychology. According to Greeno (1977), situation understanding (he used the term problem understanding) requires three conditions to achieve a useful internal representation, which is the goal of understanding. First, the representation should be coherent, in which coherence is achieved when the pattern is connected so that all the parts make sense. Second, there should be a close correspondence between the constructed representation and what is being understood. That is, although this is a constructionist framework, there needs to be a correspondence between the representation and the world. Finally, the representation must be well connected or related to background knowledge. These three requirements are achieved by skillful representation via the fitting of new information with other information and what you already know.

Information has meaning that is ascribed by a comprehender on the basis of symbolic or interpretive conventions or the idiosyncratic experience of the individual. Much of the meaning in the MI domain is domain defined and context imposed. However, when the information is understood, it has become coherent. This colors the meaning, based on relationships and interactions between that information, other pieces of information, and one's knowledge. Coherence relative to an objective within a domain reduces the degrees of freedom in interpretation. This is analogous to the way soft constraint satisfaction produces a globally stable network of relations.

Focusing on "situation understanding" rather than "problem solving" moves the emphasis to a coherent representation of the situation rather than on searching for solutions. It is in representing the situation that the largest differences are observed between novices and experts. In fact, in a domain, there is probably little need to search for a solution. Once you understand a situation, there is usually a well-established course of action. That is not to say that the course of action will not be idiosyncratic, since it may be well established by experience rather than doctrine. If the understanding fails to meet one or more of Greeno's requirements, the course of action will be less clear or may be inappropriate.

PRAGMATICS AND SOCIAL COGNITION

A Pragmatic Psychology of Understanding

Note from the previous discussion that understanding involves two kinds of "connections" beyond simple retrieval of meaning. Coherence has to do with connecting new information with what you know so that it will make sense. This implies a strong influence of experience in the coherence-making process. Since some of a person's past perceptions and experiences are idiosyncratic, some aspects of their understanding will be idiosyncratic. The second connection is to the real world as it exists now, the situated context along with the person's operational goal(s). Because the person is interacting with the world and because the goal of understanding is to enable him or her to act within and upon the situated context (even if the action is to produce information), this aspect of understanding is very pragmatic.

The term "pragmatics" is used in linguistics to refer to the study of meaning in context. Although this is appropriate, we are reaching farther back to two sources for our use of the term. The first is William James' (1907) pragmatic theory of meaning that stipulates that concepts are defined in terms of their use in experience. The second source is semiotic theory (Morris, 1955). Here, pragmatic relationships represent the relationships between signs and their users. Both ideas indicate that meaning is based on what your experience has been and what you are trying to do in the world. Paradoxically, this makes meaning both well regulated within a domain and operational context and idiosyncratic between individuals.

Cognitive psychology has often had a tendency to study cognition in the absence of actual domains and contexts, an influence of the cognitive revolution's computational roots. Because of this, its methods and theoretical frameworks have missed both ends of the paradox. However, other sub-fields within psychology, particularly those with an inherently applied agenda, have more actively pursued the pragmatic aspects of perception, cognition, and thinking.

Probably no area in psychology has developed the pragmatic tradition so richly as social cognition (Bruner, 1990; Fiske, 1993). One reason is that they view cognition in terms of the social group and interactions within the group rather than viewing cognition within an individual, outside a pragmatic context. More recently, some of these ideas have begun to flow back from social psychology to cognitive psychology as interest in domain-specific knowledge and real-world contexts has increased (Bruner, 1990; Harre & Gillett, 1994; Levine, Resnick, & Higgins, 1993).

The pragmatic perspective within social cognition (Fiske, 1992; 1993) is a constructionist perspective, as is the mainstream of cognitive psychology. Much of constructionist cognitive psychology, operating outside realistic domains, concerns itself with explaining human "errors" in reasoning (e.g., one misperceives the problem, one's heuristics bias one's judgments, etc.). In contrast, social cognition has viewed interpretation and meaning as being attuned to the pragmatic social context. Thus, accuracy of the response is not determined relative to an ideal or a statistical or probabilistic norm but rather relative to a "good enough" standard, given a goal.

According to Fiske, people are good enough perceivers. They will believe what satisfactorily fits the data and is goal relevant. Performance is accurate if it is useful and adaptive. Notice that Fiske's notion of "good enough" perception has affinity with Greeno's ideas of requirements for problem understanding. Perception can be "good enough" probably only if Greeno's requirements are met.

According to social cognition theory, people make meaning by abstracting essential structures of the situation or experience and then substituting these as "good enough" versions of the original. This idea is compatible with our earlier discussion of expertise and the use of schemas. Like an MI staff assessing a situation via incoming information, a person in a new social context relies on schemas, traits, stereotypes, and other abstracted structures to quickly assess the situation in order to know how to respond. Fiske (1993) talks about one structure in particular that is important for understanding the situation--narrative. Bruner (1990) suggests that narratives are important for linking the unexpected to the known by using human intent as a context for explanation. Thus, narrative cognitive processes establish coherence by embellishing relations between novel information and the perceiver's knowledge.

Influences of Pragmatics and Social Cognition

Finally, the pragmatic tradition within social cognition stipulates that thinking is for doing. Therefore, in studying perception and cognition, one must consider the perceiver's goals and motivations.

These pragmatic views of human cognition are emerging wherever questions need to be addressed directly within a domain, in applied settings, or via naturalistic observation. In cultural psychology and anthropological psychology, recent views suggest that to understand people, one needs to understand what motivates them. This means one needs to understand their goals and the overall interpretive system that constructs and interrelates those goals (D'Andrade, 1992). One change that seems to emerge from this view is that "culture" is treated much more as a domain or a context in which action is situated. That is, culture has psychological relevance for imparting motivational force through shared cultural schemas.

Recent themes in behavioral decision research (a field greatly influenced by Herbert Simon) emphasize that preferences for and beliefs about objects or events of any complexity are often constructed in the generation of a response to a judgment or choice task (Payne, Bettman, & Johnson, 1992). According to the research cited by Payne et al., preferences are not necessarily generated by consistent and invariant algorithms. Decision makers have a variety of methods for identifying preferences. These methods are derived from both experience and training. The information and strategies used to construct preferences highly depend on task, context, prior knowledge, and individual difference factors.

Influences of social cognition and anthropology in organizational behavior have led to pragmatic concerns with viewing the organization as a context mutually influenced by the interactions of groups and individuals. Mowday and Sutton (1993) point out that by not considering the organizational context, including individual and group constructions of that context, very little progress can be made in accounting for the actual organizational phenomenon that organizational behavior purports to study. Instead, studies refine some number of micro-theories of limited application in real-world settings (a complaint that too often applies to cognitive psychology in general).

Another applied discipline currently influenced by social cognition is consumer psychology, which attempts to assess the impact of message variables on consumer decision behavior. Tybout and Artz (1994) identify three important new trends in consumer psychology. First, there is a greater emphasis on the consumer's embellishments in his or her understanding of information. Second, more attention is being paid to characterizing the rich context in which consumer perception and decision making occurs, including the adoption of more naturalistic research methods that preserve that rich context. Finally, the view that consumers are analytical and rational decision makers has been gradually supplanted by the view that decisions are based on heuristics, good enough perceptions, usefulness relative to a goal, and even emotional responses.

Thus, there is a growing zeitgeist around taking a pragmatic perspective, not only in social cognition but in other disciplines as well. To a small extent, this perspective has begun to influence thinking in cognitive psychology. The pragmatic perspective is fueled by the need to ask and answer questions about real-world situations and contexts and naturalistic settings in which interactions are complex and normative assumptions clearly break down. This constitutes not just a theoretical shift but a methodological one as well. To the degree that pragmatics are related to the domain context in which a person is situated, studying ill-situated word meaning and structure will not answer questions about human behavior in the real world. To the degree that pragmatic relationships are idiosyncratic, as we previously indicated, comparing group means for large numbers of subjects will fail to reveal important aspects of the very phenomena we wish to explain and predict.

Since our interest in problem solving is to describe and predict how MI understands information in a real operational situation, as well as how that understanding influences decisions and judgments, such a pragmatic perspective is important for any useful model of problem solving. This shifts the emphasis to "good enough" perception, elaboration of the situation representation, and what is useful to the problem solver. By thinking of problem solving in an information production domain as understanding, coherence moves to the foreground. This changes the question from how someone navigates the problem space and what heuristics they used (a notoriously difficult question to answer for ill-defined problems) to what their response tells us about how they understood the problem.

A PRAGMATIC COHERENCE FRAMEWORK

In the following section, we present a framework and approach to address the issue of making information coherent. First, because our concern was information production and how people make information coherent to accomplish some objective, we needed a means of addressing information content more than information processing. Second, because we were concerned with pragmatic use of information content, we needed to develop the framework around real tasks, situations, objectives, and domains. Third, the framework needed to reflect the role of individual idiosyncrasies in situation understanding.

The framework attempts to characterize how people comprehend and give coherence to information in order to accomplish an objective. Because we have developed this framework within military intelligence, an information production domain, we will be concentrating on tasks in which the response is to produce new information content. We assume that this framework can be applied to other domains but must be situated in each domain separately.

The framework includes the following components: a description of the semiotic content of information, a description of the possible content transformations that information can undergo as well as the various purposes of those transformations, a description of four types of objectives that occur in any domain situation, and finally, a set of four architectures, called "paragons," that represent templates that can be overlaid on situations in a domain to make information coherent, given an objective. In what follows, we describe each of these components in detail and show how the pieces fit together in the overall framework.

Semiotic Content

This framework describes the content of information using a semiotic perspective (Morris, 1955). The content of information can be divided into three types of content relations (see Table 1). The syntactic are sign-to-sign relations. These are hierarchical or configurational relations between objects (primitives or classes), structures, events, states, or measures (e.g., "tank" might be an object; three tanks, an object quantified; or part of an independent tank regiment). The semantic concerns relations between signs and their meanings. These can be conceptual and/or causal dependency relationships between signs. For example, objects and certain events (movement of petroleum trucks) may lead to a conclusion that a particular site is a possible fuel depot. The pragmatic is concerned with the relationship between signs and the user of the signs.

Table 1

A Semiotic Description of Information Content

Type	Description	Sub-types
Syntactic	Hierarchical or configuration relationship	Objects Events Structures States Measurements
Semantic	Dependency relationships	Conceptual Causal
Pragmatic	Goal-driven relationships	Conditional (when, where) Instrumental (how) Consequential (why)

An example might be a conclusion based on states, event structures, or semantic relations about what an enemy is capable of doing. These are relationships that can be described as conditional (when and where), instrumental (how), and consequential (why). In operational environments, most syntactics and semantics are domain defined. That is, in the example, the domain definitions made the conclusion of "fuel depot" the most reasonable interpretation, given those objects and events. Pragmatics, on the other hand, being related to the user, are idiosyncratic and therefore a source of individual differences.

Content Transformations

Information content can be changed in two ways, which we call semiotic transformations: coalescence (changing the content via the fusion or forming of a whole) and analysis (changing the content via separating or decomposing information). As shown in Table 2, there are six types of each of these two transformations. For each of the six transformation types, we have associated two verbs, one for coalescence and one for analysis. These verbs characterize the purpose of the transformation.

Table 2

Types of Semiotic Transformations

Purpose of analysis	Transformation*	Purpose of coalescence
Parse	syntactic→syntactic	
	syntactic←syntactic	Aggregate
Decompose	syntactic→semantic	
	syntactic←semantic	Interpret
Interpolate	syntactic→pragmatic	
	syntactic←pragmatic	Project
Elaborate	semantic→semantic	
	semantic←semantic	Infer
Justify	semantic→pragmatic	
	semantic←pragmatic	Synthesize
Rationalize	pragmatic→pragmatic	
	pragmatic←pragmatic	Envisage

*The direction of the arrow indicates whether content is being put together (→) or taken apart (←). Coalescence refers to transforming content by putting together; analysis refers to transforming content by taking apart.

The six types of semiotic transformations are

• The transformation of syntactic information to syntactic information. **Aggregate** refers to coalescing lower level syntactic elements into higher level elements such as a structure. **Parse** refers to analyzing higher level structures into their component objects.

• The transformation of syntactic information to semantic information. **Interpret** refers to coalescing syntactic elements into a meaning or possible meaning. Meaning is generally defined by the domain. **Decompose** refers to decomposing the elements underlying a domain-defined meaning.

• The transformation of syntactic information to pragmatic information. **Projecting** refers to coalescing syntactic elements to compute estimations relevant to a specific objective. **Interpolation** refers to analyzing an estimation to determine the elements that lead to the estimation.

• The transformation of semantic information to semantic information. **Infer** refers to coalescing semantic information into information having different or richer meaning, possibly moving from one conceptualization to another. **Elaboration** refers to analyzing the relationships that may have led to a particular conclusion or conceptualization.

• The transformation of semantic information to pragmatic information. **Synthesize** refers to coalescing semantic information in order to impose the context of the operational environment. That is, one is forming a story that can be used to generate hypotheses about the current operational situation. **Justify** refers to analyzing contextually situated information or hypotheses to provide empirical support.

• The transformation of pragmatic information to pragmatic information. **Envisage** refers to coalescing contextually situated information and hypotheses in order to develop predictive scenarios regarding the operational situation. **Rationalize** refers to analyzing predicted scenarios for potential consequences.

Situational Objectives and Coherence

This framework also describes four objectives that any situation can present. That is, in any situation within any domain, there are one of four objectives a person will be trying to achieve, based on his or her perception of the characteristics of the situation. The person's objective may be appropriate or inappropriate to a particular situation, but that objective will always be one of the following four: constructive, diagnostic, reactive, or explanatory.

In addition, the framework describes four paragons (see Figures 1 through 4). These represent templates that can be imposed on a particular domain situation to provide the most efficient means of transforming information to establish coherence relative to a particular objective. The figures only capture the general use of information content transformations (coalescence and analysis). They must be situated in a domain before one can know the exact nature of those transformations (i.e., the purpose of the coalescence or analysis, as outlined previously). It is for this reason that we refer to the figures as "templates" that must be imposed on a domain situation.

The term "paragon" is used to refer to the four figures in the sense that these architectures are each an idealization. Each represents the most efficient way to disambiguate (screen information in the context) and transform information to make information coherent within an operational environment. This is not to say that a person faced with a particular situation will always impose the ideal template, nor does it mean that coherence cannot be established using a different template. However, imposing an inappropriate template will decrease efficiency, modify task difficulty, and lead to errors.

There is the temptation to view the figures as models of human information processing. These architectures are more accurately viewed as design tools you would use to help you design a machine or control a human to use information for a specific goal. Although, we refer to "cognitive structures" and "cognitive processes," these only serve to indicate how information needs to be organized and what level of processing is required. For instance, when we show comparison or integration, we identify what needs to be done with the information. We are not at all concerned with how it is done or what the actual processing might look like.

Table 3 lists the criterion for selecting the appropriate template for making information coherent to meet a particular objective. This table describes the outcome or "consequence" of using the template, the source of the "evaluation criteria" for how good an outcome is, the "purpose of disambiguation" of the situation, and the range of possible outcomes for each type of objective. Thus, the title of each paragon (constructive, diagnostic, reactive, and explanatory) refers to a type of objective that is most appropriate when the criteria in Table 3 are met. Once the criteria are identified, the most appropriate template can be selected. The template can then be used to design a program for directly controlling how a human transforms information most efficiently.

Constructive

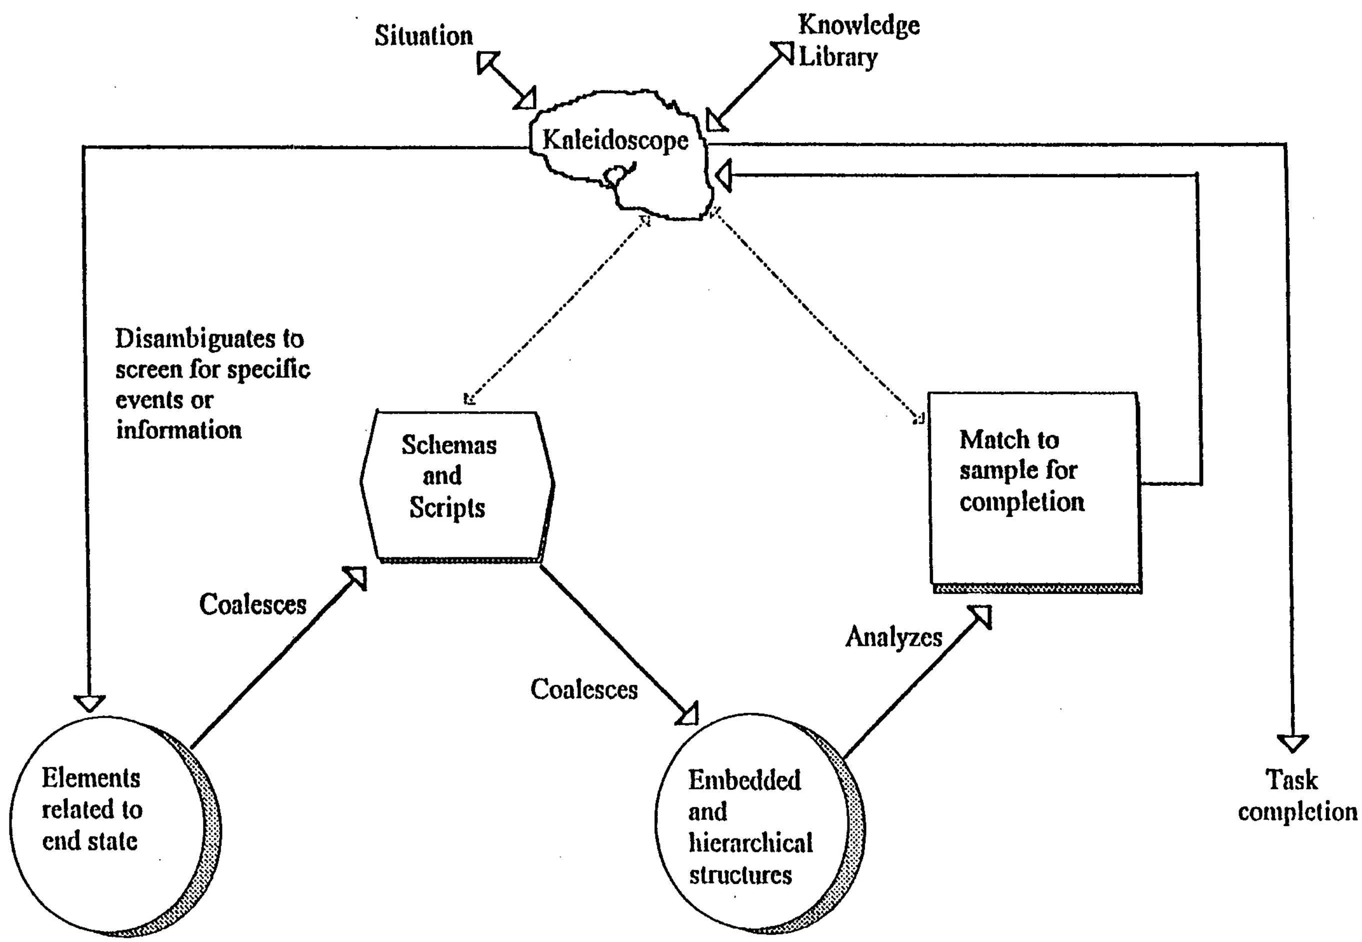

Figure 1. Template for efficiently making information coherent to meet a constructive objective.

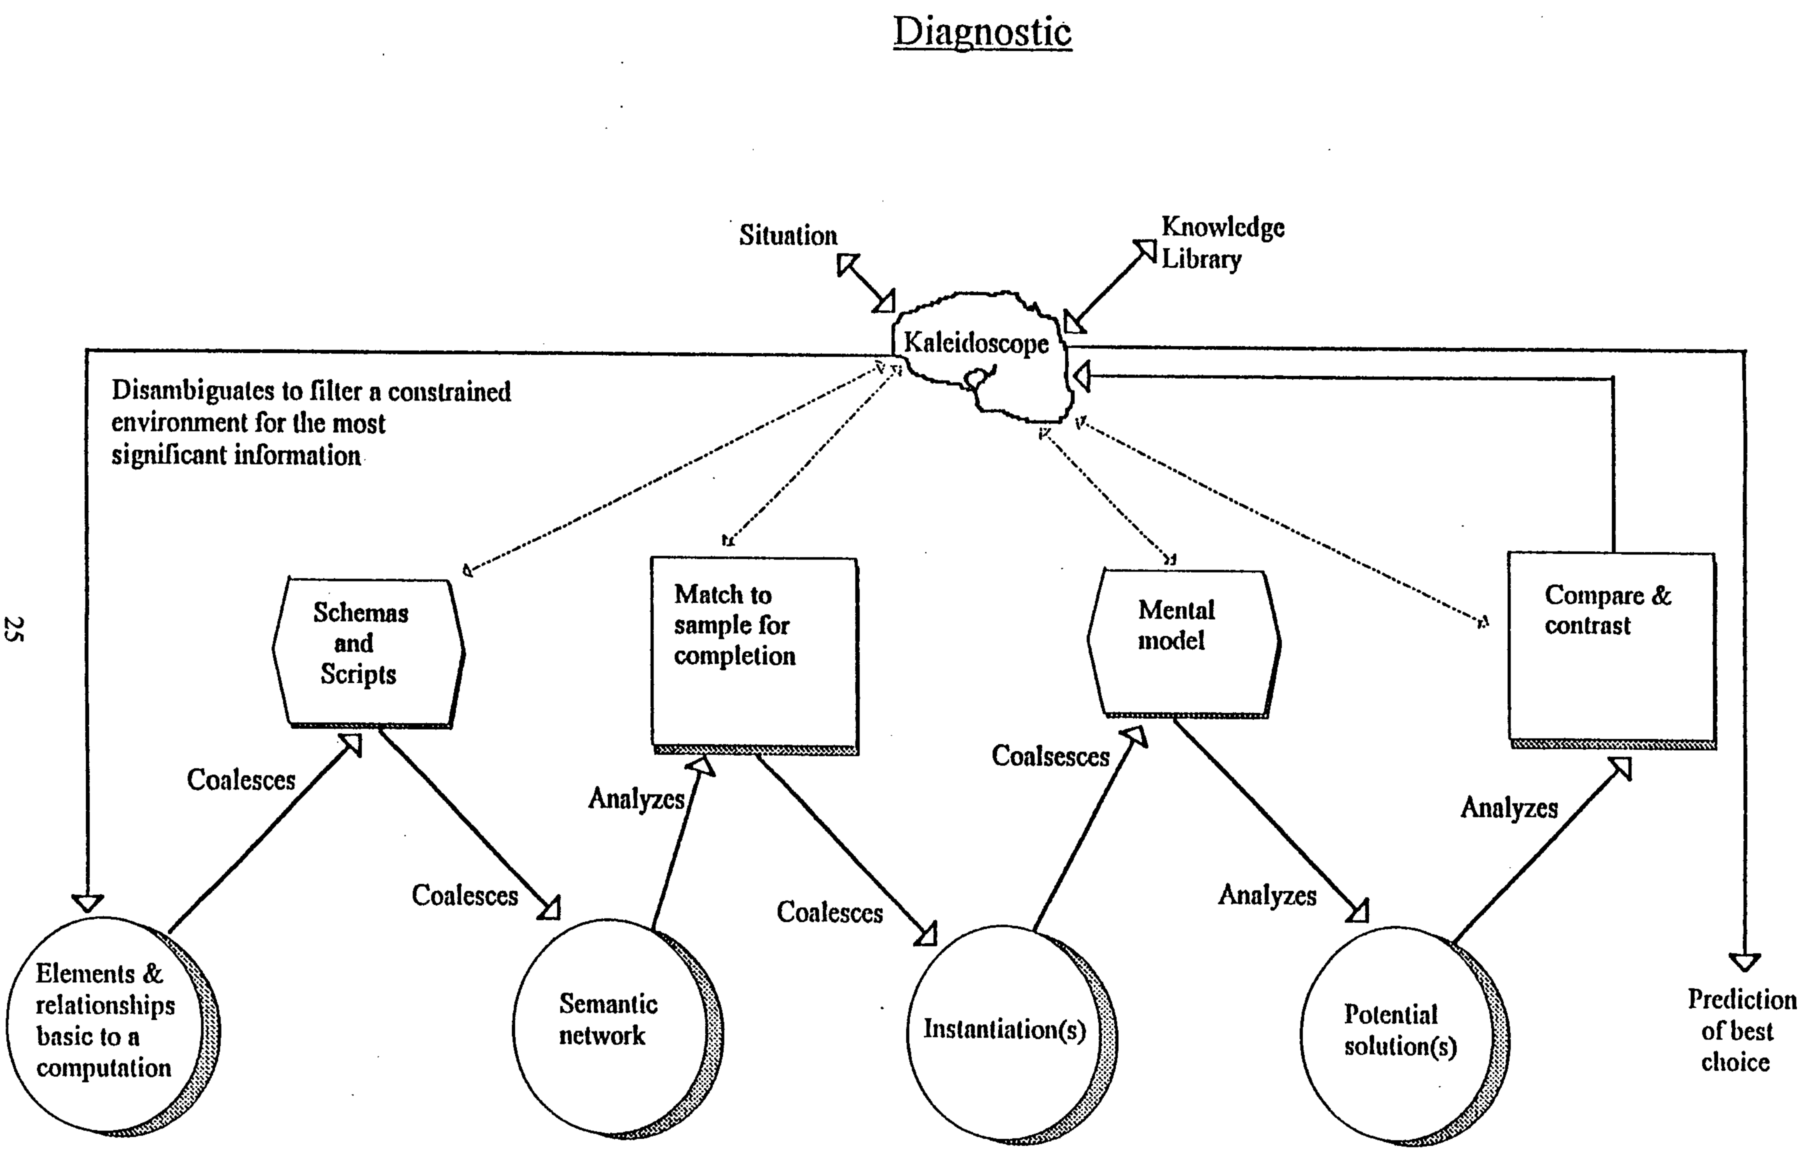

Figure 2. Template for efficiently making information coherent to meet a diagnostic objective.

Reactive

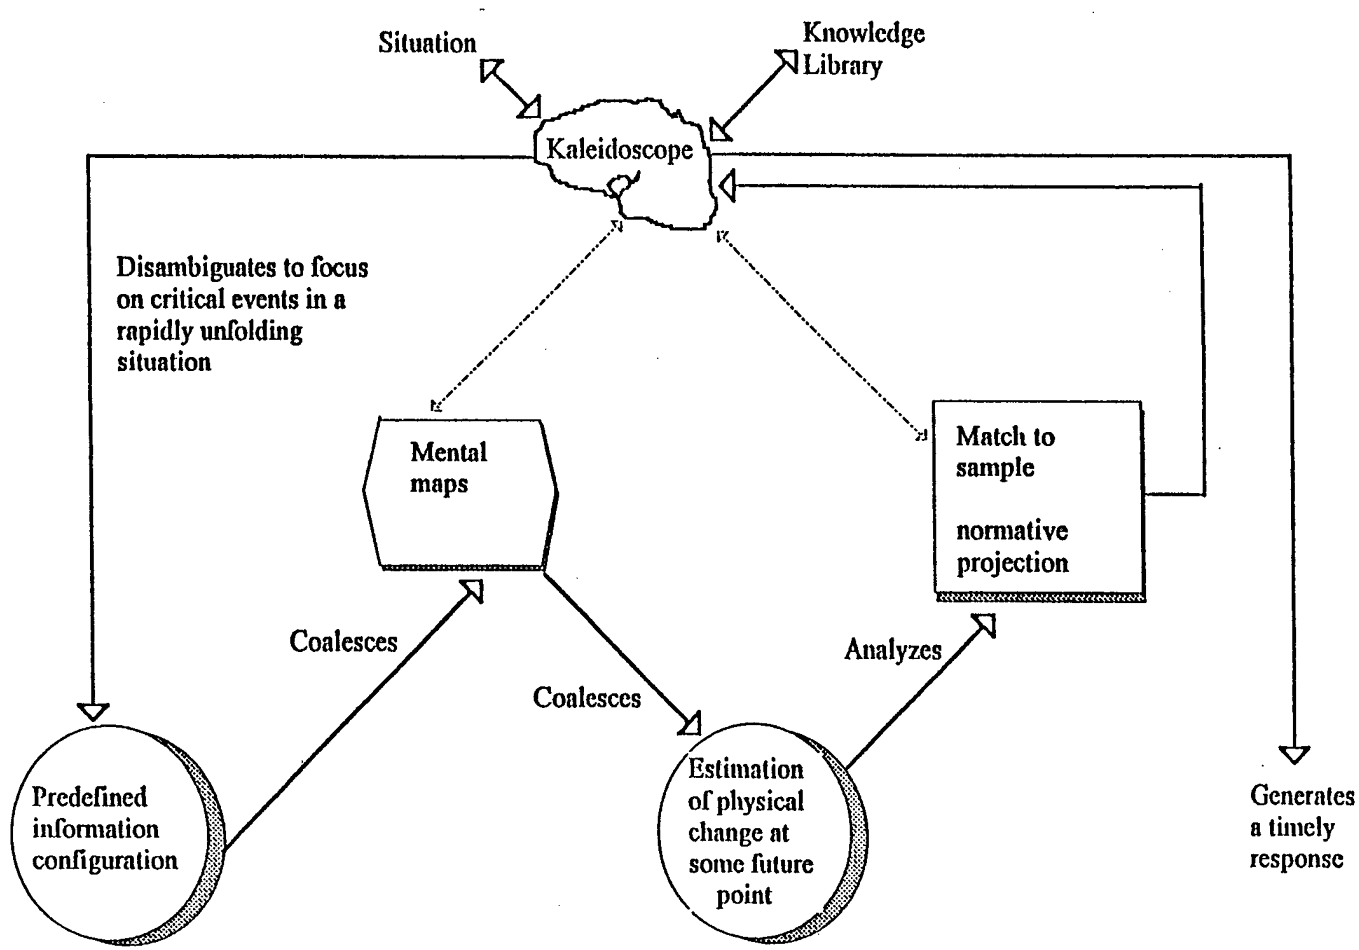

Figure 3. Template for efficiently making information coherent to meet a reactive objective.

Explanatory

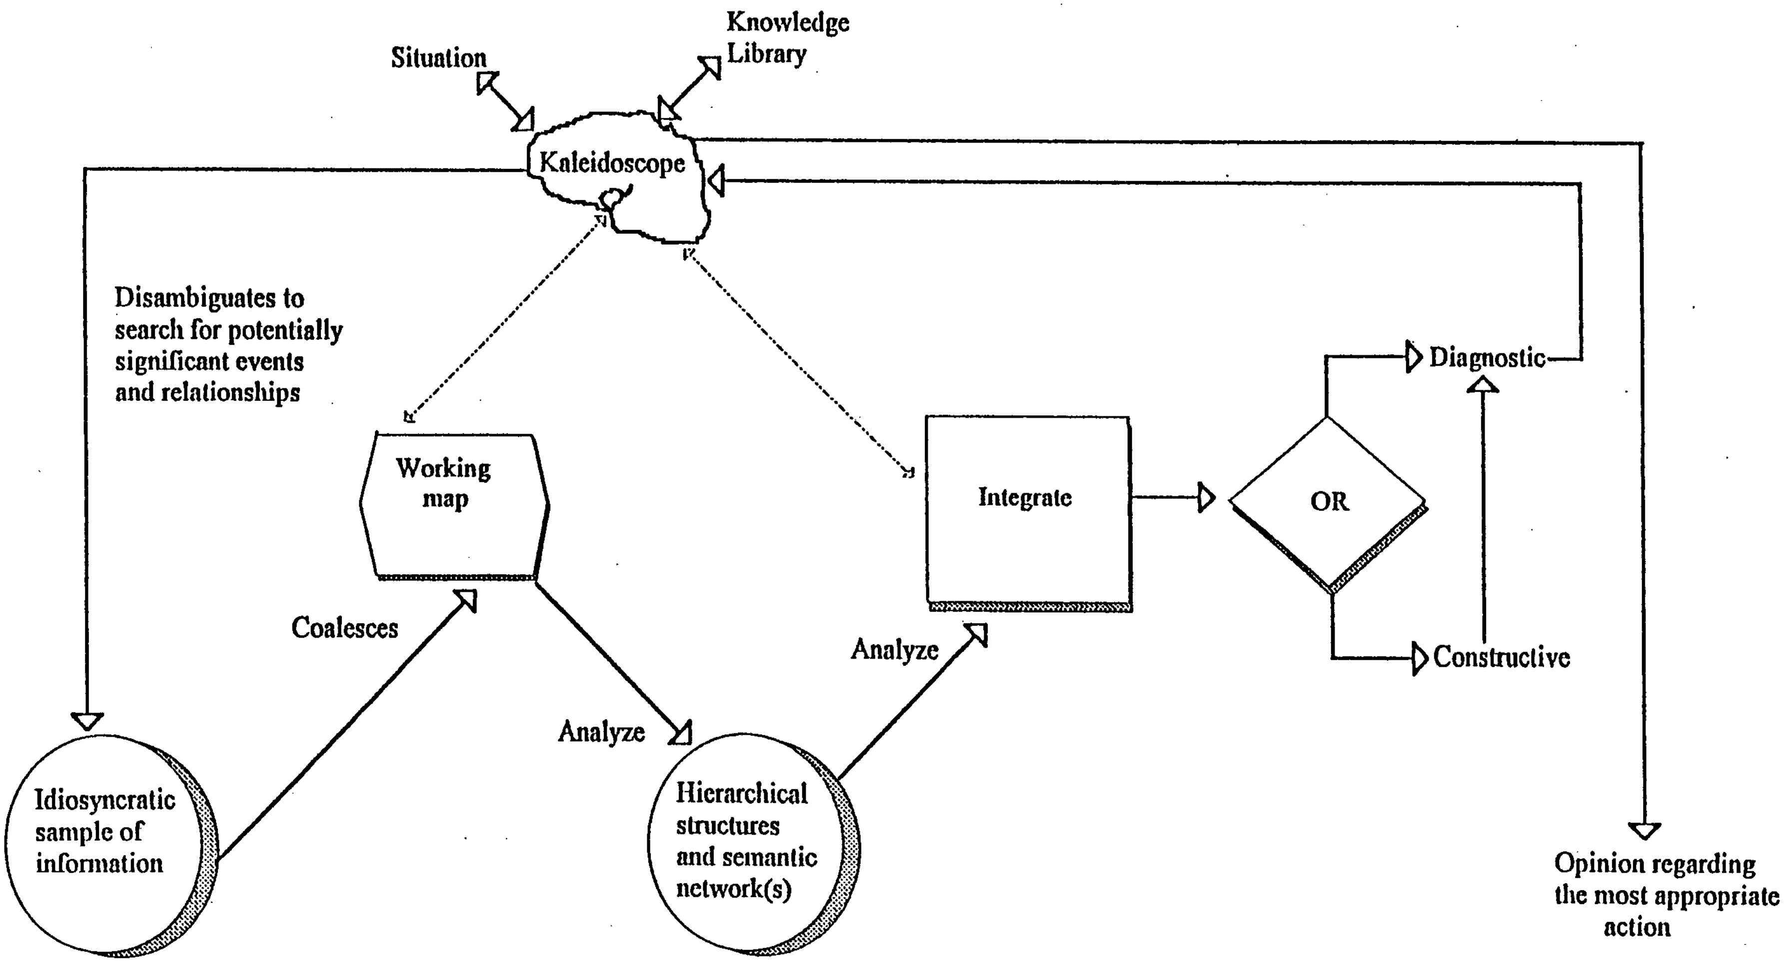

Figure 4. Template for efficiently making information coherent to meet an explanatory objective.

Table 3

Criteria for Selecting a Paragon as a Template for Making Information Coherent to Meet a Particular Objective

	Constructive	Diagnostic	Reactive	Explanation
Consequence	Task completion	Prediction of best choice	Timely response	Opinion regarding the most appropriate action
Evaluation criterion	Established by the task requirements	Established by history of consequences	Established by physical requirements	Must be derived
Purpose of disambiguation	To screen for specific events or information	To filter a constrained environment for the most significant information	To focus on critical events in a rapidly unfolding situation	To search for potentially significant events and relationships
Range of solutions	Pre-determined	Limited by the problem	Limited by the response	Unbound

The General Architecture

The figures are "read" from left to right but with the understanding that everything begins at and results emerge from the "kaleidoscope," a massively parallel, unconscious "black box" that serves as the interface between the situation, prior knowledge (the knowledge library), and ongoing conscious processing. Each figure can be viewed as having three horizontal "layers."

The bottom layer shows how the information is changing in content and form. Thus, all circles in the figures represent the current information content, given the situation, knowledge, and the results of ongoing transformations. One always begins with the left-most circle immediately after disambiguation of a situation.

The top layer represents the kaleidoscope. It represents a very dynamic, ever-changing synthesis of knowledge and new information at different content resolutions. We use the term "kaleidoscope" to capture the idea that, within that dynamism, certain types of information content and knowledge will result in particular patterns.

The middle layer represents the cognitive structures and processes that are acting upon information. These are conscious modes of operating on information, but as indicated by the dashed double arrows, they emerge from the kaleidoscope. Within the kaleidoscope are the products of all ongoing processing, including the current information content. Solid arrows show how information content is changing as well as how the person uses cognitive structures and operations in establishing coherence. However, the figures do not indicate the fact that, to get a usable result, the person or the program may have to perform several cycles of disambiguation and transformation.

Constructive Objectives

This is probably the most straightforward and common type of situation, although it is not to say that constructive objectives (see Figure 1) are always simple. These situations have the completion of some physical and/or mental task as their objective, and what constitutes completion is predetermined by the task itself. Examples of situations in which constructive objectives are appropriate are jigsaw puzzles, template-filling tasks, well-defined procedural tasks, and so forth.

The most efficient way to transform information content to address a constructive objective is to start by using the pattern in the kaleidoscope (given the current situation and prior knowledge) to disambiguate the situation to screen specific events or information. This results in information that is highly syntactic in content, with objects and events related to the end state or goal. In the jigsaw puzzle case, this is very clear. The person takes the puzzle out of the box and lays out the pieces. He or she begins to look at how those pieces can be used to complete the task (the representation of your end state being the picture on the box). In MI, for example, a person may be looking at many different messages about enemy activity. The objective is to develop a functional analysis of that activity by completing a matrix or creating a map overlay. Once the important syntactic elements have been disambiguated, he or she begins to coalesce those elements to fit them into schemas (expectations based on knowledge organization). These schemas (or scripts, in a more procedural task) result in coalescing elements into embedded and hierarchical structures. The new aggregations are compared to the end state to determine if the task is completed. If it is not complete, he or she cycles back through until it is.

Diagnostic Objectives

Diagnostic situations (see Figure 2) are also quite common but more complex. This is a situation in which the person needs to discriminate among several possibilities or outcomes. Then, based on one's experiential history, predict which will be best for minimizing serious consequences. One example is medical diagnosis; another is selecting a horse on which to bet. In each case, the person is working with a constrained situation with a few possibilities. He or she first looks for specific and important relationships and "simulates" the consequences of those relationships to evaluate possible solutions. Finally, the person predicts the best choice, given the possible consequences.

The most efficient way to establish coherence to address a diagnostic objective is to use the pattern in the kaleidoscope to disambiguate the situation. The constrained environment is filtered for the most significant information. This will result in the current information content having both syntactic and semantic relationships. These elements and relationships are then coalesced in order to use schemas. These schemas guide and further coalesce the information into a semantic network. The semantic network provides a richer understanding of the important dependencies in the situation. They are analyzed and matched to both current knowledge and the situation to be sure that the important dependencies have been considered. Keep in mind that a person may need to repeat this cycle several times before the information captures the important dependencies sufficiently to continue. The next step is to coalesce these relationships to build actual instances that can be coalesced and fed into a stored mental model or simulation of the situation. The results of this simulation are analyzed to determine the solution that the current simulation suggests. This cycle will be repeated at least once more. These solutions are analyzed, based on the consequences of each in the mental model. The eventual result is a prediction of a best choice, given an experiential history of consequences.

This is a situation whose objective is highly pragmatic, since it is concerned with consequences. For that reason, any situation that requires a choice may need to use the diagnostic template.

Reactive Objectives

Superficially, the paragon (see Figure 3) appears very similar to the paragon for constructive situations, but reactive situations are very different from constructive ones. In fact, reactive situations are unique. A person must produce a timely response to a rapidly unfolding situation in which the consequences of a response are immediate. Reactive situations can be thought of as time-driven "emergencies" within any operational environment. Examples of situations when reactive objectives are appropriate include ballistic missile defense threat tracking, using a ground-based interception system, targeting in a fighter attack aircraft, and having an accident happen ahead of you on the highway.

Because of the nature of reactive situations, the paragon in Figure 3 represents not only the efficient way to transform information, but possibly, the only way. The pattern in the kaleidoscope is used to disambiguate the rapidly unfolding situation to focus on critical events. These events are embedded within a predefined information configuration. These events are then coalesced into a mental map or "snapshot" of the kaleidoscope representing the relationships (particularly physical or configurational relations) of those events within the unfolding situation. Because of the real-time aspects of a reactive objective, a person may repeat this cycle a few times, coalescing different mental maps at different points in time. The physical relations in the mental map(s) is(are) then coalesced to produce estimations of physical change at some future point in time. These estimations are analyzed to match them to a normative projection of events. This can produce a response, or the person can repeat the cycle (this is probably a very rapid cycle) and then make a response, or he or she can keep repeating the cycle and make a series of reactive responses in time.

Explanatory Objectives

An explanatory situation (see Figure 4) is one in which a person needs to comprehend the situation and define the objective. That is, he or she has a situation or environment not yet understood. Hence, he or she goes fishing for possible important and relevant information. This permits a person to generate a story that best covers the events and relations that are found. The story allows the person to understand, assuming the explanation is correct, what objective needs to be met. Notice that the "correctness" of the story is an assumption. The explanation is based on the person's experience and on his or her perceptions of how that information will be used. Because of this, there will be the largest individual differences with explanatory objectives. Examples of situations when explanatory objectives are appropriate include predicting ENCOA, anti-submarine warfare (ASW), and observing someone's very bizarre behavior and deciding what you need to do about it.

The most efficient way to transform information for an explanatory situation is to use the pattern in the kaleidoscope to disambiguate the situation by searching for potentially significant events and relationships. The result is a very idiosyncratic sample of information. This is because "potentially significant" means significant to the user, based on his or her knowledge, experience, and perceptions. This idiosyncratic sample of information (the content of which will tend to be more syntactic and semantic) is then coalesced into what we are calling a "working map." Whereas a mental map is thought of as a snapshot in time of the patterns in the kaleidoscope, a working map is a snapshot of the initial links formed between the events and relationships in this idiosyncratic sample. The map is then analyzed to extract potentially important hierarchical and semantic relations. Based on these structures and semantic net(s), the information is analyzed and integrated. Inappropriate relations are pruned and conclusions made to build relations and modify the strength or richness of existing relations. Again, this cycle may be repeated with new information and successive working maps feeding into an ongoing integration process. You can think of the integration process as constructing a narrative to explain the situation. You also can view the working map and integration together as a "working model" in the sense that the narrative is like a rough draft of a potential mental model. In fact, this is how we create our stored mental models.

The result of this part of the explanatory paragon is comprehension or understanding of the situation. You still need to give information coherence relative to an objective and in light of your explanation. To do this most efficiently, the person needs to impose the diagnostic paragon as a template for transforming information. However, he or she also may need to impose the constructive paragon as a template first. The purpose of this would be to gather more basic syntactics from the situation. The imposition of these paragons is considered to be part of the explanatory paragon because the person is working from and attempting to verify an idiosyncratic explanation of events. The result is not a prediction of the best choice (as in the diagnostic paragon) but an opinion (again, idiosyncratic) regarding the most appropriate action, given the situation as the person understands it.

PRAGMATICS AND COHERENCE: USE VERSUS PROCESSING

In the previous section, we outlined a framework for characterizing the ways in which information content is transformed to make it more coherent. The objective and the operational situation determines what makes sense. Looking at meaning without situating it with a domain purpose may tell you something about "memory processes," but it is unlikely that it will tell you much about how people make and use meaning. Because the framework concerns "use" in real-world situations, we call it "pragmatic."

In addition, since "use" involves a conditional, instrumental, or consequential relationship between the user and the real world, based on experience, this framework is forced to take seriously the role of individual differences. For instance, the explanatory paragon tells you that a person needs to start with an idiosyncratic sample of information, and based on personal (and therefore idiosyncratic) experience, he or she will connect the information to knowledge. If you do not want that kind of flexibility and creativity, then you might either have to change the objective or provide a decision aid that can help constrain the search and/or possible narratives. On the other hand, you may need idiosyncrasies (for example, a low intensity conflict such as guerrilla warfare where you do not know what to look for). In that case, you want to ensure (via training, decision aid design, or how information is presented) that you facilitate or take advantage of those idiosyncrasies.

Because the focus is on how people use information rather than how they process information, there is a much more direct link between theory and application. The theory tells you what kinds of information content are possible and the ways in which that content can change. The theory also tells you in what situation and with what objectives you will see these changes.

The framework, as specified in this report, is not developed enough to make this leap. The last main piece will be to turn descriptions of content and transformations into reliable measures of information content and transformations. However, the direct link is that the measures used to test the theory are the same measures used for application. If you can determine "how" people use "what" kind of information in a particular situation, you can begin to make predictions and design interventions. Controlling the type of information or the experience of the user can be accomplished without developing a detailed model of the cognitive processing involved. In fact, it is possible to fit current cognitive models to the framework without disrupting the momentum of application development within the pragmatic paradigm. However, there is a much greater burden of testing the theory against the real world in a domain for a pragmatic framework than for a processing framework. This approach cannot be addressed by domain-free laboratory tests of college undergraduates. One must look at the behavior of domain experts in situated, ecological, and naturalistic contexts.

There are few frameworks within cognitive psychology for how people understand information, which take seriously the pragmatic view and address how information is used rather than how it is processed. Some discourse models, because of their concern with understanding and coherence, have some flavor of this perspective. Nonetheless, they rarely situate their models in real world contexts. Both Bruner (1990) and Harre and Gillett (1994) have attempted to develop more pragmatic frameworks for the study of cognition and have even pointed to broadly sketched possible methodologies. Still, neither source has developed a concrete framework for predictions.

One interesting model, borrowing from the comprehension literature, is Pennington and Hastie's (1993) explanation-based decision-making theory. They have applied the theory to how jurors sort evidence to reach a verdict. Their general theory takes seriously the particular domain in which the decision is being made. The reasoner's "model" of that domain influences the reasoning process. A juror (and presumably anyone making decisions outside his or her area of expertise) uses narrative story structures to give the evidence coherence and select the most plausible explanation or causal path. In fact, the framework corresponds to explanatory and diagnostic objectives as outlined here.

Probably, the most extensively developed framework in this arena is the recognition-primed decision (RPD) model (Klein, 1993). The model has two components or decision-making modes. In the first mode, decision makers recognize important and salient features in the situation. This allows them to quickly identify expectations, a reasonable course of action, and what goals make sense. In a more complex, uncertain or dynamic situation, a second mode allows for dynamically seeking more information, rechecking expectations, and modeling (internally simulating) the courses of action that are selected.

The RPD model is pragmatic because it is a framework for understanding naturalistic decision making. In fact, Klein's general framework for studying naturalistic decision making is compatible with our approach, including his questioning of the reliance on closed system type problems in psychology. Further, the RPD model can easily be linked to our paragon templates. However, the RPD model is merely a set of templates, without situated content. The major difference between the RPD model and our framework is that the RPD model focuses entirely on the information processing of experts in a domain rather than the information content and how people use that information content to meet an objective. As such, it will remain inadequate for decision aid design applications, because the focus will be on supporting purported thinking processes. It does not look at how the user needs to use information and important differences between users.

CONCLUSIONS

This report reviewed approaches to problem solving (broadly defined) within cognitive psychology as a possible framework for looking at information production and how people give meaning to and understand information. This was a first step in a much longer term goal of being able to develop a framework that can improve training and guide the design of decision support technology for intelligence analysis or any other information production domain.

Because our ultimate interests are applied, we found approaches in mainstream cognitive psychology limited in several important ways. First, the only mature and well-developed theory of problem solving applies primarily to well-defined and closed systems. Second, most cognitive theories concern themselves with the generic mechanisms of mental processing rather than with the content of what (information) is being processed. Third, because cognitive psychology relies so much on ill-situated, domain-free closed systems, they lack ecologically valid contexts that reveal how people actually use information. That is, there are too few pragmatics and too little concern with the objective of comprehension.

Drawing upon social cognition and semiotics, we have proposed an alternate framework for understanding how people use information for information production. This framework emphasizes the content of information and the transformation of that content relative to the objectives of that domain. The framework highlights coherence in understanding, i.e., how the person makes sense of information relative to the domain and the domain objectives. The framework is thus also pragmatic and focusses on how people use information in particular situations to accomplish certain objectives. This framework also considers the importance of idiosyncrasies in understanding and giving coherence to information. Finally, because the framework emphasizes use of information relative to operational situations and objectives, it suggests that research needs to be situated in naturalistic or ecologically valid settings. Furthermore, research should focus on individuals in a domain rather than statistically removing idiosyncrasy in large group comparisons.

The framework as presented here is a conceptual proposal. Future work will include evolving the framework by testing key concepts and hypotheses, detailing, and employing candidate measures of information content, and by "operationalizing" the framework. Nonetheless, assuming measurement and operationalization are achieved, we believe that the framework will lead to a better understanding of information use and close the gap between theory and application. Though this framework has been developed for and situated in the military intelligence domain, we believe that it will be equally valuable for other domains. However, an important implication of the framework is that it would have to be operationalized separately for each domain.

REFERENCES

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

Bruner, J. (1990). Acts of meaning. Cambridge, MA: Harvard University Press.

Cummins, D.D. (1992). Role of analogical reasoning in the induction of problem categories. Journal of Experimental Psychology: Learning, Memory and Cognitive References, 18, 1103-1124.

D'Andrade, R.G. (1992). Schemas and motivation. In R. D'Andrade & C. Strauss (Eds.), Human motivation and cultural models. New York, NY: Cambridge University Press.

De Groot, A.D. (1965). Thought and choice in chess. The Hague: Mouton.

Fiske, S.T. (1992). Thinking is for doing: Portraits of social cognition from daguerreotype to laserphoto. Journal of Personality and Social Cognition, 63, 877-889.

Fiske, S.T. (1993). Social cognition and social perception. Annual Review of Psychology, 44, 155-194.

Gernsbacher, M.A. (1990). Comprehension as structure building. Hillsdale, NJ: Erlbaum.

Glass, A.L., & Holyoak, K.J. (1986). Cognition. New York: Random House.

Greeno, J.G. (1977). Process of understanding in problem solving. In N. J. Castellan, Jr., D. B. Psioni & G. R. Potts (Eds.), Cognitive theory (Vol. 2). Hillsdale, NJ: Erlbaum.

Greeno, J.G., & Simon, H.A. (1988). Problem solving and reasoning. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey & R. D. Luce (Eds.), Stevens' handbook of experimental psychology. New York, NY: Wiley.

Harre, R., & Gillett, G. (1994). The discursive mind. Thousand Oaks, CA: Sage.

Holyoak, K.J. (1990). Problem Solving. In D. N. Osherson & E. E. Smith (Eds.), Thinking: An invitation to cognitive science, Vol. 3. Cambridge, MA: MIT.

Holyoak, K.J., & Spellman, B.A. (1993). Thinking. Annual Review of Psychology, 44, 265-315.

Holyoak, K.J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. Cognitive Science, 13, 295-355.

James, W. (1907). Pragmatism. Indianapolis, IN: Hackett.

Kaplan, C.A., & Simon, H.A. (1990). In search of insight. Cognitive Psychology, 22, 374-419.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. Psychological Review, 95, 163-182.

Klein, G. (1993). Naturalistic decision making: Implications for design. Wright-Patterson AFB, OH: CSERIAC.

Lesgold, A. (1988). Problem solving. In R. J. Sternberg & E. E. Smith (Eds.), The psychology of human thought. New York, NY: Cambridge University Press.

Levine, J.M., Resnick, L.B., & Higgins, E.T. (1993). Social foundations of cognition. Annual Review of Psychology, 44, 585-612.

Martindale, C. (1991). Cognitive psychology: A neural-network approach. Pacific Grove, CA: Brooks-Cole.

Matlin, M.W. (1994). Cognition. New York, NY: Harcourt Brace.

Metcalfe, J. (1986). Premonitions of insight predict impending error. Journal of Experimental Psychology: Learning, Memory and Cognition, 12, 623-634.

Metcalfe, J., & Wiebe, D. (1987). Intuition in insight and noninsight problem solving. Memory and Cognition, 15, 238-246.

Morris, C. (1955). Signs, language and behavior. New York, NY: Braziller.

Mowday, R.T., & Sutton, R.I. (1993). Organizational behavior: Linking individuals and groups to organizational contexts. Annual Review of Psychology, 44, 195-229.

Newell, A., & Simon, H.A. (1972). Human problem solving. Englewood Cliffs, NJ: Prentic Hall.

Payne, J.W., Bettman, J.R., & Johnson, E.J. (1992). Behavioral decision research: A constructive processing perspective. Annual Review of Psychology, 43, 87-131.

Pennington, N., & Hastie, R. (1993). Reasoning in explanation-based decision making. Cognition, 49, 123-163.

Rumelhart, D.E. (1989). The architecture of the mind: A connectionist approach. In M. I. Posner (Ed.), Foundations of cognitive science. Cambridge, MA: MIT Press.

Schooler, J.W., Ohlsson, S., & Brooks, K. (1993). Thought beyond words: When language overshadows insight. Journal of Experimental Psychology: General, 122, 166-183.

Smolensky, P. (1988). On the proper treatment of connectionism. Behavioral and Brain Science, 11, 1-74.

Tybout, A.M., & Artz, N. (1994). Consumer psychology. Annual Review of Psychology, 45, 131-169.

Van Lehn, K. (1989). Problem solving and cognitive skill acquisition. In M. I. Posner (Ed.), Foundations of cognitive science. Cambridge, MA: MIT Press.

Voss, J.F., & Post, T.A. (1988). On the solving of ill-structured problems. In M. T. H. Chi, R. Glaser & M. J. Farr (Eds.), The nature of expertise. Hillsdale, NJ: Erlbaum.

NO. OF COPIES	ORGANIZATION
2	ADMINISTRATOR DEFENSE TECHNICAL INFO CENTER ATTN DTIC DDA 8725 JOHN J KINGMAN RD STE 0944 FT BELVOIR VA 22060-6218
1	DIRECTOR US ARMY RESEARCH LABORATORY ATTN AMSRL OP SD TA/ RECORDS MANAGEMENT 2800 POWDER MILL RD ADELPHI MD 20783-1197
1	DIRECTOR US ARMY RESEARCH LABORATORY ATTN AMSRL OP SD TL/ TECHNICAL LIBRARY 2800 POWDER MILL RD ADELPHI MD 207830-1197
1	DIRECTOR US ARMY RESEARCH LABORATORY ATTN AMSRL OP SD TP/ TECH PUBLISHING BRANCH 2800 POWDER MILL RD ADELPHI MD 20783-1197
1	DIRECTORATE FOR MANPRINT ATTN DAPE MR DEPUTY CHIEF OF STAFF PERSONNEL 300 ARMY PENTAGON WASHINGTON DC 20310-0300
1	CODE 1142PS OFFICE OF NAVAL RESEARCH 800 N QUINCY STREET ARLINGTON VA 22217-5000
1	DR ARTHUR RUBIN NATL INST OF STANDARDS & TECH BUILDING 226 ROOM A313 GAITHERSBURG MD 20899
1	COMMANDER US ARMY RESEARCH INSTITUTE ATTN PERI ZT (DR E M JOHNSON) 5001 EISENHOWER AVENUE ALEXANDRIA VA 22333-5600
1	DEPUTY COMMANDING GENERAL ATTN EXS (Q) MARINE CORPS RD&A COMMAND QUANTICO VA 22134

NO. OF COPIES	ORGANIZATION
1	HEADQUARTERS USATRADOC ATTN ATCD SP FORT MONROE VA 23651
1	COMMANDER US ARMY MATERIEL COMMAND ATTN AMCAM 5001 EISENHOWER AVENUE ALEXANDRIA VA 22333-0001
1	COMMANDER USA OPERATIONAL TEST & EVAL AGENCY ATTN CSTE TSM 4501 FORD AVE ALEXANDRIA VA 22302-1458
1	CHIEF ARMY RESEARCH INSTITUTE AVIATION R&D ACTIVITY ATTN PERI IR FORT RUCKER AL 36362-5354
1	AIR FORCE FLIGHT DYNAMICS LAB ATTN AFWAL/FIES/SURVIAC WRIGHT PATTERSON AFB OH 45433
1	US ARMY NATICK RD&E CENTER ATTN STRNC YBA NATICK MA 01760-5020
1	DR RICHARD JOHNSON HEALTH & PERFORMANCE DIVISION US ARIEM NATICK MA 01760-5007
1	USAF ARMSTRONG LAB/CFTO ATTN DR F WESLEY BAUMGARDNER SUSTAINED OPERATIONS BRANCH BROOKS AFB TX 78235-5000
1	DR JON FALLESEN ARI FIELD UNIT PO BOX 3407 FORT LEAVENWORTH KS 66027-0347
1	USA TRADOC ANALYSIS COMMAND ATTN ATRC WSR (D ANGUIANO) WHITE SANDS MISSILE RANGE NM 88002-5502
1	STRICOM 12350 RESEARCH PARKWAY ORLANDO FL 32826-3276

NO. OF COPIES	ORGANIZATION
1	MR JOHN HUNT GE BLDG 148-301 ROUTE 38 MOORESTOWN NJ 08057
1	PURDUE UNIVERSITY SERIALS UNIT CDM KARDEX 1535 STEWART CENTER WEST LAFAYETTE IN 47907-1535
1	DR HARVEY A TAUB RSCH SECTION PSYCH SECTION VETERANS ADMIN HOSPITAL IRVING AVENUE & UNIVERSITY PLACE SYRACUSE NY 13210
1	DR ROBERT KENNEDY ESSEX CORPORATION SUITE 227 1040 WOODCOCK ROAD ORLANDO FL 32803
1	DR NANCY ANDERSON DEPARTMENT OF PSYCHOLOGY UNIVERSITY OF MARYLAND COLLEGE PARK MD 20742
1	DR BEN B MORGAN DEPARTMENT OF PSYCHOLOGY UNIVERSITY OF CENTRAL FLORIDA PO BOX 25000 ORLANDO FL 32816
1	DR LLOYD A AVANT DEPARTMENT OF PSYCHOLOGY IOWA STATE UNIVERSITY AMES IA 50010
1	DR MM AYOUB DIRECTOR INST FOR ERGONOMICS RESEARCH TEXAS TECH UNIVERSITY LUBBOCK TX 79409
1	PEO INTELLIGENCE & ELECTRONIC WARFARE ATTN AMCPEO IEW VINT HILL FARMS STATION BLDG 197 WARRENTON VA 22186-5115
1	PEO STRATEGIC DEFENSE PO BOX 15280 ATTN DASD ZA US ARMY STRATEGIC DEFENSE CMD ARLINGTON VA 22215-0280

NO. OF COPIES	ORGANIZATION
1	COMMANDER US ARMY MATERIEL COMMAND ATTN AMCDE AQ 5001 EISENHOWER AVENUE ALEXANDRIA VA 22333
1	COMMANDER MARINE CORPS SYSTEMS COMMAND ATTN CBGT QUANTICO VA 22134-5080
1	DIRECTOR AMC-FIELD ASSIST IN SCIENCE & TECHNOLOGY ATTN AMC-FAST (RICHARD FRANSEEN) FT BELVOIR VA 22060-5606
1	CDR & DIR USAE WATERWAYS EXPERIMENTAL STATION ATTN CEWES IM MI R (A S CLARK) CD DEPT #1153 3909 HALLS FERRY ROAD VICKSBURG MS 39180-6199
1	US ARMY RESEARCH INSTITUTE ATTN PERI IK (DOROTHY L FINLEY) 2423 MORANDE STREET FORT KNOX KY 40121-5620
1	CHIEF ARL HRED ARDEC FIELD ELEMENT ATTN AMSRL HR MG (R SPINE) BUILDING 333 PICATINNY ARSENAL NJ 07806-5000
1	CHIEF ARL HRED ATCOM FIELD ELEMENT ATTN AMSRL HR MI (A MANCE) 4300 GOODFELLOW BLVD BLDG 105 1ST FLOOR POST A-7 ST LOUIS MO 63120-1798
1	CHIEF ARL HRED CECOM FIELD ELEMENT ATTN AMSRL HR ML (J MARTIN) MYERS CENTER ROOM 3C214 FT MONMOUTH NJ 07703-5630
1	CHIEF ARL HRED MICOM FIELD ELEMENT ATTN AMSRL HR MO (T COOK) BUILDING 5400 ROOM C242 REDSTONE ARSENAL AL 35898-7290

NO. OF COPIES	ORGANIZATION
1	CHIEF ARL HRED AVNC FIELD ELEMENT ATTN AMSRL HR MJ (R ARMSTRONG) PO BOX 620716 BUILDING 514 FT RUCKER AL 36362-0716
1	CHIEF ARL HRED FT HOOD FIELD ELEMENT ATTN AMSRL HR MA (E SMOOTZ) HQ TEXCOM BLDG 91012 RM 134 FT HOOD TX 76544-5065
1	USARL HRED FIELD ELEMENT USAADASCH ATTN ATSA CD ATTN AMSRL HR ME (K REYNOLDS) 5800 CARTER ROAD FORT BLISS TX 79916-3802
1	CHIEF ARL HRED ARMC FIELD ELEMENT ATTN AMSRL HR MH (M BENEDICT) BUILDING 1109D (BASEMENT) FT KNOX KY 40121-5215
1	CHIEF ARL HRED USAFAS FIELD ELEMENT ATTN AMSRL HR MF (L PIERCE) BLDG 3040 ROOM 220 FORT SILL OK 73503-5600
1	CHIEF ARL HRED USAIC FIELD ELEMENT ATTN AMSRL HR MW (E REDDEN) BUILDING 4 ROOM 349 FT BENNING GA 31905-5400
1	CHIEF ARL HRED SC&FG FIELD ELEMENT ATTN AMSRL HR MS (L BUCKALEW) SIGNAL TOWERS ROOM 207 FORT GORDON GA 30905-5233
1	CHIEF ARL HRED FT HUACHUCA FIELD ELEMENT ATTN AMSRL HR MY BUILDING 84017 FORT HUACHUCA AZ 85613-7000
	ABERDEEN PROVING GROUND
5	DIRECTOR US ARMY RESEARCH LABORATORY ATTN AMSRL OP AP L (TECH LIB) BLDG 305 APG AA

NO. OF COPIES	ORGANIZATION
1	ARL SLAD ATTN AMSRL BS (DR JT KLOPCIC) BLDG 328 APG-AA
1	LIBRARY ARL BUILDING 459 APG-AA
1	CHIEF ARL HRED ERDEC FIELD ELEMENT ATTN AMSRL HR MM (D HARRAH) BLDG 459 APG-AA