UNITED STATES AIR FORCE
RESEARCH LABORATORY

DESKTOP TRAINING
FOR DECISION MAKERS

Fritz H. Brecke

SYSTEMS ENGINEERING ASSOCIATES
2851 Camino Del Rio South
Suite 410
San Diego, CA 92108-3843

Sharon K. Garcia

HUMAN EFFECTIVENESS DIRECTORATE
MISSION CRITICAL SKILLS DIVISION
7909 Lindbergh Drive
Brooks AFB, TX 78235-5352

December 1998

20010529 060

NOTICES

This report is published in the interest of scientific and technical information exchange and does not constitute approval or disapproval of its ideas or findings.

Using Government drawings, specifications, or other data included in this document for any purpose other than Government-related procurement does not in any way obligate the US Government. The fact that the Government formulated or supplied the drawings, specifications, or other data, does not license the holder or any other person or corporation, or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

The Office of Public Affairs has reviewed this report, and it is releasable to the National Technical Information Service, where it will be available to the general public, including foreign nationals.

This report has been reviewed and is approved for publication.

SHARON K. GARCIA, Ph.D.
Project Scientist

R. BRUCE GOULD, Ph.D.
Acting Chief
Mission Critical Skills Division

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE December 1998	3. REPORT TYPE AND DATES COVERED Final - February 1992 - February 1997

4. TITLE AND SUBTITLE

Desktop Training for Decision Makers

5. FUNDING NUMBERS

C - F33615-91-C-0007
PE - 62205F
PR - 1123
TA - C2
WU - 76

6. AUTHOR(S)

Fritz H. Brecke
Sharon K. Garcia

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

Systems Engineering Associates
2851 Camino Del Rio South, Suite 410
San Diego, CA 92108-3843

8. PERFORMING ORGANIZATION

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

Air Force Research Laboratory
Human Effectiveness Directorate
Mission Critical Skills Division
7909 Lindbergh Dr.
Brooks Air Force Base, TX 78235-5352

10. SPONSORING/MONITORING

AL/HR-TR-1997-0045

11. SUPPLEMENTARY NOTES

Air Force Research Laboratory Technical Monitor: Dr. Sharon K. Garcia; (210) 536-3989

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for Public Release; Distribution is Unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

This report describes the development of instructional strategies for training complex decision-making skills to Logistics Command and Control (LC2) personnel. It brings to a close an effort representing a coherent line of inquiry and prototype development conducted to identify a more affordable approach to training decision-making skills. The approach adopted consisted of training on a common desktop computer system. The report concludes with a summary of what has been learned from this effort, and what remains to be accomplished.

14. SUBJECT TERMS Logistics Command and Control (LC2) Decision Making Computer-Based Training			15. NUMBER OF PAGES 41
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL

NSN 7540-01-280-5500

Standard Form 298 (Rev 2-89) Prescribed by ANSI Std Z-39-18
298-102 COMPUTER GENERATED

Contents

Figures

Tables

Preface

This represents the final report describing the development of instructional strategies for training complex decision-making skills to Logistics Command and Control (LC^2) personnel. The work was performed in support of the Desktop Training for Logistics Command and Control (DDT/LC^2) research and development effort. The project was accomplished under Contract No. F33615-91-C-0007, with Systems Engineering Associates (SEA) of San Diego, CA. Management of this effort was provided by the Armstrong Laboratory, Human Resources Directorate, Technical Training Research Division, Instructional Systems Branch, Brooks AFB, Texas.

Summary

Since 1986 the Air Force has been concerned with the problem of finding an affordable instructional technology for training decision-making skills. This report is the final report for the second of two consecutive research and prototype development projects that pursued this goal.

The first project was conducted between November 1986 and April 1990. It addressed the problem of training decision makers working in Air Force Tactical Command and Control. The second project was conducted between February 1992 and February 1997. It addressed the training of decision makers in Air Force Logistics Command and Control.

The two projects together represent a continuous and coherent line of inquiry and prototype development that was conducted in two different decision-making domains but aimed at the same general goal: affordable training in decision-making skills delivered on common desktop computers. This report summarizes what has been learned from the efforts in both projects.

Introduction

The military has an interest in ensuring that decision makers at all levels perform in a highly proficient and reliable fashion. This is especially true during wartime when decisions can spell life or death, victory or defeat. Proficient, reliably good decision making is a high value skill in many human endeavors and effective training in this skill is likely to produce high returns. The military's urgency to provide training in decision-making skills, is driven in large part by the fact that peacetime jobs of war fighters frequently do not afford them opportunities to practice decision-making skills needed during wartime. This is because: (a) they are performing jobs completely different during peacetime; or (b) the conditions of performance are very different.

Providing decision-making skills training in an economical fashion is difficult. One reason is the lack of a solid theoretical foundation for designing such training. Decision making is a complex cognitive process heavily influenced by affective factors (Figure 1). Any of these factors can make the problem more difficult, and in most cases several (if not all of them) are at work in real time decision-making situations. Consequently, it has been a challenge to both understand this intellectual and affective process <u>and</u> to teach or train it.

Figure 1. Factors Influencing Decision Making

Nevertheless training in decision-making skills must be accomplished. To date, the primary training solution has been brute force, high fidelity field exercises within the decision-making environment. These field exercises are basically elaborate simulations of wartime scenarios. The simulations have a number of drawbacks as training grounds for decision makers (Brecke, Jacobs, & Krebs, 1988). They are expensive, infrequently used, and when used, are relatively short-term exercises. As such, not every "decision maker" has the opportunity to participate. Even if they did, a brief exposure to training for three or four days is unlikely to result in significant performance improvement. In addition, there is the issue of maintaining the decision-making skills learned over the long term.

The basic premise of the projects reported herein is the notion that the training availability problem can be solved by developing decision training software that could be delivered on the desktop computers of present and future decision makers. Such a "trainer-in-a-box" would help individuals: (a) in preparing for the decision maker's job; (b) who must make relatively critical decisions; and (b) to build, sharpen and/or maintain their decision-making skills either singly or in concert with colleagues located elsewhere.

Two consecutive projects were undertaken to build various forms of the decision-making trainer. In the first project, the emphasis was on the application of state-of-the art technological capabilities (Brecke & Young, 1990). The second project focused on instructional guidelines. An extensive theoretical foundation led to the derivation of a number of skill-specific instructional design heuristics (Brecke & Garcia, 1995). The prototype training systems developed in these two projects represent various solutions to the decision training problem. They encompass a broad range of technological approaches, differing levels of sophistication and adherence to instructional design heuristics. Each solution has drawbacks, and no one can be said to provide an ideal solution to the problem of providing economical training in decision making. However, each proposed solution does expand upon the set of tools by which such a training problem can eventually be solved.

2

This technical report builds on two earlier and much more detailed reports (Brecke & Young, 1990; Brecke & Garcia, 1995). It traces a logical sequence of issues that need to be considered when building an affordable desktop decision trainer. To do so is considered an instructional design problem, in which four design parameters are more or less well defined: the goal, the learners, the delivery system, and the environment; and where the remaining parameters, the content and the instructional strategy, are to be determined (Merrill & Wood, 1974; Frank, 1969).

The goal of both projects was to find answers to two questions:

- What content must be conveyed?
- By what instructional strategy is it to be delivered?

Ideally, the features of an intelligently designed training product should be justifiable in terms of a theoretical and empirical research base. This research base has to encompass at least three aspects: (a) the nature of the task to be performed and learned; (b) the features of human task performance; and, (c) the process of learning the task.

The process of learning the task addresses three additional questions:

- What is a decision?
- How do people make decisions?
- How do people learn to make decisions?

It is theorized that if these three questions can be answered conclusively, then the questions of content and strategy could be answered as well. Our research confirmed this to be partially true (Brecke and Garcia, 1995).

This report begins with a review of the theoretical issues surrounding the three questions: (1) What is a decision? (2) How do people make decisions? and (3) How do people learn to make decisions? This provides the basis for the next section which addresses the two instructional design issues of content and strategy. The third section features brief descriptions of the salient features of the research prototype systems that were built in the course of the two projects. Finally, the last section discusses the questions of applicability and generalizability of the work.

Theoretical Issues

The Nature of the Task

The first theoretical issue pertained to the nature of the task to be trained; namely, "What is a decision?" Simply put, a decision is a solution to a problem. The problem (in its most

general form) is to figure out what to do next, even though one does not know everything one should know or would want to know.

At the core of this deceptively simple answer is the notion that decision problems are a class of problems for which optimal solutions cannot be found a priori, due to the insufficiency of available information. The intellectual labor of solving this class of problems consists primarily of activities that aim to reduce uncertainty by gathering and assessing whatever information is available. The recognition of uncertainty reduction as one of the central features of the task led to the development of a taxonomy of uncertainties in decision-making problems (Brecke & Garcia, 1995). In many cases the labor of uncertainty reduction is strongly influenced by time constraints. This notion led to the development of a "time table" for decision-making problems (Brecke & Garcia, 1995). We have called these two analytic products the "Uncertainty Model" and the "Timeline Model" respectively. Since these models are important to the rest of this report, they are briefly described below.

Uncertainty Model

The Uncertainty Model (Figure 2) is a classification tree for the various types of uncertainty found in decision problems.

Figure 2. Uncertainty Types in Decision Making
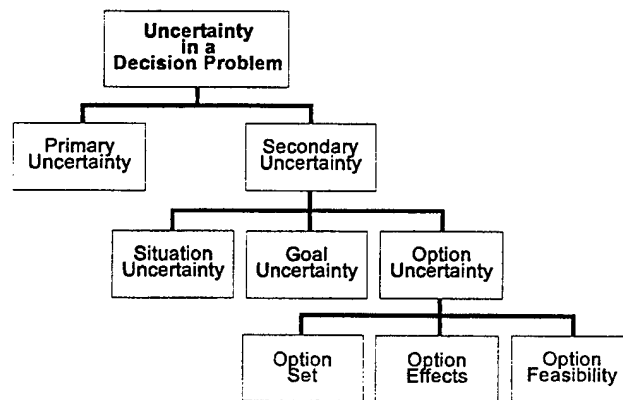
The primary function of a decision is to commit to some course of action. This can also be expressed in language borrowed from information theory: a decision reduces uncertainty about what to do to zero. We call uncertainty about "what to do" primary uncertainty. To arrive at the point where a decision can be made, the decision maker has to be both motivated by some form of pressure and has to work on reducing uncertainty about the pros and cons for any available course of action. We refer to uncertainty about "why" secondary uncertainty.

4

Secondary uncertainty includes uncertainty about the situation, goals to be achieved and courses of action or options that are available. The latter has three components: uncertainty regarding the set of possible options (e.g., "Isn't there another way?"), the feasibility of any known option (e.g., "What are the risks?"), and the effects of any known option (e.g., "How close will we get to the goals?").

Once secondary uncertainty has been reduced, commitment to a particular option can occur. As soon as that commitment occurs a decision is made and primary uncertainty is reduced to zero. The residual secondary uncertainty which had to be "overcome" by the decision (by the commitment) will only become zero after feedback about the effects of the implemented decision arrives.

Timeline Model

There is usually a strong temporal aspect to a decision problem. To describe this aspect, we have come up with a timeline model (Figure 3).

<u>Figure 3</u>. Timeline Model of Decision Making
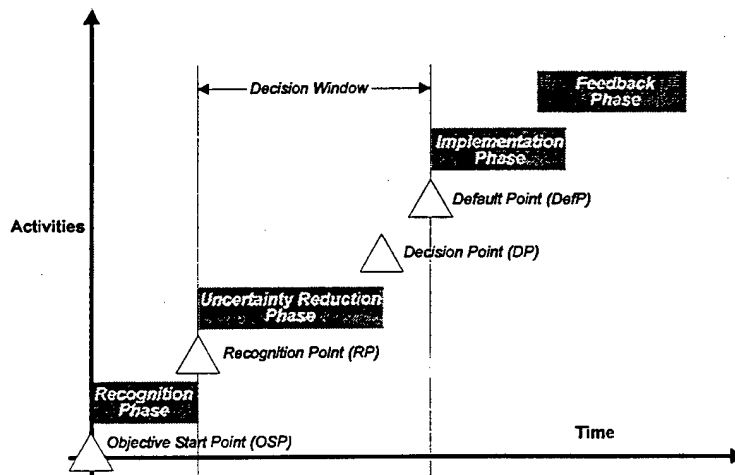
The decision-making task begins with the Objective Start Point (OSP) - the point in time in which a situation has developed and the need and opportunity for a decision has emerged. Once recognition has occurred, the decision maker begins to: (a) gather information to reduce the various aspects of secondary uncertainty; (b) makes a decision; and, (c) begins to implement the chosen course of action. All this must occur within the time limit imposed by the default point where the window for action closes. If that window is short (either because things are moving fast or recognition took too long or both), commitment may have to occur even though a great deal of secondary uncertainty remains.

5

Significance of the Models

The models, together with the definition, provide conceptual tools that make the decision training problem more tractable. For example, one might examine a type of decision task where options must be developed, and contrast it with another type of decision task where the options are clear but the goals are not. The task type requiring option generation relies on creativity during task performance. Training such a task would require a training system capable of evaluating novel approaches to a class of decision problems - a capability not easily developed in a computer-based system.

In the second task type, the decision maker has a clear set of options (e.g., there is certainty that all options are being considered, how feasible they are and what their effects will be) but the decision maker simply doesn't have the goals firmly articulated and prioritized. This task type requires no creativity or originality and a training system would merely have to be able to recognize a correct choice from existing options - a capability easily realized in a computer-based training system.

These two types of tasks obviously require fundamentally different cognitive activities to reduce secondary uncertainty as well as different learning and instructional strategies during training. It therefore seems fair to say that qualitative differences in decision-making tasks have significant impact on training design. The same is true for quantitative differences. For example, a decision task that is easy under normal circumstances, may become a substantially different and much more difficult task when the amount of uncertainty suddenly rises above normal levels (i.e., when a telephone line gets cut). Howard (1968) and Brecke (1975) point to the amount of uncertainty as an important dimension determining the nature of a decision-making task (Howard) or of a judgment task (Brecke). These researchers see uncertainty as one of three dimensions that together determine the nature of a decision-making task. The other two dimensions are identified as complexity and time constraint (Figure 4)
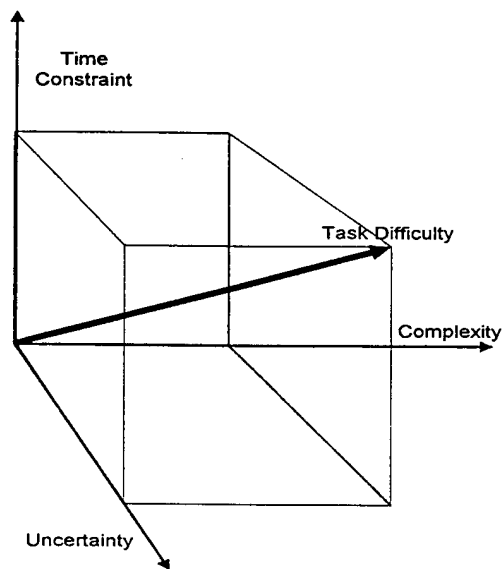

Figure 4. Dimensions of the Decision Task

6

The Uncertainty and Timeline models are guidelines for analyzing tasks to be trained. They result in distinctive profiles that express qualitative and quantitative features to be emulated by training tasks. The models are tools for ensuring that training tasks are structured like real world tasks to facilitate a positive transfer of training. We can use them and other factors to select tasks for training and once selection has occurred, to increase the chances that training will be effective.

The second significant aspect of the models is their use for training content. It is assumed that a person who is knowledgeable about the general features of the decision-making task will be a better decision maker earlier and will have an easier time achieving higher levels of performance. In this context the models represent meta-knowledge and are assumed to function as facilitating foundations for the accelerated formation of cognitive strategies.

Feedback in Decision-Making Tasks

The role of feedback in decision making presents a particularly interesting research problem. Feedback in real decision making can only occur after the implementation of a decision, and is therefore called "a posteriori" feedback. In training tasks it can occur before the start of implementation and is referred to as "a priori" feedback. A priori feedback addresses the issue of whether the decision maker has made a decision that is both optimal and logically consistent with the information that was available prior to implementation. A posteriori feedback provides information regarding the effectiveness of the decision (i.e., whether things actually worked out the way the decision maker intended). Since a priori feedback does not occur in the real world it has been called "artificial feedback " as opposed to natural, a posteriori feedback (Brecke, Hays, Johnston, McGarvey, Peters, & Slemon, 1989). Natural feedback may or may not actually occur in the real world, and if it does it is often confusing and difficult to attribute to a particular prior decision. The generation of natural feedback (i.e., feedback that has all the salient characteristics of a posteriori feedback as it occurs in the real world), is therefore usually difficult in a simulation environment.

The question then is what kind of feedback might be most useful in training? Nickerson and Feehrer (1975) contend that in real world situations, decision makers are often evaluated on the basis of a posteriori feedback (i.e., on the basis of results). This appears to be less than reasonable, since decision outcomes in real world domains are usually subject to many factors that are beyond the decision maker's control. Real world domains are basically open loop systems and in open loop systems the only appropriate manner to evaluate a decision maker is on the basis of a priori feedback. A posteriori feedback is perfectly fine in closed loop systems where all factors are under the decision maker's control, but such systems are rarely encountered in the real world and especially in the real world of armed conflict.

It is frequently the case that people have to learn to make high risk, high frequency decisions in a real world domain (i.e., an open loop system) without the benefit of instruction. The only type of feedback available to them is a posteriori feedback and when such feedback

arrives it may not always be clear what it means and/or which earlier decision it belongs to. It is quite possible that a person under this set of circumstances may never enter a learning process, much less complete it successfully. This situation can be improved through instruction which provides: (a) a priori feedback; (b) strategies to discover and properly assess a posteriori feedback; or (c) both.

Human Task Performance

The next step is to include the human performer in the equation. It is assumed that if it is possible to uncover the natural human performance methods for this type of task, then one can also identify the strengths and weaknesses of these methods. If this can be accomplished, then it may be possible to devise training guidelines that enhance the strengths while modifying, avoiding, or extinguishing the weaknesses.

This assumption is diametrically opposed to the one represented by the prescriptive school of thought that dominated the research up to 1975, when Nickerson and Feehrer published their landmark report. This view held that the real human performer is inadequate when compared to an abstract ideal performer and that training should endeavor to teach methods and procedures of task performance that are directly deduced from abstract analyses of ideal performance. This is of course a hopeless endeavor. Humans are not built to function that way and no amount of training will make a difference in this respect.

The new school of thought is represented by the work of a number of researchers investigating the concept of "Naturalistic Decision Making" (Klein & Calderwood, 1990). The fundamental notion of naturalistic decision-making research is that humans adapt their cognitive apparatus to the requirements of the decision-making task.

Between 1985 and 1988, Klein and Calderwood (1990) explored decision making in operational settings using a combination of field studies and experiments to test specific hypotheses. The decision domains included urban fire ground commanders, wildland fire incident commanders and U.S. Army tank platoon commanders. Their most important findings were in many cases at odds with the traditional, prescriptively biased, concept of the decision making process, and added considerable depth to a purely descriptive view of decision making.

During the same time frame, Noble and his associates (Noble, Grosz, & Boehm-Davis, 1987) conducted a series of studies designed to examine the development and use of schemata in decision making. Starting with a general model of decision making proposed by Lawson (1987), they distinguish two modes of decision making which they call "Rational Outcome Calculation" (ROC) and "Recognition-based Decision Making" (RB). The former is defined as a "rational process of explicitly comparing options and choosing the optimal alternative" while the latter is a mode where "the decision seems to follow directly from a recognition of the type of situation and a recollection of what actions usually work well in this kind of situation." Noble et al., see these modes as the two extremes of a continuum of decision-making strategies where the space between these poles is occupied by what they call "hybrid" strategies.

Klein's Recognition Primed Decision Making and Noble's Recognition Based Decision making are clearly related. Hammond (1986) makes essentially the same distinction when he refers to "analytic" versus "intuitive" modes of decision making. A similar distinction is also made by Rasmussen (1988) who differentiates between rational, knowledge-based decision making and heuristic, rule-based decision making and indicates that the former process is used by novices (and by experts facing unfamiliar situations) and that the latter is "applied by skilled actors."

The picture of the human decision-making process that arises from the combined results of this research can be summarized as follows:

- There are two distinct and different pure forms of the decision-making process which form opposite ends of a continuum that is populated by mixed or "hybrid" processes. One form is the rational outcome calculation (ROC) process which involves conscious, analytical, knowledge-based processing within an overall logical, systematic framework. The other form is the recognition-based (RB) process which involves unconscious, intuitive, heuristic-based processing.

- The ROC process is a weak, general, rational problem-solving procedure which is probably minimally adapted to the gross features of any particular decision-making domain and which is to some extent modified by prior experience in related or similar domains. The initial concern is reduction of situation uncertainty which, in Nickerson and Feehrer's (1975) terms, probably involves processes of information gathering, data evaluation, problem structuring, hypothesis generation and hypothesis testing. The second concern is the reduction of uncertainty regarding goals to be achieved by the decision. This probably involves either a more or less explicit clarification of personal goals or consultation of external, "public" sources of guidance or both. The third and really central concern is the reduction of uncertainty concerning options. This involves generation of the option set, assessment of the possible option effects in view of the goals to be achieved and finally an assessment of the feasibility of each option. The ROC process is generally performed with the intent of achieving a decision that is logically consistent with the available information and optimal as far as goal achievement and feasibility is concerned.

- The RB process is characterized by a rapid or near instantaneous aggregation of situational data into a situational assessment. The situation is identified as a member of a problem class which is defined by a prototype (e.g., a schema). Simultaneous with the situation recognition appears the prototypical solution (also a schema) for members of this situation class. Deviations of the actual situation from the prototype situation are recognized and provide cues for adapting the prototypical solution to the current actual situation. In general, experts are first concerned with the reduction of situation uncertainty. This initial problem solved, they "see" the problem and a fairly specific solution for it at the same time and spend the remainder of the available decision time primarily with reducing uncertainty associated with the feasibility of the

one option they are considering. The primary mechanism for reducing this type of uncertainty appears to be something like imagery-based simulation using some sort of "runnable" mental model. Uncertainty regarding goals to be achieved by the decision commonly does not appear to be a concern. Experts are familiar with the general goals of their job and "know" what goals can or cannot be achieved in a specific situation. The RB process is generally performed with the intent of satisfying goals and requirements rather than optimizing their achievement.

The degree to which one or the other process is used on a given occasion is a function of decision task characteristics and decision maker characteristics. Keeping task characteristics constant, novice decision makers will operate closer to the ROC end of the continuum and experts will operate closer to the RB end of the continuum. Keeping experience constant, decreasing time pressure and increasing task complexity, increasing task risk (consequences of bad decisions) and increasing uncertainty will increase the use of ROC processes over RB processes and vice versa.

This is certainly a richer, more comprehensive, and realistic view of the decision-making process than that produced by the earlier, prescriptively contaminated research, but it is not so much an alternative view as it is a complementary one. The concept of option assessment and choice which was central to earlier views of decision making is confirmed as a feature of the process engaged in by novices. It is also central when the task is very important and when there is ample time to perform it. But when experts are performing under time pressure option assessment gives way to concerns of situation recognition and feasibility assessment[1].

Given these performance mode characterizations, the first aspect that training designers must consider is whether the desired end performance is closer to the ROC end or the RB end of the performance spectrum. The second is to factor into the training design, the deficiencies and limitations of the human performer.

If the target performance is RB decision making, then the focus of training must be on: (a) situation recognition; (b) the heuristic matches between situation and solution classes; (c) rapid identification of satisficing rather than optimizing solutions; and, (c) the skills of mentally rehearsing the implementation of solutions. This kind of training must by necessity be very

---

[1]   However, the process is still not described in its entirety. Little is known about how a decision maker recognizes the need for a decision, what factors influence the duration and reliability of the recognition process, and what happens after the decision is made (after the DP or Decision Point). For example, very little, if anything, is known about how the "embedded" decision on when to start implementation is influenced by task or skill factors. It is also still somewhat of a mystery how people use feedback during or after implementation, what constitutes useful feedback, and how feedback is employed in decision chains as opposed to singular, independent decisions. However, the crucial mid-phase, the phase that encompasses the central task of decision making, has become better understood.

domain specific. It also must provide large numbers of practice cases in highly realistic settings (including stress inducing factors) to achieve an effective transition from an initial ROC mode of performance to a solidly established RB performance mode.

If the target performance is ROC decision making, then the focus of training must be on: (a) general and domain related methods of uncertainty reduction; (b) generation of option sets; and, (c) classical option comparisons. This kind of training is less bound to the specific technical domain in which the decision-making function must be exercised. Training can be much shorter especially when the target environment permits the use of job aids.

Finally, training must address the issue of human deficiencies and limitations. Deficiencies, such as the tendency to be overly conservative in the use of probabilistic information, can presumably be corrected or "trained out." Limitations, such as the inability of most people to weigh more than some small number of factors cannot be corrected and thus can only be "trained around" or helped by job aids.

Task Learning and Performance Improvement

We now turn to the issue how this type of task is learned. Here, as in the preceding section, the ideal answer is a comprehensive descriptive model that accounts for the entire learning process from the state represented by people who are novices in a particular domain to the state represented by domain experts.

Multi-Stage Skill Acquisition Theories

The volume of existing theoretical and empirical research relative to skill acquisition theory is exceedingly large. We concentrated on multi-stage skill acquisition theories because identifiable stages in the learning process offer the possibility to precisely tailor instructional treatments to each stage. This results in faster, more efficient learning than with a single undifferentiated treatment.

We examined a number of multi-stage theories (Forbus and Gentner, 1986; Rasmussen, 1986; Anderson, 1982; Dreyfus and Dreyfus, 1980; Gentner, 1980; Siegler, 1978; McDermott and Larkin, 1978; Fitts and Posner, 1967) and found commonalities, both in the dimensions that were used and in the changes that were described for these dimensions. In other words, the commonalities went beyond the common theme of a focus gradual process marked by stages. It was therefore believed that one could, without excessive distortions, fuse the essential themes in each theory into one common, unified model, which would come closer to descriptive truth than each of the constituent models alone.

This work eventually led to a synthesis of two concepts: the concept of a prerequisite order in the learning process and the concept of multiple stages. From our point of view, these two concepts are logically related. The prerequisite concept holds that learning of a particular skill can only begin after other skills which are prerequisites to that skill have been learned at

11

least to some extent[2]. The multi-stage concept also represents that very same notion but goes beyond the prerequisite idea by postulating qualitative and quantitative changes not just in what type of information can be processed but in how it is processed. For example, processing changes from serial to parallel, from conscious and analytical control to non-conscious, intuitive control, and from very laggard and laborious to very fast and facile.

The prerequisite notion has been captured in a widely accepted two-dimensional matrix by Merrill (1983). This matrix relates a "Type of Content" dimension to a "Level of Performance" dimension and consists, for the sake of operational utility, of 10 discrete boxes. The boxes are simplifying constructs with somewhat fuzzy borderlines which make it practically and theoretically easier to associate particular instructional treatments with particular intersections of type of content and level of performance. This matrix can be seen as a map of intellectual skills or a "learning surface."

When viewed from this perspective, one can see that this learning surface has to be traversed along a central diagonal vector that originates in the "REMEMBER-FACTS" cell and points towards the "FIND-PRINCIPLE" cell. A learner cannot acquire a concept on the remember level unless they have first acquired its constituent facts. Further, a learner cannot learn to use a concept unless they have first acquired the concept on the remember level, and so forth (see Figure 5). This central vector leads from simple, disjointed knowledge structures to more and more complex and integrated knowledge structures, and from surface knowledge to increasingly deep knowledge.

---

[2]We avoid the term "mastered" in this connection, because we are of the opinion that complete mastery, however defined, is not necessary for a successful start in learning the next higher level skill.

Figure 5. Simplified Multi-Stage Learning Model
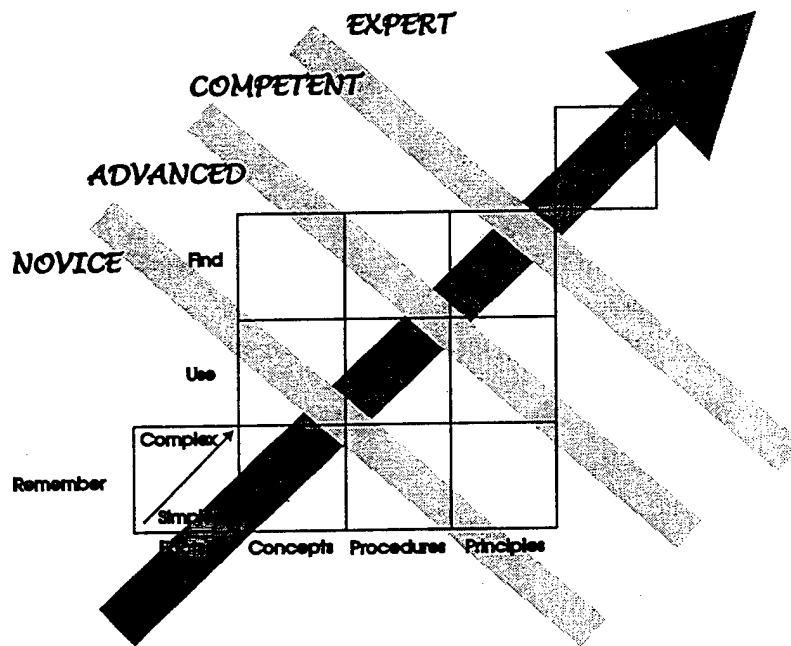
The key to the learning model is the notion that stages in learning must be defined as partitions of the learning surface. These partitions are not sets of cells but bands of skill aggregates which run orthogonal to the central vector of the learning surface. Secondly, within these bands and within the cells included in these bands, one can further postulate a complexity dimension that runs parallel to and in the same direction as the central vector. This dimension simply indicates that simple skills must be mastered before more complex skills. Finally, a third essential notion is that the performance-content matrix does not include what is generally known as expertise. We agree with Dreyfus and Dreyfus (1980) that the expert "no longer needs principles" and assert that expertise is in it's own category, its own box, which we have attached to the upper right hand corner of Merrill's matrix.

Based on these ideas the learning process for decision-making skills can be divided into four stages. These stages synthesize the prerequisite notion with the multi-stage notion. They are, like the cells of Merrill's matrix, simplifying constructs, which should allow us to associate instructional design guidelines with them. The stages represent levels of decision-making skill that are achievable by the learner given that they have learned or are learning certain types of content to certain levels of performance. In general, as the learner proceeds through these stages in the direction of the central vector, knowledge structures become increasingly complex and integrated, and knowledge processing changes from a slow, serial, analytical, rational outcome calculation mode to a fast, parallel, intuitive, recognition-based mode. In more specific terms, we can characterize decision-making performance on each of the four stages as follows:

- At the NOVICE stage, decision-making performance is constrained by the lack of domain specific knowledge. The novice is able to solve simple decision problems by applying weak general methods to a domain where they are just beginning to acquire some surface knowledge and are able to use some of the simpler concepts of the domain. The novice learner is slow and unreliable in recognizing decision problems and cannot judge the available decision window. The ability to discriminate between salient and non-salient features of the situation or context is rudimentary. The novice cannot prioritize uncertainty reduction requirements and has no domain-specific strategy for uncertainty reduction. Only incomplete option sets can be generated and there is as yet no concept of hierarchical classes of options (i.e., the option sets are not ordered in any sense). Work is performed exclusively in a ROC mode.

- At the ADVANCED stage, decision-making performance centers around the Use-Procedures cell. The learner is now more proficient in using concepts and in finding and defining concepts. They begin to invent their own way of doing things (Find-Procedure cell) and develop a surface understanding of the principles governing the domain (Remember-Principle cell). Recognition of decision problems is timely and reliable and judgment of available decision time is usually correct. The advanced learner begins to recognize and utilize salient situation or context features. They are able to solve moderately complex decision-making problems in a manner where weak general methods are increasingly supplemented and/or replaced by low level, domain specific methods. Advanced learners begin to prioritize uncertainty reduction requirements and use low level, domain specific uncertainty reduction strategies. They develop complete option sets and can identify major classes of options. The learner still works predominantly in a ROC mode, but there are instances of the RB mode being employed with very simple and very frequent decision problems.

- At the COMPETENT stage, the learner has become a reliable performer. They put the finishing touches on their ability to use principles and achieve competence in inventing/finding procedures or rules. They are now fully capable of focusing on situational or context features that are germane to their problem and in modifying options to fit these features. The competent learner recognizes decision problems quickly and reliably, knows exactly how much time they have and manages time well. They develop well ordered and complete option sets and become efficient in pruning branches or classes with low feasibility, undesirable effects, and excessive uncertainty early on. The learner can solve complex decision-making problems and do so by strong, domain-specific methods. They are fully competent in prioritizing uncertainty reduction efforts and use coherent, fully developed, domain specific strategies to reduce uncertainty. The predominant working mode is still the ROC mode, but instances of RB-decision making are becoming more frequent, especially with simple and moderately complex decision problems that are familiar, have tight time constraints and involve relatively low stakes.

14

- At the EXPERT stage, learning is no longer significant. The degrees of automaticity, speed and reliability continue to rise, albeit more slowly than before. The learner has essentially turned into an accomplished performer. Recognition and timing have ceased to be issues to be concerned with. They are able to judge and apply contextual features intuitively and instinctively and thus can prune an option tree down to the "first, best" option without ever explicitly constructing the entire tree in the first place. In most cases, a satisfactory option simply comes to mind at the same time as the decision problem is recognized. The learner is extremely efficient in uncertainty reduction and can "see" entire situation-goal-option patterns based on minimal informational cues. Uncertainty reduction for the expert consists essentially of probing for information to validate a perceived pattern, and of mental simulation (or rehearsal) to ascertain feasibility of the primary option, which is frequently the only option under consideration. The decision maker now works predominantly in the RB mode. They may resort to the ROC mode if time permits, if the stakes are high, and if the problem is new and unusual.

Given the usual time and resource constraints under which training has to occur, the objective of training must be to accomplish the transitions from one stage to the next in the most expedient, efficient and effective way possible.

Transition 1, from NOVICE to ADVANCED, is primarily concerned with establishing a foundation of domain specific knowledge and with the development of domain specific procedural skills. This is not a challenging problem for training design and development. Transition 2, from ADVANCED to COMPETENT, on the other hand, is concerned with refining, integrating and accelerating domain specific procedural skills and with the development of the performer's ability to apply knowledge of underlying domain principles to practical decision problems. This is a fundamentally different learning requirement and the primary challenge that both projects had to address. Transition 3, from COMPETENT to EXPERT, requires massive amounts of practice in the actual job environment or in a very high fidelity simulation of that job environment. This transition is usually either left alone, or facilitated with some form of on-the-job training (OJT). On rare occasions (i.e., flight training), high fidelity simulations are used to get the transition started, but even in those cases, cost prohibits full support of this transition in a training environment.

Each transition presents different learning requirements and therefore different instructional treatments or methods. In both projects, the choice of instructional methods was constrained by an a priori media choice (i.e., the requirement to provide instructorless training on a desktop computer).

Computer-based instructional media can provide instructional treatments in two basic forms. The first is what is commonly referred to as Computer Assisted Instruction or CAI. The second is simulation. CAI implies an instructional environment where a carefully structured sequence of lessons and lesson segments provides training in the building blocks of some terminal skill. CAI also implies that explanations, cues and memory aids are provided and it implies more or less extensive practice and feedback in component skills. Simulation on the

15

other hand implies an emphasis on part or whole task practice in a job-like environment. Explanations, cues and memory aids are usually absent in simulation per se[3].

Given the learning requirements for the first transition, we can broadly assert that some form of CAI would be the instructional method of choice, and with equally general reasoning we can assign some form of simulation as the method of choice to the second transition. Simulation would not work for the first transition because of the need to accommodate REMEMBER-level objectives and because of the need to train PROCEDURES. CAI would work to some extent for the second transition, but it would be very difficult if not impossible to provide a realistically dynamic decision-making environment with this method.

Summary of Theoretical Issues

For our purposes, the work on theoretical issues performed during the two projects was productive. In 1975, Nickerson and Feehrer found a large body of work concerned with decision making, but very little of it provided useful guidance for a training designer. It appears that the situation has changed during the last 20 years. Strong guidance in the design of training tasks can be obtained from the two models presented earlier. The recognition that the prescriptive school of thought presents unrealistic demands on the human performer and the clear recognition and distinction between two natural modes of human performance provides guidance in the definition of achievable and task-appropriate training outcomes. Multi-stage learning models provide guidance in focusing on the critical training design issues. The application of this new guidance to the instructional design problem is described in the following section.

Instructional Design Issues

As previously indicated, the crucial issues in solving the instructional design problem are the definitions of the content and strategy parameters of the training system to be designed. The results of the theoretical analyses performed during the two projects were helpful in bringing these issues closer to a definitive resolution.

Content Issue

The instructional content is the knowledge the student requires to perform the task and/or to learn to perform the task. The issue of what that content should be can be formulated as a choice between domain-specific heuristics, meta-heuristics, and/or some combination thereof.

Domain-Specific Heuristics

The term domain-specific heuristics denotes the specialized task knowledge employed or deployed by domain experts in the act of decision making. For example, an expert fire

---

[3]Although they may be present in briefings before and after the actual simulation sessions.

commander employs "rules of thumb" about the use of fire retardants to combat various types of fires. A physician deploys knowledge about the efficacy of procedures and medicines to decisions on what therapy should be employed to various types of pathologies.

One way to formulate this type of knowledge is in the form of production rules where the expert essentially makes a match between a problem situation class and a solution class. Clancey (1987) calls this a heuristic match. To make this match the expert must first abstract from the particular situation to a class of problems, match this class of problems with a class or type of solutions, and then instantiate the solution class with a particular solution tailored to the particular situation.

If this type of content is to be used in training, then it must exist in explicit form. The problem is that this type of task knowledge is usually not written down anywhere. This means it has to be extracted from experts and that is always a difficult, time-consuming and therefore costly undertaking, even if one has plenty of access to willing and articulate experts.

Meta-Heuristics

The second content choice is meta-heuristics: knowledge about decision making in general. This type of knowledge consists of guidelines or principles that can be employed to manage the deployment of one's own cognitive or material resources in the service of a decision-making task.

This is fundamentally different content from domain-specific heuristics: it deals not with rules of thumb for some specialized domain but with the problem of managing resources for decision making in any domain. It is feasible to adapt very general meta-heuristics to a particular domain and to render them more domain related, however the fundamental difference remains: domain-specific heuristics are task oriented while meta-heuristics are resource oriented. The examples provided in Table 1 illustrate this point. The domain-specific heuristic addresses the operational situation in the decision-making domain. The meta-heuristics address the situation in the decision maker's head whether they are very general or more closely adapted to a particular domain.

Table 1. Content Examples

Content Type	Content Examples
Domain-specific heuristic:	If the time is too short for normal sourcing and lateral sources are available, use the least impacted lateral source.
Meta-heuristic:	If the situation is not clear to you, acquire any missing information on the causes and scope of the problem.
Meta-heuristic adapted to the domain:	Missing causal and scope information can be obtained from requesting agencies, item managers, manufacturers etc.: from any agency that uses, services, stores, buys, or develops the type of item in question.

The question is which of these types of content is most helpful to a student who needs to learn to make decisions in a specific domain? To some degree this is an economical issue: extracting domain-specific heuristics from experts is difficult, time consuming, and therefore usually quite expensive. It also has other, very significant problems: it is full of local differences, i.e., it is very much tied to particular places and positions, and it is very unstable, i.e., it changes very frequently. If the training objective is recognition-based performance, then the content must be domain-specific and planning for the development of such training must factor in a sizable expense for the acquisition of that content and for its continual update.

However, if the training objective is the kind of performance characterized as rational outcome calculation, then a set of meta-heuristics that are more or less adapted to the target domain might suffice. This type of content is much less expensive to come by not only because it can be leveraged over many domains, but also because the two models described earlier provide useful starting points.

If the acquisition of domain-specific content is out of the question for budgetary reasons, then one might consider using a "proxy" domain for training. The choices are either a domain that is very stable and well known or one that is artificially created. The advantage of a stable, well known domain is that the students might already possess the prerequisite domain knowledge. The question is whether training in a proxy domain, be it a natural or an artificially created one, would transfer to the target domain?

The argument is that people can be trained to be better decision makers in the real world through training in a proxy world, as long as the training tasks in the proxy world and the real world tasks have the same critical features. The Uncertainty and Timeline models help ensure that training tasks and job tasks are indeed rigorously similar in all essential aspects. As long as this similarity can be maintained, there is no reason to expect anything other than positive transfer of training from a proxy world to the real world. The question is not the direction, but the degree of transfer and thus the cost vs. benefit ratio.

Instructional Strategy Issue

The second instructional design issue concerns strategy. In order to design a training regimen we require an instructional strategy, more specifically an organizational strategy (Reigeluth, 1983) articulated on two levels: the macro and the micro level.

Macro Instructional Strategy

For the macro aspects of organizational strategy we rely on Reigeluth's Elaboration Theory (ET) as the primary organizational principle. This theory integrates and synthesizes instructional sequences proposed by a number of other researchers into an internally consistent set of prescriptions that are guided by the goal of building stable cognitive structures in a meaningful, subsumptive (Ausubel, 1968), or assimilative (Mayer, 1977) way (Reigeluth, 1987). It provides an admirable level of operational guidance for instructional designers in selecting, sequencing, synthesizing and summarizing instructional content. It explicitly addresses interrelationships within instructional content and is the only theory that "specifically allows for some learner control over the selection and sequencing of the content" (Reigeluth, 1987, p. 246).
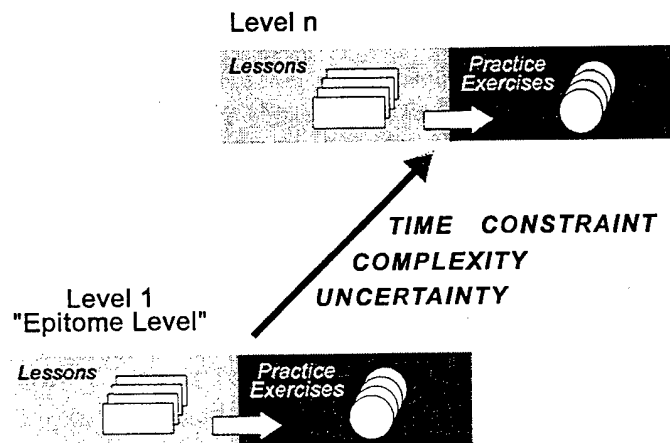
18

The key prescription of Reigeluth's ET is the elaborative sequence. This prescription makes intuitive sense, has considerable theoretical and empirical support (Reigeluth, 1983), and fits the instructional problem at hand. The entire content of a course of instruction in decision making can be organized in levels where the student learns on the first or epitome level to make the simplest types of decisions that are possible in a given domain and where they learn to make more and more complex decisions on subsequent levels of elaboration. The notion of levels of elaboration is compatible with the multi-stage learning model presented earlier.

The macro aspects of organizational strategy can be organized such that learners progress on each level from the NOVICE stage to the COMPETENT stage. The learner first becomes a competent decision maker in the simplest and most "benign" version of a particular domain and then progresses on to the next level, where they are exposed to a more complex and less benign version of the same domain, and so forth.

On each next higher level, the learner is once again a NOVICE with respect to new facts, concepts, procedures and principles that they must first acquire and then apply to the more complex, more uncertain and more time-constrained decision-making situations on that level, until they are COMPETENT on that level.

The strawman macro organizational strategy for a course in logistics decision making therefore is based on two organizing principles. The first is the principle of levels of elaboration where complexity, uncertainty and time constraints rise with each level. The second organizing principle is a division of each level into two successive sections, where the first section is designed to accomplish Transition Nr. 1 in a CAI environment and the second section is designed to accomplish Transition Nr. 2 in a simulation environment. This macro strategy concept is presented graphically in Figure 6.

Figure 6. Macro Instructional Strategy for Decision-Making Training

Level n

Lessons   Practice Exercises

TIME CONSTRAINT
COMPLEXITY
UNCERTAINTY

Level 1
"Epitome Level"

Lessons   Practice Exercises

This defines a high level macro strategy for a course in decision making. The next concern in defining organizational strategy is the micro strategy level. Micro strategy concerns the organization of instruction within lessons and the organization of instruction within exercises.

Micro Strategy within Lessons: Lessons are employed to facilitate the acquisition of content in declarative form. To formulate a micro strategy for use within lessons we begin with Merrill's Component Display Theory (Merrill, 1983), and its three successive "phases" of presentation, practice and performance (i.e., testing). This scheme leaves out instructional elements that should occur prior to presentation, which other theorists, notably Gagne, Keller and Ausubel, consider important.

An alternative "phase" scheme, which includes a pre-instructional phase, an instructional phase and a post-instructional phase, is more comprehensive. The types of instructional micro-elements that are either mandatory or optional components of each phase are shown in Table 2.

Table 2. Micro Strategy Elements in Lessons

Phases and Elements	Function	Theory Support
PRE-INSTRUCTIONAL PHASE	**Prepare student for instruction**	
Attention Grabber	*get the student's attention*	*Keller, Gagne*
Advance Organizer	*provide cognitive scaffold*	*Ausubel*
Recall Stimulator	*activate memory of prerequisite knowledge and skills*	*Gagne, Keller*
Objective	*direct attention to desired outcome*	*Gagne, Keller*
INSTRUCTIONAL PHASE	**Instruct student**	
Generality	*provide the information required to learn and/or control the desired performance*	*Merrill, Gagne*
Help	*facilitate acquisition and/or retention of generality*	*Merrill, Gagne*
Example	*facilitate acquisition and/or retention of generality*	*Merrill, Gagne*
Demonstration	*facilitate acquisition and/or retention of generality*	*Merrill, Gagne*
Practice and Feedback	*elicit the performance and provide feedback*	*Merrill, Gagne*
POST-INSTRUCTIONAL PHASE	**Solidify results of instruction and evaluate success**	
Summarizer	*facilitate retention of generality*	*Reigeluth*
Synthesizer	*facilitate connections to prior learning*	*Reigeluth*
Test	*evaluate the results of instruction*	*Merrill, Gagne*

Micro Strategy within Exercises: Scenario-based exercises afford the student the opportunity to apply the knowledge acquired during the lessons and to practice decision making. Exercises are divided into three phases: an **Orientation Phase**, an **Operations Phase** and a **Debriefing Phase** as shown in Figure 7.

Figure 7. Micro Organizational Strategy within Exercises
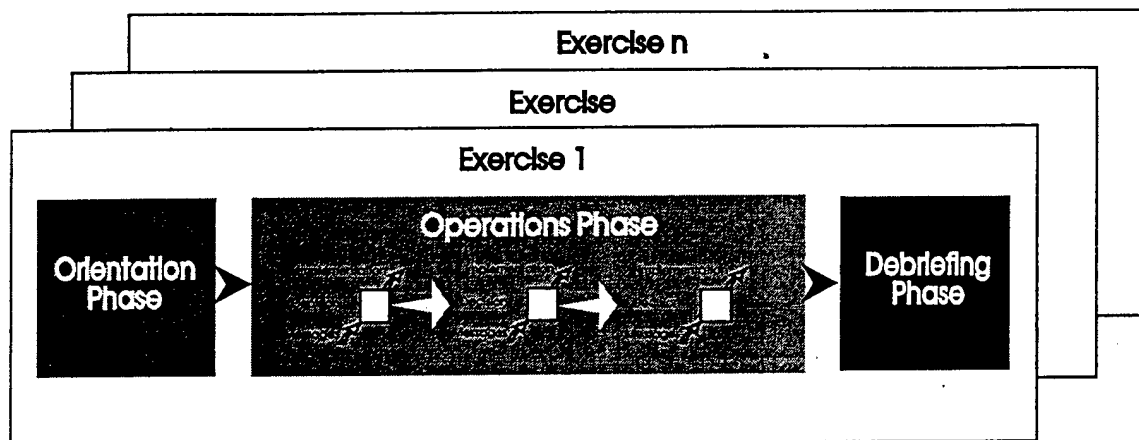
The purpose of the orientation phase is to introduce the student to the basic situation. The tools used for this purpose are dictated by the target domain and include such things as briefings, orders, status boards, message logs, etc. Complexity of these items can be controlled either by means of scripting guidelines or by means of designing different templates for different levels.

During the **Operations Phase** the student receives messages to which they must respond. Some of these messages are designed to trigger decisions; others provide distractions of various kinds. In order to make decisions the student has to reduce secondary uncertainty within the available decision window by consulting information resources that are available to them. Typically these resources include status boards, other members of the team, reference materials, outside agencies, etc. Once the decision is made it can be implemented by whatever means are appropriate to the domain. The environment in which these activities occur can be anywhere from completely quiet to extremely distracting and noisy. The density with which messages arrive can be anywhere from "once in a while" to "hot and heavy" (i.e., the task load can vary over a wide spectrum).

During the Operations Phase, a priori feedback should be available immediately after a decision is made and should provide evaluation and corrective aid with regard to the solution of the decision problem itself and with regard to performance in terms of the underlying model of decision making (if meta content is taught explicitly). Artificial or a priori feedback should be faded during later stages of learning as the student moves increasingly towards a recognition based decision mode. Natural or a posteriori feedback should occur as a consequence of changes introduced into the domain by the implementation of decisions.

Besides feedback, several other types of instructional strategy elements can be introduced into exercises. One type consists of suggested solutions to the decision problem at hand. Another type are prompts of various kinds which direct the student's attention to salient features of the current decision problem and thus provide some degree of assistance. Such assistance can address either domain content concerns, meta-content concerns or both, and may be introduced during early stages of the learning process and faded as the learners performance improves.

21

A particularly interesting type of learning assistance may be provided by using immediate, a priori feedback as a device to improve decisions prior to implementation. In this mode the student formulates a decision and has it evaluated before they commit themself to it. If the feedback from the evaluation indicates problems with the decision (either by a low score on some parameters or by pointing to specific deficiencies), the student can then modify their decision until feedback is satisfactory. This exploratory optimization can provide the student with a unique opportunity to explore and try out specific variations prior to making a commitment to a particular decision. It not only highlights the role of commitment, but also provides instant remediation for gaps and misconceptions in the learner's current knowledge base. If this mode is used, it should be used during early stages and it should be used initially without time constraints. As the learner progresses time constraints can be introduced such that the learner can do all the exploring he wants to do up to the default point. During later stages the mode should be withdrawn entirely.

During the **Debriefing Phase** the learner receives cumulative feedback on an entire exercise from its start to its termination. This feedback should address three issues: (1) a set of "bean counting" scores that indicate to what extent the initial situation has improved or worsened; (2) statistics on how many decision problems were presented, how many decisions were actually made and the quality of these decisions; (3) statistics that indicate how well the student did in terms of decision-making strategy (i.e., did they reduce uncertainty in the right priority, did they access the appropriate information resources, how many decisions were made before and after the default point, etc.).

Levels of Simulation

The practice environment for decision making must be some type of simulation. Each of the five prototypes that were built in the course of the two projects used a different kind of simulation. These simulations can, in retrospect be categorized into three distinct levels shown in Table 3.

Table 3. Levels of Simulation

	Level 1	Level 2	Level 3
Generative	✔	✗	✗
Reactive	✔	✗	✗
Fidelity	high	low	lower
Free Play	✔	limited	✗

The categories on the left side of Table 3 mean the following:

- A generative simulation is one that generates the decision-making problems automatically. The converse of automatic generation is to develop practice problems the hard way: by hand as it were. We call that prescripted problems.
- A reactive simulation provides plausible decision consequences automatically. The converse of that is either prescripted consequences or no reaction at all.
- Fidelity is a term everyone is familiar with.
- A simulation that provides free play is both reactive and generative and provides an evolution of events that depends on the interaction of simulated objects and student interventions. The converse of free play is a prescripted plan of events.

Looking across the three levels of simulation it is easily seen that Level 1 provides something akin to a real life environment and that Level 3 is basically a prescripted linear environment quite like conventional CAI. Level 2 is a compromise between the two.

Content-Strategy Matrix

The content choices together with the levels of practice environment form a matrix that can serve to characterize the spectrum of possible design solutions for training decision-making skills (Table 4).

Table 4. Spectrum of Design Solutions for Decision Training

PRACTICE ENVIRONMENT

CONTENT	Level 1: generative, reactive, high fidelity, free play	Level 2: non-generative, non-reactive, low fidelity, limited free play	Level 3: basic, low fidelity CBT environment w/prescripted problems
Domain specific heuristics	**SuperKEATS**		
Meta heuristics		**DDT**	
Some Combination			**GAIDA**

23

The five prototypes are shown in the cells of this matrix. SuperKEATS was a training package for Fighter Duty Officers in Air Force Tactical Command and Control. This package clearly represented the domain specific approach to training decision makers.

DDT is a desktop trainer for personnel in Logistics Command and Control. This trainer reflects equally distinctly the meta-content approach coupled with a simulation environment that is explicitly fictitious but does have parallels to the current Air Force Logistics system. GAIDA stands for Guided Approach to Instructional Design Advising. It is part of the Air Force's effort to provide automated assistance to instructional design. We designed a sample CAI lesson in Decision Making for this system. In this lesson we used essentially two levels of meta-content: the first level was domain free and the second level was tailored to the domain - which is why GAIDA occupies the "Some Combination" row.

In the following section of the report, each of the five prototype solutions is briefly described and its relative merits and drawbacks are assessed.

Prototype Implementations

Prototype: KEATS

KEATS stands for Knowledge Engineering And Training System. The system addressed the decision training requirements for Fighter Duty Officers in an Air Support Operations Center, an Air Force Command and Control cell that coordinates air support for an Army Corps and is attached to and collocated with an Army Corps Headquarters. KEATS provided automated support for the acquisition of expert knowledge (Figure 8).

Figure 8. Training Tactical Decision Makers: Exercise Flow in KEATS
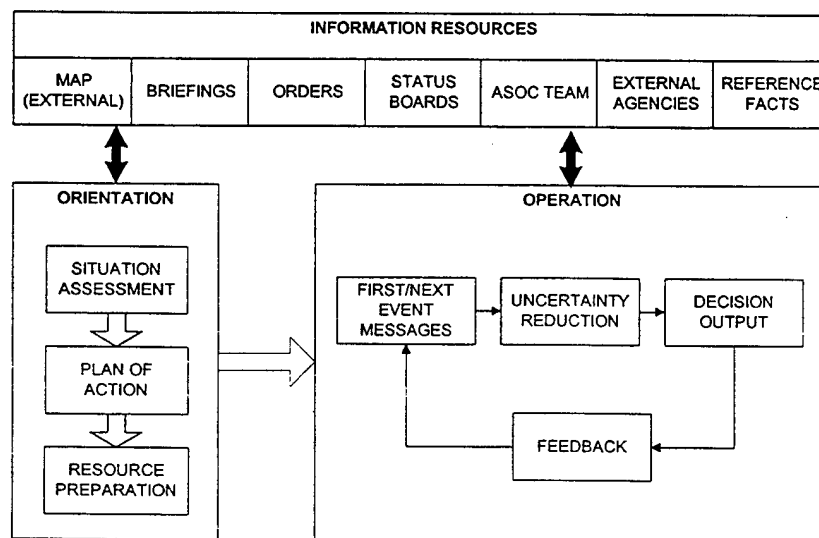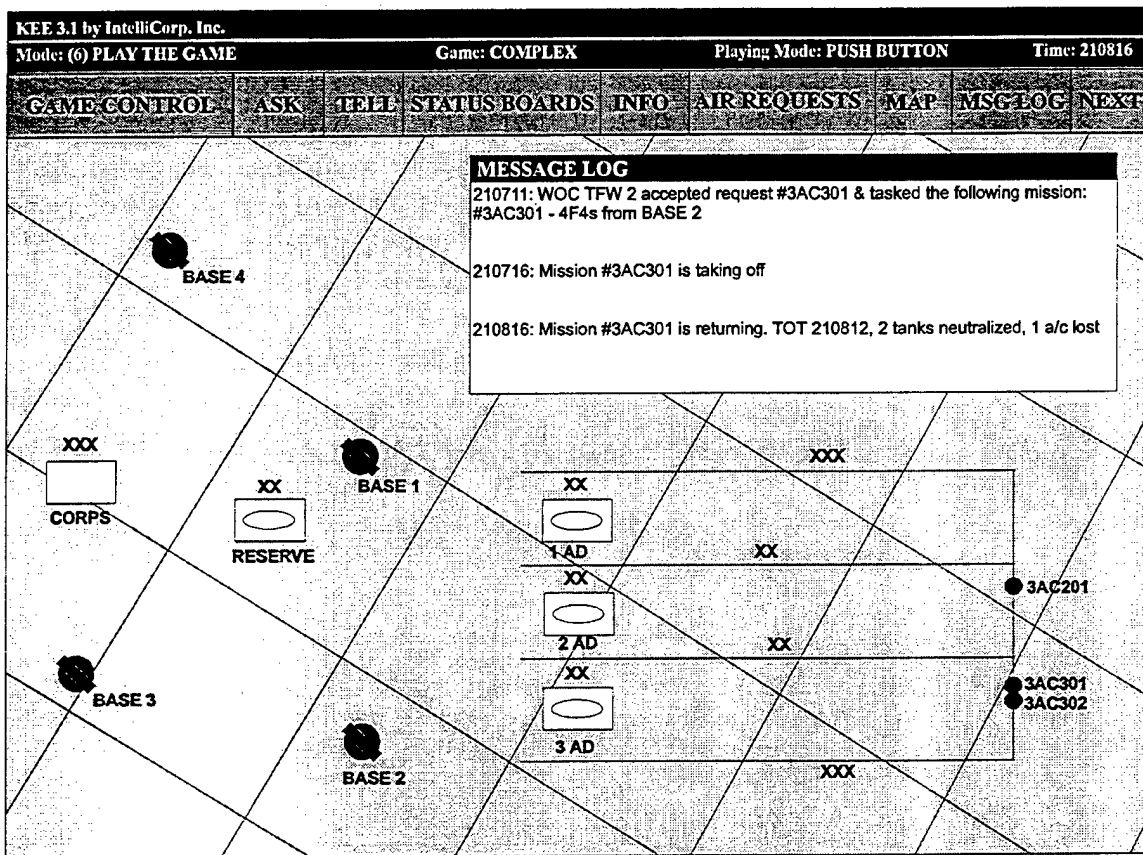

24

Figure 9. Training Tactical Decision Makers: SuperKEATS, Main Screen

KEE 3.1 by IntelliCorp. Inc.
Mode: (6) PLAY THE GAME | Game: COMPLEX | Playing Mode: PUSH BUTTON | Time: 210816

GAME CONTROL | ASK | TELL | STATUS BOARDS | INFO | AIR REQUESTS | MAP | MSG LOG | NEXT

MESSAGE LOG
210711: WOC TFW 2 accepted request #3AC301 & tasked the following mission:
#3AC301 - 4F4s from BASE 2

210716: Mission #3AC301 is taking off

210816: Mission #3AC301 is returning. TOT 210812, 2 tanks neutralized, 1 a/c lost

Prototype: SuperKEATS

Figure 9 above illustrates the main screen of the SuperKEATS prototype. The operational situation is depicted by means of a stylized, simplified map. Decision problems are posed by Army requests for air support on targets that are plotted as red dots. In each case the Fighter Duty Officer has to decide whether they will accept or reject the request. If they accept it they have to decide which of the available fighter resources they will task against the request. In so doing they can query the wings that own the fighters and obtain information from status boards.

The battle waxes and wanes in response to the decisions made. If the decision maker uses their fighter resources judiciously and makes the most of what they have, targets will be increasingly in enemy territory. Conversely, if the decisions do not demonstrate efficient use of resources, targets will begin to wander deeper and deeper into friendly territory. This practice environment was generative, reactive, had high functional fidelity, and allowed free play. Content was domain specific and supplied by a priori feedback. Although it was a computationally expensive solution, experienced Fighter Duty Officers were enthusiastic about it and surmised that the system would provide a high training value.

Prototype: GAIDA Lesson

On the other end of the spectrum is, a lesson in decision making (Figure 10) designed for the GAIDA system (Merrill & Spector, 1991).

The intent of the GAIDA system is to provide advice to instructional designers (ID) by means of instructional design heuristics and examples. Given a type of instructional objective, the ID can consult the system for advice on how to construct each of Gagne's nine events of instruction and can review fully executed examples of these events.

Figure 10. Training Logistics Decision Makers: Sample Screen from GAIDA

What you see on the screen is from the first event: "Gain the student's attention." In the crucial fourth event: "Present the material to be learned" we provide the kind of meta-heuristics and domain adaptations of them that you have seen under content examples.

The student is guided through the solution of a decision problem where they can apply these heuristics in a fairly interactive series of frames; and then another less guided practice example and a test problem are provided. All practice is pre-scripted as is the entire lesson.

This then is an example of the meta approach combined with the least possible effort computer simulation.

Prototype: DDT #1

Somewhere between the GAIDA and the SuperKEATS examples are the DDT prototypes we produced, which are both desktop trainers for decision makers in Air Force logistics. They exemplify the macro strategy shown in Figure 6 by providing levels of elaboration with each level containing both CAI lessons and a practice environment that allows limited free play. DDT #1 (Figure 11) was the first implementation. It featured facilities for the student to provide an assessment of the situation and to formulate a plan of action. The decisions made during the operational phase were evaluated in terms of their compliance with the plan of action and other factors. The user interface of this system was challenging. The situation assessment and plan of action features were considered to be suitable for advanced training. In basic, initial training these features were distracting.

Figure 11. User Interface in DDT #1

Prototype: DDT #2

The lessons provide pure meta-content. The practice environment is explicitly different from the operational environment: it has been placed 300 years into the future, but at the same time it provides the same kinds of problems and the same kinds of tools to solve them as the operational environment. We thus avoid the inherent brittleness of a high fidelity environment while still retaining all the essential features of the decision problems that occur in reality.

DDT #2 (Figure 12) uses prescripted problems, which affords a high level of control over the features of the problems. Automatic problem generation in logistics would be excessively expensive, because it would require simulation of problem causes - and in logistics possible problem causes are legion. We can thus confine the development requirement to a surface simulation of the interaction level, but that is necessarily bought with much labor in scripting problems.

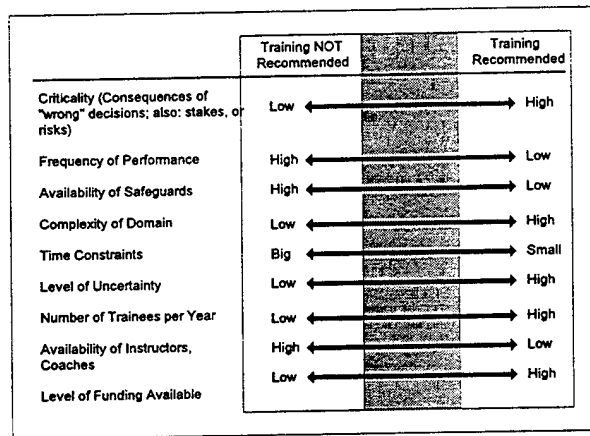Figure 12. Lesson Interface in DDT #2
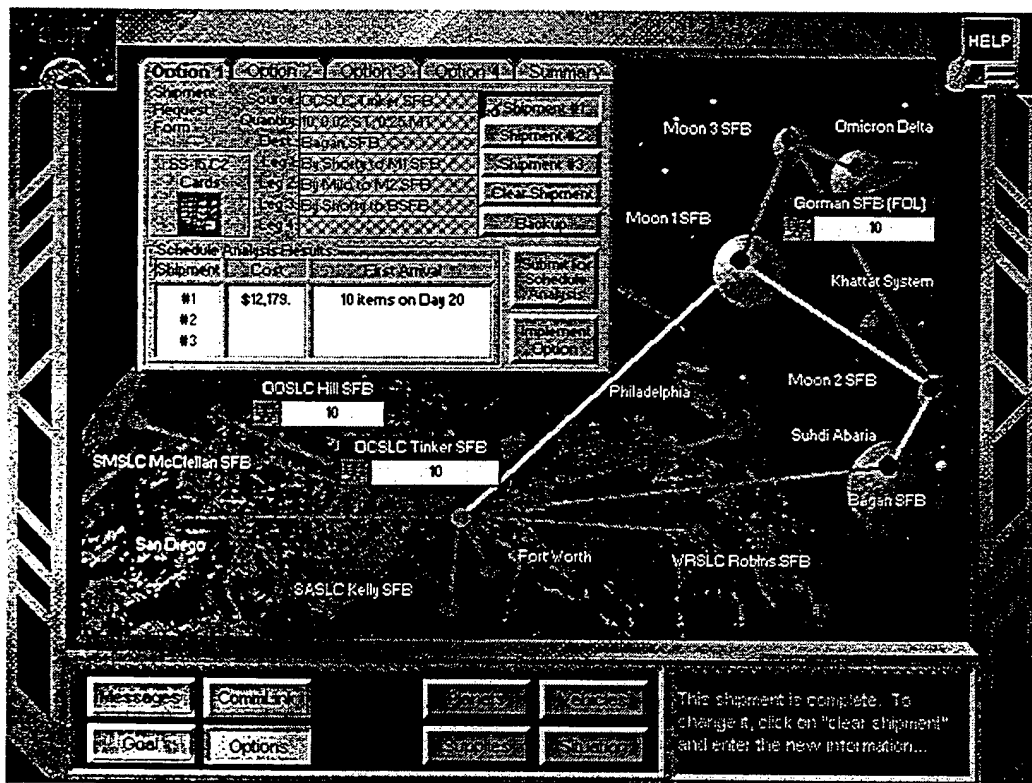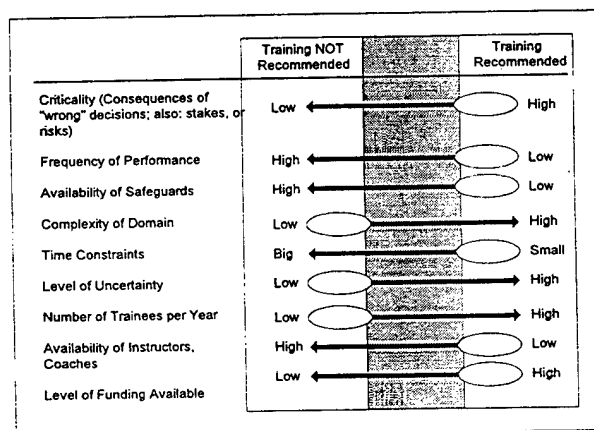
Conclusions

To Train or not to Train

Whether a particular group of people requires training or not depends on the values of a number of variables, some of which are identified in Figure 13.

Figure 13. When to Train Decision-Making Skills

	Training NOT Recommended		Training Recommended
Criticality (Consequences of "wrong" decisions; also: stakes, or risks)	Low	⟵———⟶	High
Frequency of Performance	High	⟵———⟶	Low
Availability of Safeguards	High	⟵———⟶	Low
Complexity of Domain	Low	⟵———⟶	High
Time Constraints	Big	⟵———⟶	Small
Level of Uncertainty	Low	⟵———⟶	High
Number of Trainees per Year	Low	⟵———⟶	High
Availability of Instructors, Coaches	High	⟵———⟶	Low
Level of Funding Available	Low	⟵———⟶	High

The key variable appears to be something we have called "criticality." If the decisions to be made do not have a "high" criticality level, it is generally not even worth considering training. When decisions are "mission-critical" (as they often are in the military) then the remaining variables shown in the figure come into play. If they all weigh in on one side or the other, the training/no training decision is easy. However, if the variables combine into a profile such as the one pictured in Figure 14, the decision is more difficult and might very well require the consideration of additional variables (that are not shown in this line-up) to shift the weight to one or the other side.

Figure 14. A Special Case

	Training NOT Recommended		Training Recommended
Criticality (Consequences of "wrong" decisions; also: stakes, or risks)	Low	⟵———	High
Frequency of Performance	High	⟵———	Low
Availability of Safeguards	High	⟵———	Low
Complexity of Domain	Low	———⟶	High
Time Constraints	Big	———⟶	Small
Level of Uncertainty	Low	———⟶	High
Number of Trainees per Year	Low	———⟶	High
Availability of Instructors, Coaches	High	⟵———	Low
Level of Funding Available	Low	⟵———	High

In the case pictured above critical decisions need to be made rarely but quickly and without safeguards (i.e., nobody is double checking; there is no warning system). However, the domain is relatively simple; the uncertainty level is generally low; there aren't many trainees per year; there are few personnel who could function as coaches but there is a lot of money available for training development. Would one consider the development of some form of automated training in this case? The answer depends as much on the logic of the argument for either side as it depends on such factors as the biases of the people in charge.

The point is that the question of the applicability of automated (or desktop) decision training cannot be answered in any general way. Custom development of such training is invariably a non-trivial expense and therefore must be justifiable in terms of some set of operational benefits. The set of variables illustrated above might be used in the deliberations for any specific case where such training is contemplated.

There are decision tasks that are clearly very important, complex and absolutely buried in uncertainties, but if they occur only one or twice in a lifetime (i.e., deciding on a partner in matrimony), there is a real question whether any task-specific training would or could have beneficial effects (i.e., the probability of divorce) that could counteract long term, pervasive influences in the general environment of the prospective trainee. On the other hand, there are decision-making tasks which occur very frequently, but are so simple and of such trivial importance that any effort spent on providing formal training for them would have to be considered a waste of time.

In between these extremes lie tasks which are the "bread and butter" of decision makers everywhere: mission critical tasks that occur frequently, that represent a qualitatively and quantitatively significant portion of what a given job performer has to do and that are complex enough to require non-trivial learning time and effort. Those tasks are clearly the ones that should receive priority attention of training developers and of managers and it is along these lines of reasoning that one can make intelligent allocations of training resources.

## When is a Desktop Computer the "Right" Medium for Decision Training?

As previously mentioned, each of the two projects addressed a different domain. The first project targeted tactical decision making in an ASOC ; the second addressed decision making tasks in logistical command and control. These domains were of course first and foremost "targets of opportunity" but it is illuminating to assess them in the light of the discussion above. Were these decision-making domains good or bad candidates for the development of automated training?

The tactical decision-making task in an ASOC was an excellent choice for researching the problem, primarily because it was a relatively well bounded task and because it was not overly complex. As a candidate for the development of automated training it was less suitable, because the target population was very small and typically can (and does) learn the task through on-the-job coaching.

Decision making in logistics command and control on the other hand is a totally unmanageable domain unless one carves out some boundable subdomain. The decision tasks in this arena are mission critical but there are usually many safeguards in place, the decision windows permit adequate deliberation, and the coaching mechanism works here as well. This domain was neither a good research target, nor a good training development candidate.

The ideal target for training on a desktop computer is represented in Figure 15. The tasks are clearly critical, they are - at least in peacetime - not performed frequently enough to allow skill acquisition and maintenance on the job, they have to be performed relatively rapidly, and without safeguards against errors. However, care must be taken not to attack a domain that is too complex for cost-effective simulation or one that presents problems where the overall level of uncertainty is so high that any decision becomes just a "wild guess."

Figure 15. The Ideal "Desktop Candidate"
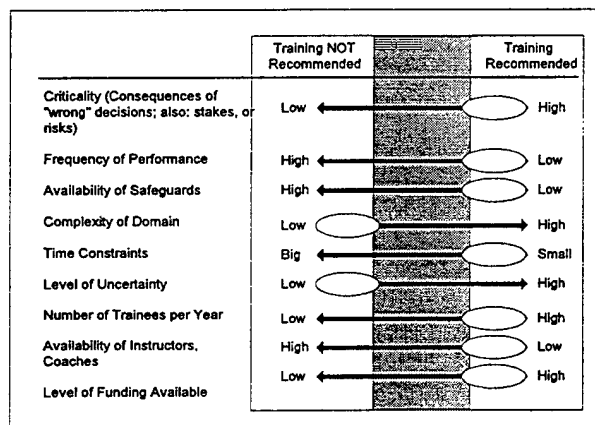
Clearly, the required funds have to be available and their expenditure should be justifiable in terms of decision criticality, size of the target population, and availability (or non-availability) of instructor or coaching personnel.

The Accomplishments

The two research projects have accomplished the following:
1. They have produced the Uncertainty and Timeline Models which contribute to a better understanding of the decision-making task.
2. They have developed a comprehensive training strategy that is well grounded in the existing research base.
3. They have brought the dream of the desktop trainer closer to reality with five different prototype implementations.

The Need for Research

Significant questions remain. For example the issue of emotional engagement during decision-making practice. Decision making is simply a different problem when the stakes are real. One can of course create stress during training, there can even be personal consequences for good and bad decision making, but real war is a very different environment. What can the training designer do to engage the learner as if they were at war? This is one question that requires at least some level of empirical research. Others are the efficacy of meta-content and the question of transfer from an artificial domain. These and more are listed in Table 5 which represent a starter list of empirical research issues that must be addressed in order to optimize the training methodology that has been developed and to provide a basis for generalizability across a variety of decision-making domains.

Table 5. Starter List of Research Issues

Nr.	Issue Title	Type	Research Questions
1	Artificial Domain	Content	To what extent does decision making in a real domain benefit from training in an artificial domain where the same types of problems are solved?
2	Meta Content	Content	Does the inclusion of meta content increase training effectiveness and/or efficiency? What type of meta content is most helpful?
3	Presentation Form	Content Form	Does the form of content presentation affect training effectiveness?
4	Macro Partitions	Organization	What is the optimal number of elaboration levels for given entry and exit skills?
5	Facilitative Elements	Organization	What are the effects of including facilitative elements on student achievement and instructional efficiency?
6	Feedback Forms	Organization	What are the effects of a priori feedback, a posteriori feedback , or combined feedback on student achievement and instructional efficiency?
7	Exploration Practice	Organization	Can forms of exploratory practice increase student achievement and/or instructional efficiency?
8	Learner Control	Management	What are the effects of various levels of learner control on training effectiveness and efficiency?
9	Criterion Levels	Management	Where should one set the criterion levels for student advancement for the optimal compromise between effectiveness and efficiency?
10	Engagement	Management	What measures are available to generate war-like levels of emotional engagement in decision making?

References

Anderson, J. R. (1982). Acquisition of cognitive skill. Psychological Review, 89, 369-406.

Ausubel, D. P. (1968). Educational psychology: A cognitive view. New York: Holt, Rinehart and Winston.

Brecke, F. H. , & Garcia, S. K. (1995) Training Methodology for Logistic Decision Making. AL/HR-TR-1995-0098. Human Resources Directorate, Technical Training Research Division, Brooks AFB, Texas.

Brecke, F. H. (1982). Instructional design for aircrew judgment training. Aviation, Space, and Environmental Medicine, 53 (10): 951-957.

Brecke, F. H., Jacobs, F. , & Krebs, J. (1988). Training of battle staff and commanders assigned to tactical command and control (C^2) systems. AFHRL-TP-87-38. Air Force Human Resource Laboratory, Logistics and Human Factors Division, Wright-Patterson AFB, Ohio.

Brecke, F. H., Hays, P. J., Johnston, D. L., McGarvey, J. M., Peters, S. M., & Slemon, G. K. (1989). KEATS: A system to support knowledge engineering and training for decision making skills. AFHRL-TR-89-25. Air Force Human Resource Laboratory, Logistics and Human Factors Division, Wright-Patterson AFB, Ohio.

Brecke, F. H., Hays, P. J., Johnston, D. L., McGarvey, J. M., Peters, S. M., & Slemon, G. K. (1990). SuperKEATS: A system to support training for decision making skills. AFHRL-TR-90-2. Air Force Human Resource Laboratory, Logistics and Human Factors Division, Wright-Patterson AFB, Ohio.

Brecke, F. H., & Young, M. J. (1990). Training tactical decision making skills: An emerging technology. AFHRL-TR-90-36. Air Force Human Resource Laboratory, Logistics and Human Factors Division, Wright-Patterson AFB, Ohio.

Clancey, W. J. (1987). The knowledge engineer as student: Metacognitive bases for asking good questions. STAN-CS-87-1183, Stanford University, Stanford, California.

Dreyfus, S. E., & Dreyfus, H. L. (1980). A five-stage model of the mental activities involved in directed skill acquisition. ORC 80-2, Operations Research Center, University of California, Berkeley.

Fitts, P. M., & Posner, M. I. (1967). Human performance. Belmont, California: Brooks Cole.

Forbus, K. D., & Gentner, D. (1986). Learning physical domains: Toward a theoretical framework. UIUCDCS-R-86-1247, Department of Computer Science, University of Illinois, Urbana, Illinois.

Frank, H. G. (1969). Kybernetische Grundlagen der Paedagogik. Agis Verlag, Baden-Baden.

Gagne, R. M., & Briggs, L .J. (1974). Principles of instructional design. New York: Holt, Rinehart and Winston.

Gentner, D. (1980). The structure of analogical models in science. TR No. 4451, Cambridge, Massachusetts: Bolt, Beranek, and Newman.

Hammond, K. R. (1986). A theoretically based review of theory and research in judgment and decision making. Report No. 20, Center for Research on Judgment and Policy, Institute of Cognitive Science, University of Colorado.

Howard, R. A. (1968). The foundations of decision analysis. IEEE Transactions on Systems Science and Cybernetics, 4 (3), 211-219.

Keller, J. M. (1983). Motivational Design of Instruction. In C.M. Reigeluth, Instructional design theories and models: An overview of their current status. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Klein, G. A., & Calderwood, R. (1990). Investigations of naturalistic decision making and the recognition-primed decision model. ARI Research Note 90-59, Alexandria, Virginia: United States Army Research Institute for the Behavioral and Social Sciences.

Lawson, J. S. (1987). The art and science of military decision making. Keynote Address, 49th Symposium, Military Operations Research Society, Albuquerque, New Mexico.

Mayer, R. E. (1977). The sequencing of instruction and the concept of assimilation-to-schema. Instructional Science, 6, 369-388.

McDermott, J., & Larkin, J. H. (1978). Representing textbook physics problems. In Proceedings of the Second National Conference, Canadian Society for Computational Studies of Intelligence. Toronto, Canada: University of Toronto.

Merrill, P. F. (1978). Hierarchical and information processing task analysis: A comparison. Journal of Instructional Development, 1 (2), 35-40.

Merrill, M. D. (1983). Component display theory. In C.M. Reigeluth, Instructional design theories and models: An overview of their current status. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Merrill, M. D., & Spector, J. M. (1991). <u>Designing an Advanced Instructional Design Advisor: Transaction Shell Theory (Volume 6 of 6)</u>. Utah State University, Logan, UT.

Nickerson, R. S., & Feehrer, C. E. (1975). <u>Decision making and training: A review of theoretical and empirical studies of decision making and their implications for the training of decision makers</u>. Technical Report: NAVTRAEQUIPCEN 73-C-0128-1. Cambridge, Massachusetts: Bolt, Beranek, and Newman, Inc.

Noble, D., Grosz, C., & Boehm-Davis, D. (1987). <u>Rules, schema and decision making</u>. Office of Naval Research Technical Report R-125-87. Vienna, Virginia: Engineering Research Associates.

Rasmussen, J. (1986). <u>Information Processing and human-machine interaction: An approach to cognitive engineering</u>. Series Volume 12. New York: North Holland.

Rasmussen, J. (1988). A cognitive engineering approach to the modeling of decision making and its organization in: Process control, emergency management, CAD/CAM, office systems and library systems. In William B. Rouse (Ed), <u>Advances in Man-Machine SystemsResearch</u>, 4, 165-243. Greenwich, Connecticut: JAI Press, Inc.

Reigeluth, C. M. (Ed.) (1983). <u>Instructional design theories and models: An overview of their current status</u>. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Reigeluth, C. M. (Ed.) (1987). <u>Instructional design theories in action: Lessons illustrating selected theories and models</u>. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Siegler, R .S. (1978). The origins of scientific reasoning. In R.S. Siegler (Ed.), <u>Children's Thinking: What develops?</u> Hillsdale, New Jersey, Lawrence Erlbaum Associates, 109-149.