TRAINING CRITICAL THINKING FOR THE BATTLEFIELD

VOLUME III: MODELING AND SIMULATION OF BATTLEFIELD CRITICAL THINKING

by

Marvin S. Cohen, Bryan B. Thompson, Leonard Adelman, Terry A. Bresnick, Lokendra Shastri

Cognitive Technologies, Inc.
4200 Lorcom Lane
Arlington, VA 22207
(703) 524-4331

and

Sharon L. Riedel

U.S. Army Research Institute
Field Unit-Leavenworth
Ft. Leavenworth, KS 66027

June 2000

Contract No. DASW01-97-C-0038

TECHNICAL REPORT 00-2

TRAINING CRITICAL THINKING FOR THE BATTLEFIELD

VOLUME III: MODELING AND SIMULATION OF BATTLEFIELD CRITICAL THINKING

by

Marvin S. Cohen, Bryan B. Thompson, Leonard Adelman, Terry A. Bresnick, Lokendra Shastri

Cognitive Technologies, Inc.
4200 Lorcom Lane
Arlington, VA 22207
(703) 524-4331

and

Sharon L. Riedel

U.S. Army Research Institute
Field Unit-Leavenworth
Ft. Leavenworth, KS 66027

June 2000

Contract No. DASW01-97-C-0038

REPORT DOCUMENTATION PAGE			FORM APPROVED OMB NO. 0704-0188
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing the reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1214 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington DC 20503.			
1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE June 2000	3. REPORT TYPE AND DATES COVERED Final, March 1997-September 1999	
4. TITLE AND SUBTITLE Training Critical Thinking for The Battlefield Volume III: Modeling and Simulation of Battlefield Critical Thinking			5. FUNDING NUMBERS Contract No.: DASW01-97-C-0038
6. AUTHOR(S) Marvin S. Cohen, Bryan B. Thompson, Leonard Adelman, Terry A. Bresnick, Lokendra Shastri, and Sharon Riedel			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Cognitive Technologies, Inc. 4200 Lorcom Lane Arlington, VA 22207			8. PERFORMING ORGANIZATION REPORT NUMBER Technical Report 00-2
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Institute Fort Leavenworth, KS 66027			10. SPONSORING/MONITORING AGENCY REPORT NUMBER
11. SUPPLEMENTARY NOTES Dr. Sharon L. Riedel was the Contracting Officer's Representative.			
12A. DISTRIBUTION/AVAILABILITY STATEMENT Distribution authorized to U.S. Government Agencies only; report contains proprietary data produced under SBIR Contract No. DASW01-97-C-0038. Other requests shall be referred to the performing organization in Block 7 of this form.			12B. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 words): The three volumes of this report describe research findings on three closely related fronts: (1) Development of a theory of the cognitive skills that individuals need to make effective decisions in fast-paced and uncertain environments; (2) development and testing of methods for training critical thinking skills on the battlefield; and (3) development of an advanced system architecture to support adaptive instruction and feedback in critical thinking training. Theory development focused on mental models and critical thinking about mental models in a team context, where initiative might be necessary. Training addressed mental models and critical thinking on three major themes: purpose, time, and maneuver. The training was utilized and successfully tested in an advanced tactics course at the Command and General Staff College. Finally, algorithms were developed to simulate both rapid recognitional responding and more reflective reasoning when time is available.			
14. SUBJECT TERMS Command, Training, Critical thinking, Decision making, Mental models, Maneuver warfare, Tactics			15. NUMBER OF PAGES 104 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 None

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI std. Z39-18

ACKNOWLEDGMENTS

We have benefited enormously from the advice and comments of General Leonard Wishart (U.S. Army Ret.), General William Richardson (U.S. Army Ret.), and General Dave Mead (U.S. Army Ret.), who served as consultants on this project. We are grateful to COL Douglas Tystad, COL Soby, and LTC William Hewitt of Command and General Staff College, Fort Leavenworth, KS, for their generous feedback and support. We particularly appreciate the collaboration of LTC William Hadfield in improving and testing the training, and the students in his course for participating in the evaluation. We thank LTC David Lacomb, MAJ Gary Jones, MAJ(P) Jeffery Engbrecht, LTC James Lunsford, and other instructors at CGSC for their feedback and cooperation. Last but not least, we wish to acknowledge Anthony Harrison and Adam Thornton for their prodigious work on the MEntalMOdeler software.

TRAINING CRITICAL THINKING FOR THE BATTLEFIELD

EXECUTIVE SUMMARY

Research Requirement:

Instructors at Army schools and officers in the field agree that current Army education and training do not adequately address decision making skills. What is lacking is a system of training that combines advanced instruction in flexible thought processes (going well beyond doctrinal publications), immediate relevance to Army applications, opportunity for practice in realistic scenarios, and detailed, individualized feedback (not available in current simulators) – and that accomplishes all this despite severe limits of costs, and time and availability of both instructors and students.

The present research had three main objectives:

(1) Develop and extend a theory of the cognitive skills that individuals need to function effectively in fast-paced and uncertain domains.

(2) Develop methods for training those skills in the context of Army battlefield decision making. Improve the ability of Army tactical staff officers to grasp the essential elements of a complex, uncertain, and dynamic situation, visualize those elements in terms of their organization's goals, and take action in a timely and decisive manner.

Test the effectiveness of the training. Does the training improve critical thinking skills? Does it improve the quality of decisions?

(3) Develop a system architecture to support adaptive instruction and feedback in critical thinking training. The architecture should be able to simulate both rapid responses to familiar situations and more reflective responses to novel and uncertain situations.

The training method, like the theory of cognitive skill it is based on, should be readily applicable to a wide spectrum of domains where individuals work in uncertain and dynamic organizational contexts.

Procedure:

Work proceeded on three parallel and closely related tracks: (1) cognitive theory and research, (2) critical thinking training and training evaluation, and (3) advanced modeling and simulation of critical thinking A separate volume of this report addresses the methods and findings of each of these tracks.

In the first track, previous theoretical work was extended in several ways to meet the needs of critical thinking training development: A review and analysis of existing literature on uncertainty handling, additional analysis of interviews with Army staff

officers, and extension of a theory of critical thinking to support algorithm development and to address initiative in teams.

In the second track, we developed and evaluated critical thinking training. We laid the groundwork for training development, by surveying Army training needs and identifying relevant skills for training. We then developed training content and incorporated it into a training delivery system. The training was evaluated in two stages, at Army posts around the country and in a class on advanced tactics at the Army Command and General Staff College, Leavenworth, KS.

In the third track, we developed a computer architecture and algorithms to simulate human critical thinking. These algorithms can serve as the basis for adaptive feedback in future training development.

Findings:

The project introduced innovative statistical methods for discovering the cognitive structure and thinking strategies utilized by decision makers, and employed these methods to analyze several dozen interviews with active-duty Army officers. The Recognition-Metacognition model of critical thinking was extended to address mental models and critical thinking in a team context in which initiative may be required.

A training package was developed with approximately 500 screens. The training addresses three major battlefield thinking themes (purpose, time, and maneuver) and looks at both mental models and critical thinking for each – making a total of six major modules. The training utilizes conceptual instruction, practice in exercises, and historical examples. Graphical interactive techniques were developed to train officers to use both the knowledge structures and decision making strategies characteristic of more experienced decision makers. The training was incorporated into a delivery system that is accessible either through CD-ROM or over the World Wide Web, and is suitable for classroom instruction, training in the field, or distance learning.

The training was tested with active-duty officers in Army posts around the country and at the Command and General Staff College. A very short period of training has been consistently found to significantly affect on both (1) variables related to critical thinking processes and (2) participants' decisions in a military scenario. With respect to critical thinking processes, training increased the frequency with which participants used both proactive tactics and contingency planning, and the frequency with which they referred to the higher-level purposes of the mission. The effect on decisions was dramatic. Participants significantly increased their use of three key tactical elements after training, and also increased their use of combinations of those tactical elements to counterbalance problems with the individual elements.

An advanced computer architecture was designed and partially implemented to support adaptive feedback in critical thinking training. The architecture consists of two interacting components: a reflexive subsystem, which simulates rapid recognition and retrieval of appropriate responses in familiar situations, and a reflective subsystem, which identifies critical uncertainties in the reflexive system and implements strategies for resolving them.

Utilization of Findings:

This project represents an unusually high degree of success both in terms of original research, successful practical application, and commercial potential. The project introduces, develops in detail, and tests a variety of methods for improving decision making skills (i.e., the derivation of training objectives from expert decision processes, a theory of those processes, research techniques for developing training content by modeling expert mental models and decision processes, graphical interactive techniques for conveying this type of content, flexible computer and web-based media, and highly adaptive feedback and guidance. The project addresses immediate Army needs for effective and economical methods for improving the battlefield decision making skills of officers at every level of command, in the schools, in the field, and at home. Its products are already being put to use by instructors in advanced courses at the Command and General Staff College. The training methods have demonstrated enormous commercial potential in a large number of fields, including business, medicine, and aviation. The underlying mental model and decision making technology has even wider potential, for web-based intelligent information retrieval and evaluation.

TRAINING CRITICAL THINKING FOR THE BATTLEFIELD

CONTENTS

CHAPTER 1
INTRODUCTION

The Persistence of Uncertainty

A U.S. military handbook published in 1939 states, "The art of war has no traffic with rules, for the infinitely varying circumstances and conditions of combat never produce exactly the same situation twice." Though perhaps slightly exaggerated, this precept sounds a useful warning, at least in the short and middle terms (and probably much longer), against the persistent dream of achieving "near-perfect knowledge and information of the battlefield" (Ullman & Wade, 1996, p. 9).

Uncertainty in military operations has many causes, not simply the "fog and friction" of combat described by Clausewitz, or deliberate enemy deception, but also novel missions and mission environments, on the one hand, and the unexpected effects of new technology, on the other. Recent military missions have involved operations other than war, joint and multinational regional theaters, and littoral operations. U.S. military personnel have had to navigate between competing and sometimes inconsistent diplomatic, civil, and military objectives in ill-defined missions, and to work within unclear or highly restrictive rules of engagement. "Situation assessment" in such missions means keeping track of blurred and shifting distinctions between friend and foe, guessing the ambiguous intent of armed "bystanders," and ferreting out guerilla fighters in urban or mountainous terrain. In these missions, military personnel have had to overcome communication difficulties and cultural clashes, work with both unstable governments and dissident groups, and to undertake many traditionally non-military tasks, such a police work. Coordinating among own troops, allies, and assisted populations is often more of a challenge than dealing with the "enemy."

Another driver of uncertainty is the expansion of the battlespace through increases in both force dispersal and operational tempo. The last century saw the introduction of motorized, armored, airborne, undersea, unmanned, and space-based platforms. These developments could not have occurred without parallel improvements in sensor and communication technologies. Yet information technology has not fully offset the effects of increasing dispersal and independent action. There is an inescapable tradeoff between amount of information collected and transmitted versus the time it takes for the appropriate human operator to receive it, comprehend it, and react. The unintended consequence has been increasing uncertainty, if not about the enemy, then about the status and even the intent of one's own forces. New high-bandwidth communication technologies (such as the Force XXI Battle Command, Brigade and Below Program) will almost certainly continue this trend, by passing more initiative and decision-making responsibility further down the levels of command.

New technology and new ways of operating have also increased uncertainty in the business world. In the internet economy, the cost of producing an additional copy of an information product is miniscule, and potential customers are overwhelmed by information options. The result is fierce competition for customers' attention, leading to drastic price cutting or free distribution. These investments will pay off in future profits only if a stable base of customers can be created, but such a base is constantly threatened by the possible entry of new competitors and rapidly evolving new technologies.

Technology-based businesses must choose between reliance on open standards to attract a base of customers and to increase the overall size of the market, and development of proprietary products to lock customers in and retain control. Technologies that were intended to increase the accuracy and timeliness of information have shaped a business environment in which uncertainty has increased dramatically.

In the Army as well as business there is a need for training that supports the human's ability to handle uncertainty under time stress. Despite this need, instructors at Army schools and officers in the field agree that current Army education and training do not adequately address decision making skills. What is lacking is a system of training that combines advanced instruction in flexible thought processes (going well beyond doctrinal publications), immediate relevance of the training to Army applications, opportunity for practicing skills in realistic scenarios, and detailed, individualized feedback (which is not available in current simulators). Moreover, all this must be accomplished despite severe limits of costs, and time and availability of both instructors and students.

The present research had three main objectives:

(1) Develop and extend a theory of the cognitive skills that individuals need to function effectively in fast-paced and uncertain domains.

(2) Develop methods for training those skills in the context of Army battlefield decision making. Improve the ability of Army tactical staff officers to grasp the essential elements of a complex, uncertain, and dynamic situation, visualize those elements in terms of their organization's goals, and take action in a timely and decisive manner.

Test the effectiveness of the training. Does the training improve critical thinking skills? Does it improve the quality of decisions?

(3) Develop a system architecture to support adaptive instruction and feedback in critical thinking training. The architecture should be able to simulate both rapid responses to familiar situations and more reflective responses to novel and uncertain situations.

The training method, like the theory of cognitive skill it is based on, should be readily applied in a wide spectrum of domains where individuals work in uncertain and dynamic organizational contexts.

Overview of the Report

This report is divided into three volumes, corresponding to the objectives described in the last section:

Volume I	Basis in Cognitive Theory and Research
Volume II	Critical Thinking Training Training Evaluation
Volume III	Advanced Simulation System for Training

In this introduction, we briefly describe each part of the report. For convenience, the introduction is repeated at the beginning of each volume.

Volume I: Basis in Cognitive Theory and Research

CTI's critical thinking training has several key features:

(1) Unlike many other approaches, it is not based exclusively on formal models of how people ought to think, but on observed differences in decision making strategies between more and less experienced decision makers.

(2) Instruction does not present a set of abstract, disembodied thinking strategies, but trains the targeted skills in a concrete way, embedded within the specific decision making domain.

(3) Training does not simply focus on the individual, but includes an emphasis on decision making within a group context, in which communication is often imperfect or impossible.

In Volume I, we trace the theoretical and research background for the development of such a critical thinking training strategy. Chapter 2 contrasts different views on decision making strategies and strategy selection. Recommendations for handling uncertainty have been dominated until recently by general purpose rules derived from the formal axioms of decision theory. From this point of view, researchers have tended to interpret actual human performance in terms of biases, or systematic deviations from decision theory's formal constraints. In the past 15 years, however, a critical mass of empirical and theoretical work has accumulated that focuses more directly on the knowledge and skill that experienced decision makers apply in real-world tasks, and on strategies that enable them to exploit that knowledge (Cohen, 1993). Chapter 2 traces some of the research threads that have contributed to this development, and which have influenced the present work.

Chapters 3, 4, and 5 describe the way that we have extended that research background in order to build a foundation for the present training. CTI has collected empirical data over several previous research projects that examined decision making in both Army and Navy battlefield environments (Cohen, Adelman, Tolcott, Bresnick, & Marvin, 1993; Cohen, Thompson, Adelman, Bresnick, Tolcott, & Freeman, 1995; Cohen, Freeman, & Wolf, 1996; Cohen, Freeman, & Thompson, 1998). In the Army, we interviewed nearly a hundred officers prior to the present project, occupying a variety of positions and ranks and possessing varying amounts of experience. The present report examines these data from a new point of view, focusing on insights that pertain specifically to *initiative* in a *team* context. This approach was well-suited to an opportunity to develop training for an advanced tactics course at the Army Command and General Staff College entitled *Initiative-based fighting* (developed by LTC Billy Hadfield).

Chapter 3 describes an innovative methodology for identifying knowledge structures, or mental models, from critical incident interview protocols. The methods categorizes judgments or decisions and then analyzes the correlations among the categories across incidents. Mental models are defined as co-occurring categories of information. The influence of other variables, such as level of experience, terrain, and unit type, on the use of these mental models can then be examined. This chapter emphasizes the use of mental models pertaining to *organizational purpose*; the *intent* not just of the enemy but of others in the same organization; *initiative* as an orientation of action to time; and team member *reliability*.

Chapter 4 describes a model of the cognitive strategies that tend to distinguish more effective from less effective officers in battlefield situations (Cohen et al., 1993; Cohen, Freeman, & Thompson, 1998). The model is based on the combination of rapid recognition of familiar situations together with the ability to think critically about the results of recognitional processes. Critical thinking, from this point of view, is not the use of abstract formal rules of thought, but is pragmatic and time-constrained reflection on the uncertainty in the immediate situation and plan. Critical thinking strategies include the identification of qualitatively different types of uncertainty (i.e., incompleteness, conflict, and unreliable assumptions), and the use of different uncertainty handling responses for each. Although the underlying principles of critical thinking are general across domains, the skills themselves are best-acquired in a specific application context, building on previously acquired domain knowledge of the decision makers.

Chapter 5 uses a (newly analyzed) military incident to illustrate how the theory applies to real-world decision making in a team context. The example emphasizes the ability to *think critically* about mental models in situations that require balancing the benefits against the risks of taking initiative. Critical thinking is not just an individual decision making skill. When exercised by a *team leader* and/or *team-members*, it can profoundly alter group dynamics and have important organizational implication.

<u>Volume II: Critical Thinking Training and Training Evaluation</u>

Volume II describes the transition from theory and research to the development of a training strategy (Chapter 6) and training content (Chapter 7), and the incorporation of that content into a computer-based training system (Chapter 8). It then describes the results of two empirical tests of the training system (Chapters 9 and 10).

Chapter 6 reports the results of a survey of Army training needs, and lays out the critical thinking skills to be targeted by the training based on the data, cognitive theory, and student needs survey. It lays out a training strategy based on this analysis, including such methods as instruction, practice, and feedback. Finally, it outlines the theoretical rationale for the training strategy, and contrasts it with training based on other conceptualizations of decision making skill.

Chapter 7 summarizes the training content itself. The training addresses both mental models and critical thinking about three major battlefield themes: purpose, time, and maneuver. It includes six major segments:

(i) mental models to represent the *purposes* of superordinate, subordinate, and coordinate units in an organization

(ii) critical thinking about organizational *purpose*,

(iii) use of action schemas called *time stances* to achieve the proper balance of initiative in achieving those purposes,

(iv) critical thinking about *time stances*,

(v) mental models used in *maneuver warfare*

(vi) critical thinking about *maneuver warfare*.

Chapter 8 describes an integration technology for incorporating the training content within a distributed learning environment. This technology permits distributed sharing of training system resources, interactive exercises, and collaborative, asynchronous learning. The chapter also describes an automated web-capable tutor that we used for testing and evaluation. The system, called *Training to Think Critically on the Battlefield*, can be distributed on compact disc for use on a personal computer or can be accessed over the World Wide Web. It can be used by instructors in the classroom, can be assigned as homework, and can support distance learning and learning in the field. In addition, we developed an authoring tool that permits the construction of new training sequences and interactive exercises, and developed a more advanced prototype system that provides adaptive feedback to trainees regarding critical thinking strategies.

The bottom line question regarding the training is, does it work? Does it improve critical thinking processes as intended, and do such improvements result in enhanced decision making? Training concepts were tested informally with active-duty Army officers at several different Posts, and at a variety of levels of rank and experience, on a continuous basis throughout the development process. Findings from these tests guided training development in an iterative fashion. A more formal test of the training was conducted with over 50 students of an advanced tactics course at the Army's Command and General Staff College. In both cases, training was delivered by computer running software from a CD-ROM.

Interim evaluation results are summarized in Chapter 9. Participants developed courses of action for a combat scenario prior to receiving training, and then revisited the scenario at several points during the training. Exposure to the training helped participants identify and fill information gaps in their plan, expose and evaluate hidden assumptions, and in many cases change their course of action.

Chapter 10 describes experimental tests of the training system with students at the Center of Army Tactics, Army Command and General Staff College. Training was associated with significantly more attention to higher-level purposes (e.g., regarding the larger spatial and temporal context of the unit's own mission), with a greater use of proactive tactics to achieve those higher-level purposes, with a greater ability to identify uncertain assumptions, and with a greater use of contingency plans or branches to handle those assumptions. Training also lead to significant changes in the courses of action that participants adopted. In sum, training influenced both critical thinking processes and the decisions to which they led.

Volume III: Advanced Modeling and Simulation System for Training

Volume III describes the development of an advanced computer architecture to simulate critical thinking performance and to support critical thinking training. The architecture has two interacting components:

(1) a reflexive subsystem, which simulates rapid recognition and retrieval of appropriate responses in familiar situations, and

(2) a reflective subsystem, which identifies critical uncertainties in the reflexive system and implements strategies for resolving them.

Chapter 11 provides an overview of how these two subsystems, working together, can provide the basis for adaptive instruction and feedback in critical thinking training.

The starting point of the reflexive subsystem was a system called *Shruti*, developed by Lokendra Shastri (Shastri & Ajjanagadde, 1993). Shruti combines speed, scalability, and representation of subtle but crucial relational aspects of real-world decision making. To accomplish this, Shruti utilizes rapid, parallel, neural processing, along with temporal synchrony for tracking the identities of objects and roles through relational inferences.

Chapter 12 describes Shruti and extensions of Shruti developed in this project. The extensions were necessary both to improve its representation of reflexive reasoning and to make it work in conjunction with the reflective subsystem. Among the extensions that we worked on were the following:

- integration of utility and belief so that Shruti can simulate decisions as well as inferences;
- mechanisms required for shifting attention, such as temporarily storing and integrating results through a series of attentional shifts; and
- implementation of supervised learning of link strengths through backpropagation.

Chapter 13 describes work performed in this project on a *reflective* subsystem, which critiques the conclusions of reflexive processing and guides its subsequent progress. Features of the reflective subsystem include:

- methods for identifying qualitatively different types of uncertainty based on activation patterns in the reflexive system;
- methods for identifying beliefs most likely to be responsible for different types of uncertainty;
- strategies for shifting attention to beliefs most likely to be responsible for uncertainty.

Uncertainty handling strategies include both domain-specific and more general methods for diagnosing possible causes of the uncertainty and the use of attention and assumptions to stimulate the activation of new information in long-term memory that might resolve the uncertainty.

For convenience, this Introduction is reproduced in all three volumes.

Guide for Readers

Happily, there are alternative paths through this report for readers who have specialized interests, or who wish to get the main points without all the detail. An abbreviated tour through the report that touches on the main areas might consist of the following:

Volume I		
	Chapter 4	Cognitive model of critical thinking that underlies the training design
	Chapter 5	A military decision making example to illustrate the cognitive model
Volume II		
	Chapter 7	Training Content
	Chapter 10	Evaluation of the training at Command and General Staff College
Volume III		
	Chapter 11	Overview of the advanced simulation model for support of adaptive feedback

Another way to break the report down into smaller chunks is by topic or by the reader's primary interest. For example:

Primary Interest	**Most Relevant Sections**
Army training	Chapter 5, to get a flavor of the research basis for the training from a concrete example Volume II
Cognitive Theory	Volume I Chapter 7, for application of the cognitive model to training Chapter 11, for a computational implementation of the cognitive model
Computational models of decision making	Chapter 4, for overview of the cognitive model Volume III

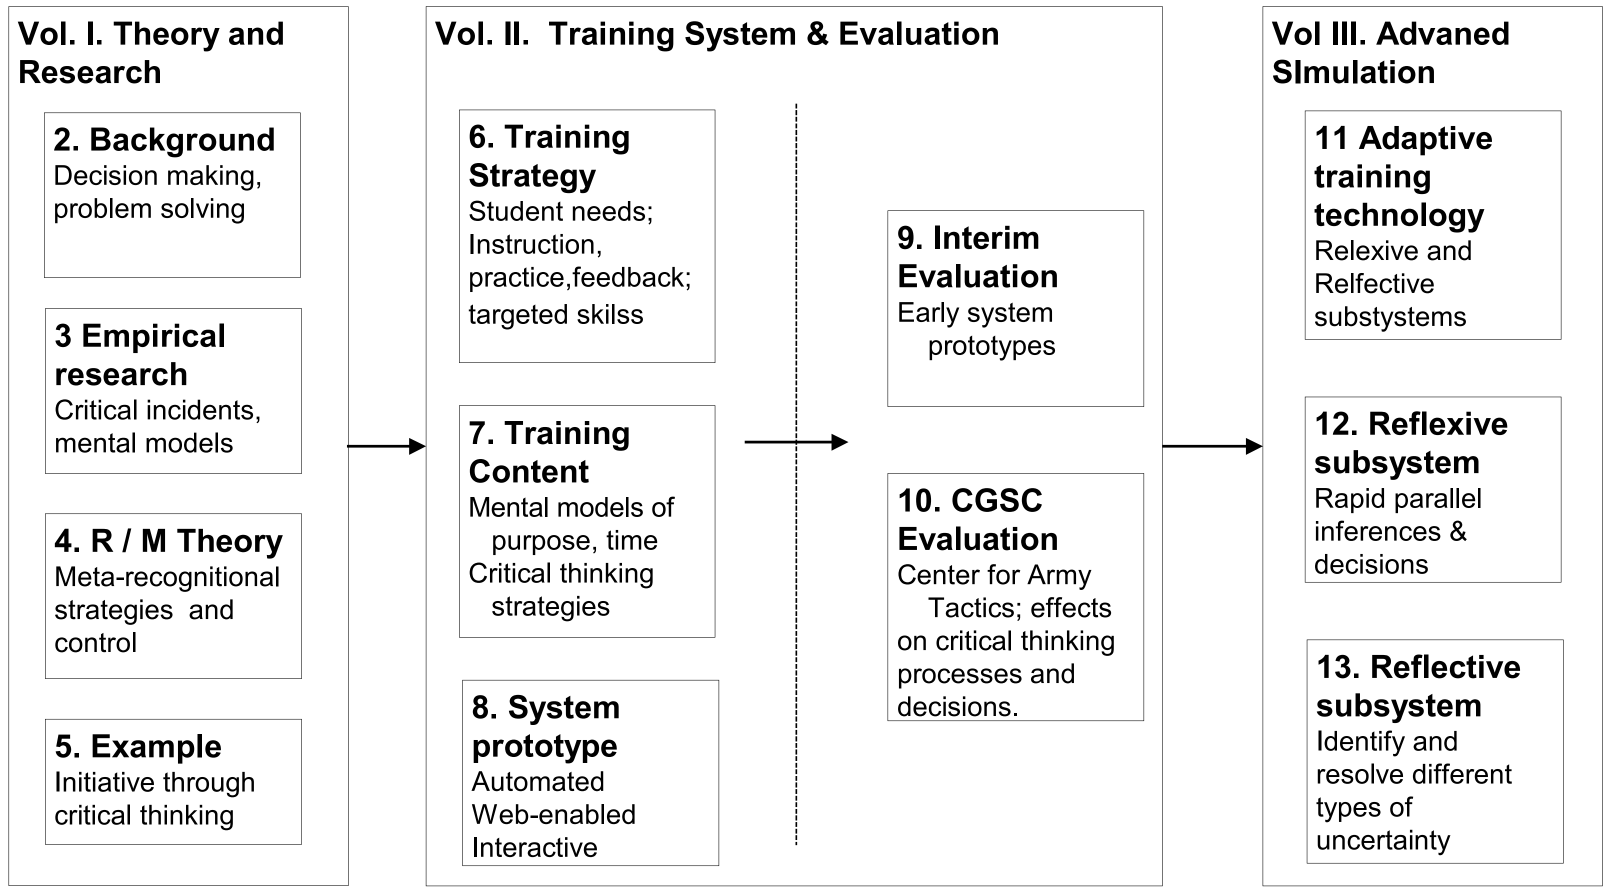

Figure 1. Structure of the report.

CHAPTER 11
SIMULATION TOOL OVERVIEW

Computational Limits and Attention Shifting

The development of training system technology in this project is part of a larger research effort at CTI, in which we are trying to understand and simulate the structure and dynamics of human cognition and decision making processes. In doing this, we are exploring the linkages among several fields and building new models that are informed by data from each of these fields. Our aim has been to integrate high-level cognitive modeling, connectionist knowledge representation, and methods from optimal control (specifically, approximate dynamic programming).

This work represents a synthesis of the Recognition / Metacognition theory described in Part I with work on computer-based inference within the Shruti system, by Professor Lokendra Shastri, of the Computer Science Department, University of California at Berkeley.[1] The immediate goal of this work was to develop a tool that (1) can perform rapid recognitional inferences and planning within a large (expert) belief network, (2) exemplify human limitations on computational resources and attention, and (3) implement metacognitive control process that regulate recognitional processing, help overcome computational limitations, and deal with uncertainty. Such a tool could form the basis, in subsequent research, for the development of an adaptive training system.

According to the Recognition / Metacognition theory (Part I), decision makers structure complex and voluminous knowledge about their world into causal models that enable them to rapidly generate coherent interpretations and plans in response to an influx of new evidence and observations. We model these rapid recognitional processes using Shruti, a connectionist architecture for *reflexive* inference. Critical limits on dynamic access to long term memory (LTM) emerge naturally from the computational structure of Shruti and the neuro-biological constraints it respects. These limits effectively insist that not all information known by the agent can be brought to bear at the same time. One of the key differences between experts and novices is in *how* they structure knowledge to manage these resource limitations and apply the appropriate information during reasoning.

The existence of such limits means that inference and planning processes must be capable of (a) dynamically determining the scope of active human memory from which they draw at any given time, and (b) of remaining coherent within those limits. These changes of scope underlie the fluidity with which a reasoner is able to *focus* limited computational resources at different levels of spatial and temporal abstraction (the chunking problem in AI), and extend planning horizons from moments to years and back to moments. At the same time, this need for fluid changes in focus introduces the necessity for an adaptive dynamics of executive *attention*. The mechanisms of attention

---

[1] Work by CTI in this area has been independently funded by the Office of Naval Research (Contract N00014-95-C-0182 and Contract No. N00014-00-M-0070) and the National Science Foundation (Contract DMI-9861411).

shifting, in turn, form a developmental basis for acquiring skilled *metacognitive* behaviors, which monitor and regulate recognitional processing.

Some of the metacognitive processes that guide the focus of attention within active memory were described in Part I. Studies by the project team and others suggest that these metacognitive processes include (i) monitoring recognitional results for different kinds of uncertainty, including gaps in knowledge, conflicting evidence or goals, and unreliable assumptions; (ii) attempting to fill gaps, resolve conflicts, and evaluate assumptions, e.g., by generating and considering alternative hypotheses to explain evidence and alternative plans to achieve goals; and (iii) regulation of the time taken for reflection versus immediate action, based on the costs of delay, the stakes, and the degree of uncertainty. Our computational research suggests that these metacognitive processes may develop naturally as extensions of skilled attention-shifting behaviors within a resource-limited active memory. Metacognition enables reasoning about highly mediated relationships in long-term memory, i.e., interdependencies that are implicit in long-term memory, but which are too distant to combine to influence decisions. Metacognitive processes thus forge more distant connections and introduce a wider perspective within a computationally constrained reasoning process.

We use Shruti to model both rapid reflexive processes and the reflective processes that monitor and regulate them. In this way, the resource limitations that Shruti implies are shared across reflexive and reflective processing. Therefore, a *reflective* decision maker achieves less in any particular cycle of *reflexive* processing, but may receive a net benefit by extending the span of reflexive processing across multiple cycles of attention shifting.

Application for Training

The great promise of computer-assisted training is its potential to track the progress of individual students in real time. Feedback and training content might be adapted to individual students on at least three different levels: (1) at the lowest frequency, to enduring personal cognitive styles and overall goals, (2) at an intermediate rate of change, to current level of ability, and (3) at the most transient, high-frequency level, to the momentary state of strategy execution, fatigue, attention, or stress. A key technical hurdle at all of these levels is flexibility. Once the knowledge base for a particular problem situation has been coded, the feedback tools and adaptation policies should be able to recognize and evaluate a range of unacceptable and acceptable *variations* in student responses, at various levels of abstraction, and over long and short time periods, without requiring that all variations and their significance be explicitly anticipated by training designers and scripted in advance in the training system. *Flexible*, adaptive training of this kind requires an advanced computer-based model of the targeted decision making skills, going beyond procedural rules for predicting and/or tracking molecular responses. A bonus of this kind of flexibility will be its extensibility to new exercises within the same situation and to new situations in the same domain.

In the rest of this chapter, we provide an overview of some of the distinctive features of the simulation tool, on both the recognitional (reflexive) and metacognitive (reflective) sides, and indicate how they support one another's functioning. We will illustrate how they work in a simple example from the Sanna's Post scenario (described in more detail above in Volume II Chapter 10). We will also briefly outline how these features lend themselves readily to visualization within a graphical user interface. Chapter 12 discusses the reflexive system in more detail, while Chapter 13 provides a more detailed description of the reflective system.

Reflexive and Reflective Processing: An Example

Recall (Volume II, Chapter 10) that in the Sanna's Post scenario, you are the commander of a reinforced rifle company (which we shall call Company A), whose mission is to provide flank security for the other companies in the battalion and to be prepared to become the main effort if necessary. At present, the other companies are heavily engaged to the east, supporting the brigade fight, and your company must prevent reinforcement by the enemy from the west. You receive a scout report that enemy supply vehicles and fuel trucks, as well as some armored vehicles, are in the vicinity of the town of Sanna's Post. There is a road through Sanna's Post that runs eastward to a ford over the Modder River, and from there to the area of the battalion fight (Figure 2). What do you do?

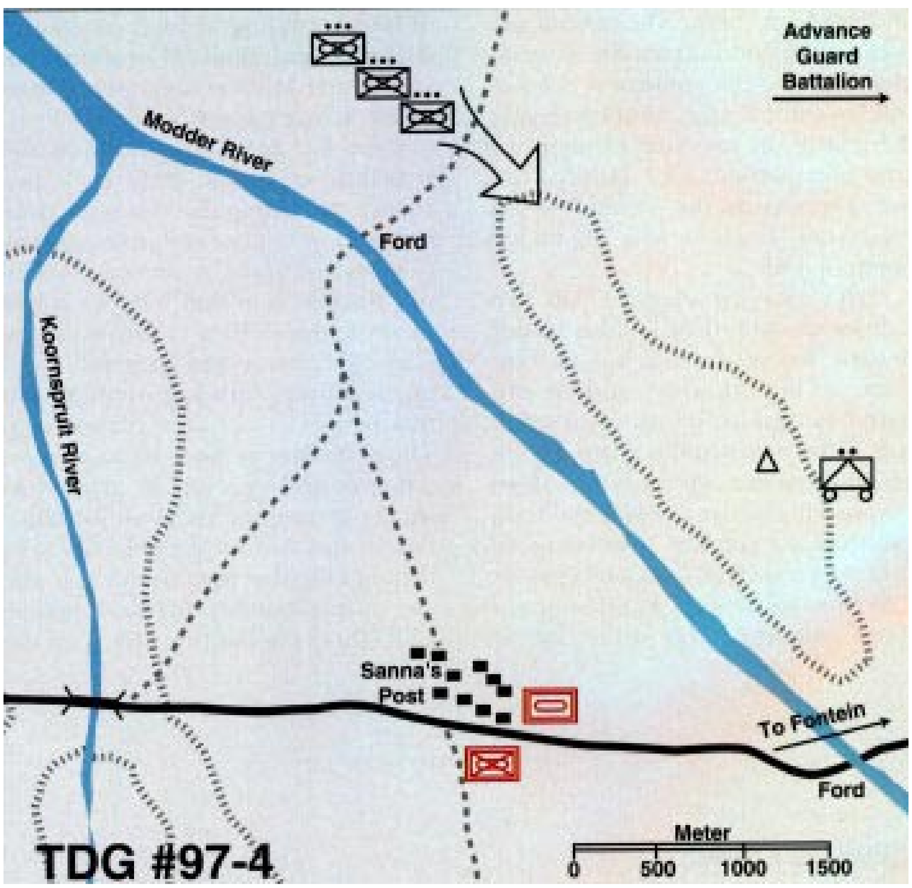

Figure 2. Sanna's Post scenario map.[2]

---

[2] MacIntyre, Capt Douglas J. "Tactical Decision Game #97-4: Battle of Sanna's Post." *Marine Corps Gazette*. April 1997. Quoted with permission by Steve M. Crittenden, Managing Editor, *Marine Corps Gazette*, Box 1775, Quantico, VA 22134, 4 Feb 99.

Figure 3 represents a possible initial recognitional response to this situation. It will help us introduce some important basic features of Shruti's model of recognitional processing. Figure 4 and Figure 5 represent the results of shifting attention under the control of the reflective system, and will help us discuss how reflective, metacognitive processes monitor and regulate recognition. We provide more detail on how these features are *implemented* in Chapters 12 and 13.

Basic Features of the Reflexive System

1. *Shruti enables a connectionist representation of causal mental models.* In Figure 3, the perception of trucks at Sanna's Post activates knowledge in the commander's long-term memory. Some of this knowledge is in the form of an *enemy intent* mental model with four active components: purpose, opportunity, intent, and actions (Volume I, Chapter 3). The long term memory of the commander includes a belief that the enemy's *purpose* is to reinforce the fight against the battalion, and that a likely *opportunity* to do this is along the road through Sanna's Post. For these reasons, the enemy *intends* to use Sanna's Post as a logistics base, and the supply vehicles indicate that the enemy has already taken *action* to implement the intent. Additional *actions* that might be expected based on this intent are use of the road by enemy armored columns heading east.

Weights in the reflexive system can be adapted to the statistical properties of the environment through experience. The Shruti simulator tunes network weights and rule-strengths via supervised learning, using a form of backpropagation. These weights reflect the co-occurrences of concepts that define *mental models* (Chapter 3).

2. *Reasoning proceeds both backward, to find explanations, and forward, to generate predictions.* For example, the perception of trucks at Sanna's Post activates *explanatory* beliefs regarding intent, purpose, and opportunity. These in turn activate a *prediction* of future events, i.e., the appearance of armored vehicles on the road. Combination rules in Shruti capture predictive reasoning and abductive reasoning, as well as taxonomic / semantic reasoning. As a result of the latter, rules framed in terms of general categories can be activated by information about instances of those categories.

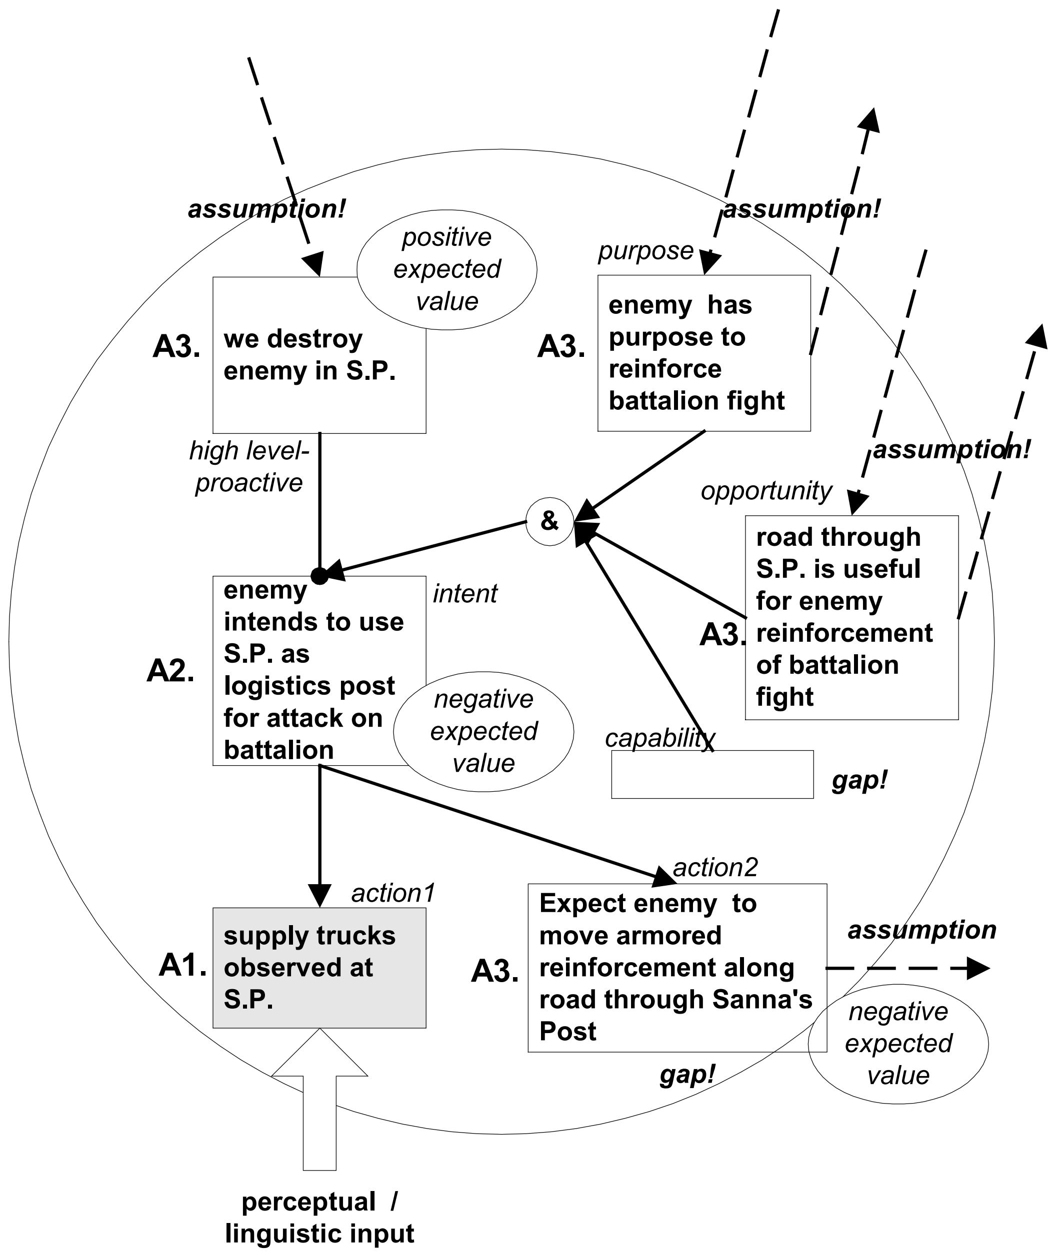

Figure 3. Initial recognitional response to perceptual input in the Sanna's Post example. Attention is focused by perceptual inputs on supply trucks at Sanna's Post (shaded box). The circle represents the spread of activation, and labels A1, A2, and A3 represent the order in which nodes receive activation. Intent mental model components are also labeled (*purpose, intent, opportunity, capability, action*).

3. *The system propagates values as well as belief, and settles on actions at the same time as it settles on a situation interpretation.* Changes in belief lead to the activation of goals, and the activation of goals influences the direction of attention to other beliefs, and ultimately the release of action. As we have seen, in Figure 3, the observation of trucks leads to a prediction (i.e., enemy use of the road to reinforce the battalion fight), which has negative expected utility. The utility of that predicted event

changes the salience of beliefs; in particular, it heightens the degree and persistence of activation of all the causes of the event, including enemy intent, purpose, and opportunity. Thus, *belief* propagated from an event to its causes is heightened by the (positive or negative) utility of the event. In parallel, *utility* itself propagates to causally related events over which the decision maker might have some control. For example, the negative expected utility of enemy reinforcements on the road leads to activation of positive expected utility for any friendly action that can prevent it from being carried out. One such action is destroying the enemy in Sanna's Post. Metacognitive processes, as we shall see, are likely to shift attention to this possible action.

4. *The reflexive system uses dynamic variable binding to keep track of objects and the roles they play in relations.* This feature is virtually unique among rapid, parallel systems. In this example, it enables the system to know that the same entities (i.e., Sanna's Post, the road, the enemy) recur in different parts of the inference, i.e., are bound by different predicate nodes of the inference network. These identities are necessary for the validity of the conclusions. [3] Traditional models of associative processing support "associations of ideas," but do not support the specific, relational inferences that people quickly and accurately arrive at. For example, without object tracking and enforcement of identity constraints by rules, it would be possible to conclude that because trucks were observed at *Sanna's Post*, the enemy intended to use *Fontain* as a logistics post. It might be possible to infer that because this force intended to use Sanna's Post for logistics, that another force intended to use a different local road to move armored vehicles.

5. *The reflexive system uses parallel processing to achieve scalability and speed.* Human recognitional reasoning is extremely rapid over a very large knowledge base, on the order of 500 milliseconds or less. The parallel character of reflexive processing is illustrated in Figure 3. The node labeled A1 is the first to be activated, by perceptual inputs, and A2 is activated next. However, the four nodes labeled A3 are all activated simultaneously, because they are all two layers removed from the initial activation. Because computations at both the reflexive and reflective levels are parallel, time increases as a linear function of the size of the network.[4]

6. *Predictions regarding resource constraints derive naturally from the representational and computational features of the system.* These features account for limits on the amount of long-term memory that can be active in working memory at one time. In Figure 3, for example, activation beginning at A1 spreads to two additional

---

[3] In Shruti, object identity is represented by temporal synchrony of firing at nodes in different places in the network, and identity constraints on the application of rules are enforced by temporal pattern matching.

[4] In addition to versions of Shruti that run on serial platforms, a version of Shruti has also been implemented on a parallel machine, the CM-5, (Mani 1995; Shastri & Mani, 1997). The resulting system can encode knowledge bases with over 500,000 (randomly generated) rules and facts, and yet respond to a range of queries requiring derivations of depth five in under 250 milliseconds. Even queries with derivation depths of eight are answered in well under a second. We are actively exploring the possibility of mapping Shruti onto a more readily available parallel processor, such as a network of workstations (NOW), Beowulf cluster, or array of StrongARM processors. Parallel processor solutions for Shruti have been studied extensively in the context of the CM-5. The results of that research should apply especially well to a StrongARM array, which shares the extremely low message latency of the CM-5.

layers (A2 and A3), but no further.[5] In order to bring more knowledge into play, the commander will have to shift attention.

7. *Shruti acquires and stores aggregated information about what lies beyond the current edge of the active network.* Whenever a node is on the edge of the currently active network (i.e., the A3 nodes in Figure 3), aggregated information stored at that node comes into play, representing the *average historical effects* of currently inactive information that is linked to that node in long-term memory. Instances of aggregated historical information are labeled as *assumptions*, because the validity of inferences within the active part of the network depends on (at least implicitly) *assuming* that the aggregated, historical information in fact fits the *present* situation. Acting on such average information is a big part of what meant by "acting on habit," i.e., failing to mindful of the particulars of the situation. In this model, assumptions of this sort are a principal target of reflective, metacognitive monitoring.

Three kinds of aggregated information, or assumptions, are shown in Figure 3:

(i) *Prior probability of a causal explanation (e.g., purpose & opportunity)*. Two of the A3 nodes in Figure 3 are possible causes of the trucks' presence at Sanna's Post, e.g., enemy purpose and enemy opportunity. The observation of trucks provides some support for beliefs about these causes, specifically, that the enemy's purpose *is* to reinforce the fight and that the road through Sanna's Post *is* a likely opportunity for doing so. To infer the new strength of these beliefs after observing the trucks, the impact of that evidence must be combined with the prior degree of belief in those causes. This prior belief is based on the past frequency with which an enemy of this type had a purpose (or opportunity) of this type. This aggregated information does not take into account specific features of the current situation, which might make it different from the historical average. For example, there may be further causal links that suggest that this will (or will not) be the enemy's purpose, e.g., aspects of enemy doctrine or the historical practice of the enemy commander,. Nodes representing these specific possibilities, even if they do exist in long-term memory, are not currently in "working memory," i.e., they have not been activated in the current network.

(ii) *Expected utility of an event (e.g., moving armored reinforcements along the road through Sanna's Post)*. One of the A3 nodes in Figure 3 is the predicted event, moving reinforcements along the road through Sanna's Post. Much of the negative utility of this event is indirect, inherited from its historical association with further events (to which it causally contributes) that are *more directly* undesirable: i.e., an increase in friendly casualties and a higher chance of the enemy's prevailing. Over past experience, the positive and negative utility from these and other subsequent events has propagated back to the precursor events. As a result, aggregated information about the expected utility of moving armored reinforcements toward the fight is stored at this event node. However, as with

[5] Because Shruti uses temporal synchrony for object identify, finite bandwidth means that (i) only a limited number of objects can be tracked at any given time, and (ii) if jitter increases with the length a signal travels in the long-term memory network, accurate inference about objects (as opposed to coarser associations of ideas) is limited in depth.

prior probabilities, this information does not take into account the specifics of the present situation. For example, it might even be desirable for the enemy to try to move reinforcements on that road in this situation, if expected rains are likely to make the road impassable. This information, if it exists in long term memory at all, is not currently active.

(iii) *Feasibility of an action (attacking Sanna's Post)*. Any event is a potential action or goal, when it is actually within a decision maker's power to bring about or prevent the event. Thus, expected utility propagates from a predicted event back to other events that can causally affect it and over which the decision maker might have some control. Positive or negative expected utility (unlike belief) is propagated to an event only to the degree that action to influence the occurrence of the event is *feasible*. Feasibility information is stored at an event node and consists of aggregated information about the results of trying to accomplish or prevent that event in the past. For example, in Figure 3, the enemy's intent to use Sanna's Post for logistics receives negative expected utility if it is feasible to prevent it, e.g., by destroying the enemy at Sanna's Post. Since it is on the edge of the active network, destroying the enemy at Sanna's Post receives positive expected utility only if the average outcome of attempting to destroy enemy posts of this kind has been success. Once again, this aggregated information does not take the specifics of the situation into account. There may be factors that make destroying the enemy less or more feasible than usual in this particular situation. This information, if its exists in long-term memory, is not active in the current network.

<u>Interaction Between Reflexive and Reflexive Systems</u>

8. *Shruti provides a mechanism for shifting attention, and for the activation of additional information in long-term memory by means of such shifting*. Figure 4 shows the result of shifting attention, under a metacognitive control process, from the observational inputs to one of the four *assumptions* in Figure 3. The nodes labeled B1, B2, and B3, represent spreading activation during the second attentional cycle. In this example, the decision maker focuses on the recognitional response, destroying the enemy in Sanna's Post. This shift brings into view knowledge that was previously dormant. The newly activated knowledge concerns the feasibility of destroying the enemy in Sanna's Post. Attention shifting disaggregates the historical information stored at the node (destroying enemy in Sanna's Post), by exploring the contents of the network beyond that node. In doing this, it allows the pattern of activation to adapt to the facts of the present situation that are represented in the newly activated part of the network. Attention shifting thus removes some of the reliance in decision making on "habit", i.e., historically aggregated information.

It turns out that the facts in this example are not as clear cut as the average. On the one hand, Sanna's Post is a logistics post, which is typically a weakly defended target. On the other hand, armored vehicles have been spotted there, which outgun an infantry company. This represents a conflict of evidence about the feasibility of destroying the enemy at Sanna's Post. Here we have a typical result of metacognitive critiquing and correcting (Volume I, Chapter 4): The solution to one problem (confirming or disconfirming the reliability of an assumption) leads to another problem (conflicting evidence).

Determining when and where to shift attention so as to obtain a more detailed verification of current recognitional conclusions is a key function of metacognitive strategies. In Figure 5, the commander has initiated a third attentional cycle, by shifting attention to one of the sources of the conflict: the presence of armor at Sanna's Post. The nodes labeled C1, C2, and C3 represent spreading activation in the third attentional cycle.

9. *Shruti exhibits priming effects that preserve and integrate the results of reflexive reasoning during successive shifts of attention.* The new information activated by an attention shift must be combined with information that was active before the shift. Once the decision maker has shifted attention to the armor at Sanna's Post in Figure 5, most of the initially active part of the network (in Figure 3) slides beyond the range of activation. In Figure 5, destroying the enemy at Sanna's Post is now on the edge of the active network again, but the consequences of this action (Figure 3) are now inactive. Its expected utility, however, is not simply an average over many occasions of destroying such posts. It reflects details of the present situation that were very recently active in Figure 3 and that are now *primed*. This primed information shows the relevance of Sanna's Post to the battalion fight to the east. As a result, the expected utility of destroying the enemy at Sanna's Post is far higher than a historical average would indicate. Without priming, reflexive reasoning would be a succession of hermetically sealed windows, uninfluenced by context except through crystallized historical averages. With priming, the context of each active window exerts a real-time influence through primed information, even when it is not fully attended.

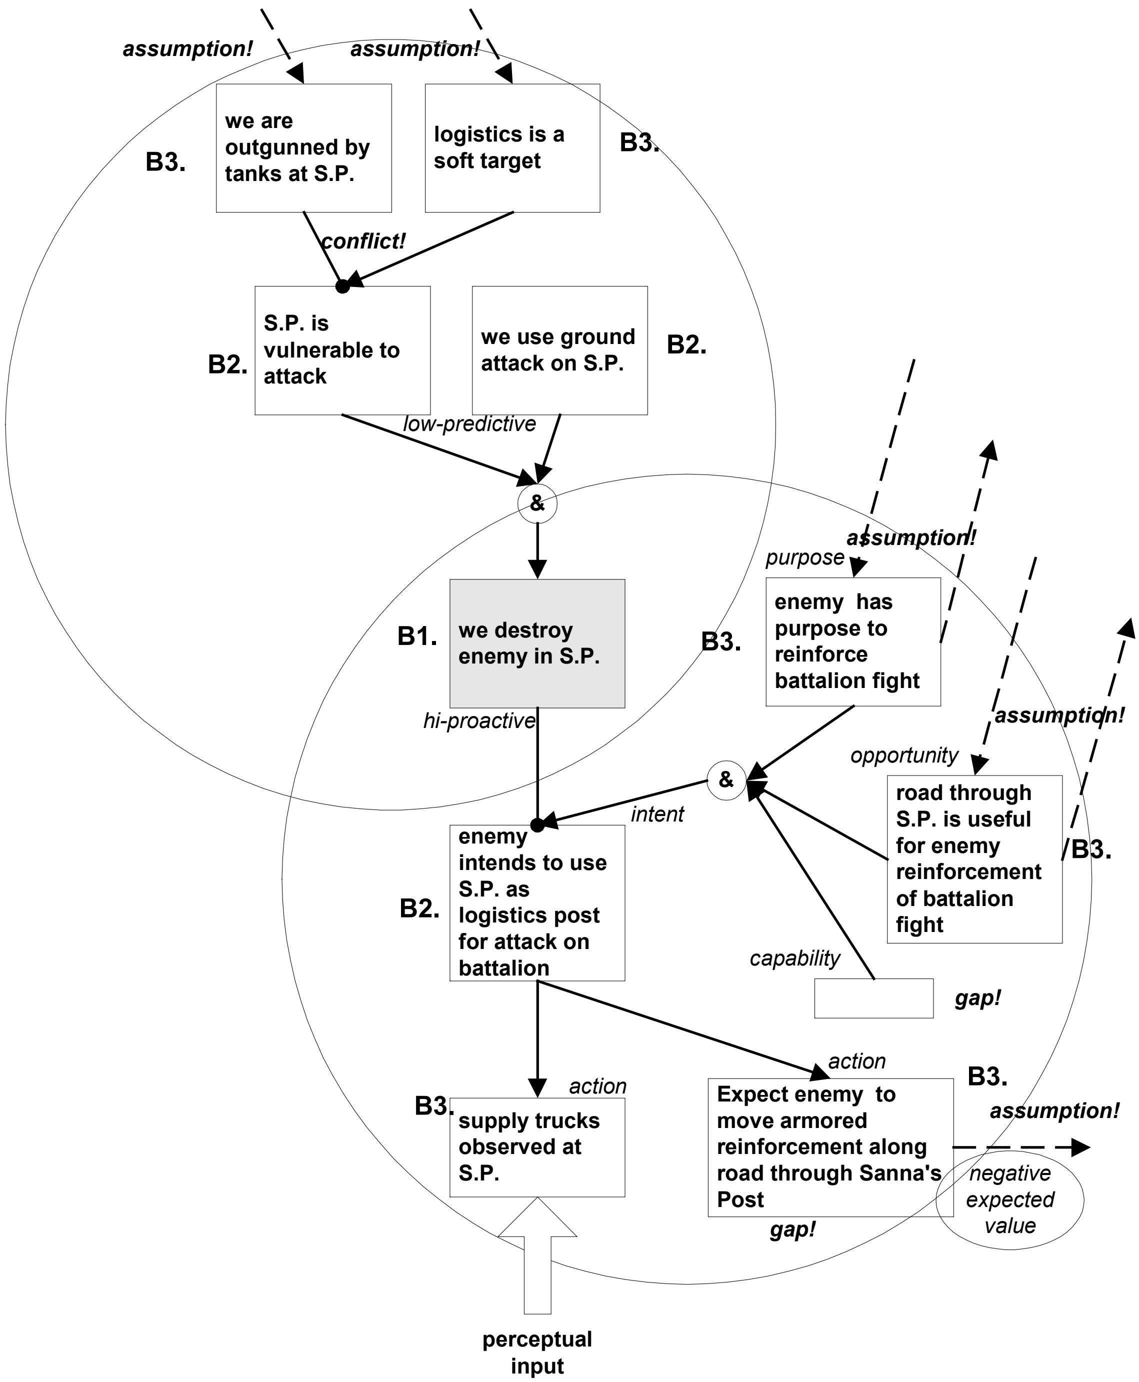

Figure 4. The result of shifting attention to the action of destroying the enemy at Sanna's Post (shaded node). Nodes labeled B1, B2, and B3 reflect the order of activation after the attention shift.

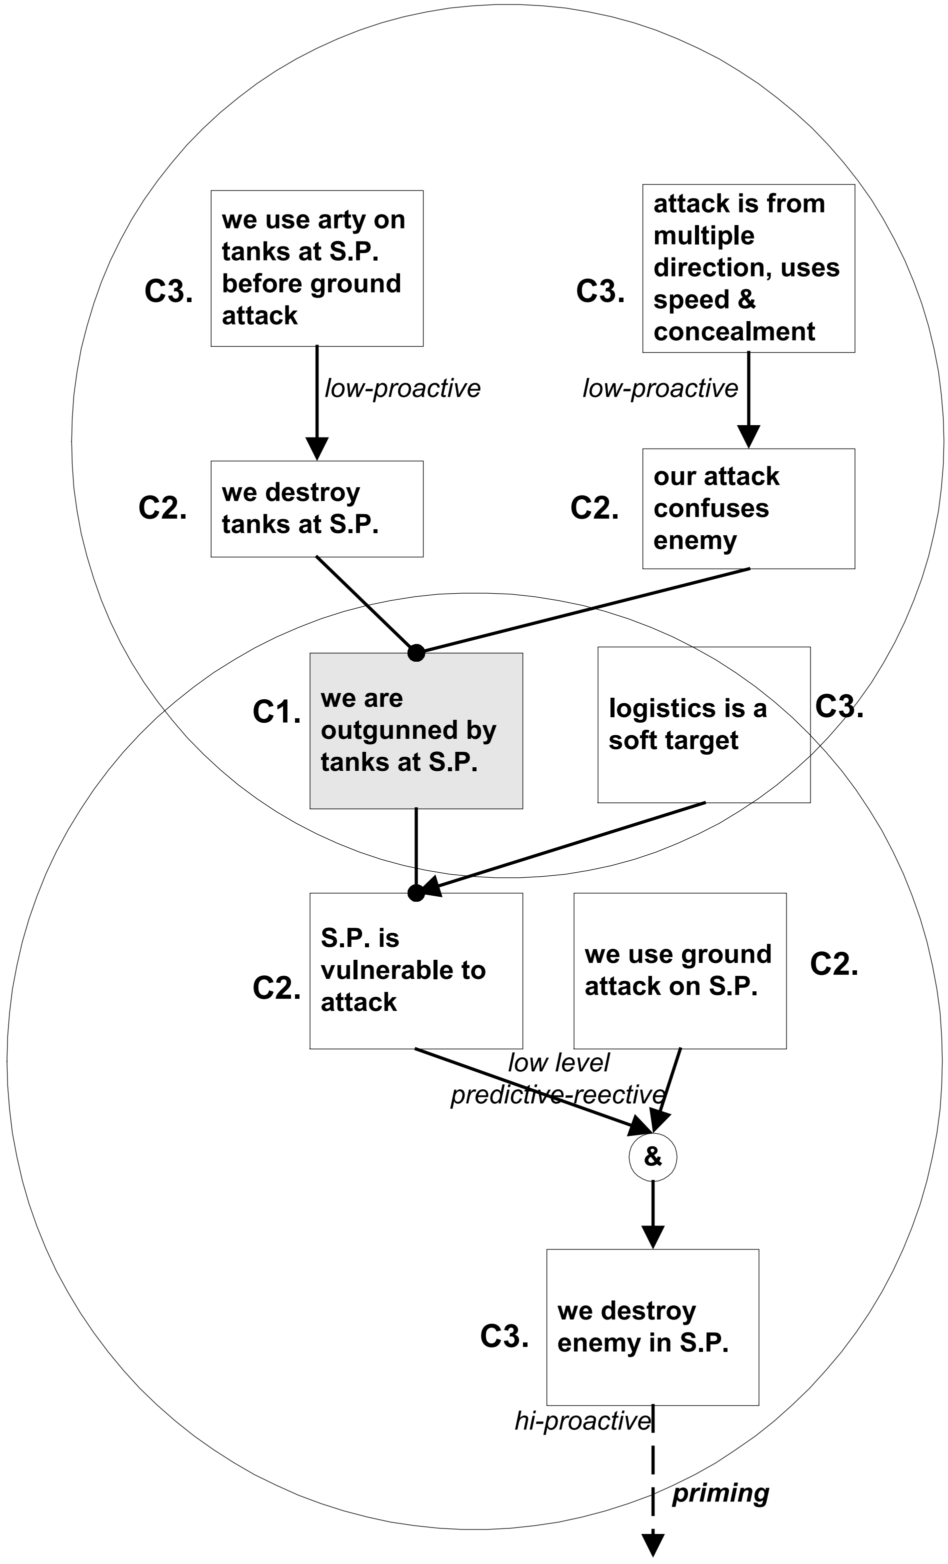

Figure 5. Result of shifting attention again, to belief that we are outgunned by tanks at Sanna's Post (shaded node). Labels C1, C2, and C3 represent ordering of activation of nodes after this attention shift.

10. *Reflective and reflexive attitudes can co-exist although they compete for some of the same resources.* In order to detect uncertainty in a mental model, the decision maker must adopt a *reflective stance*. This means simply that he or she is treating at least one of the components of the currently active reflexive network (e.g., instantiated relations) as an *object*. Shruti constrains the total number of objects that can be discriminated. As a result, the more reflective a decision maker becomes, the less able he or she is to sustain dynamic discrimination of the objects in the domain, such as ships and countries. Because of this constraint on the total number of objects, during routine processing reflection will tend to be limited to the predicate in focal attention. In novel or uncertain situations, reflection can expand to include more parts of the active network and to represent evidence-conclusion relationships among these objects. Thus, the reflective stance (which objectifies statements) competes with, but can also co-exist with, the reflexive stance (which objectifies objects in the domain).[6]

11. *The metacognitive system recognizes qualitatively different patterns of uncertainty in the active network.* Decision makers respond differently when the probability of events is known (e.g., 50% chance of heads and 50% chance of tails in a coin toss) than when the probabilities are totally unknown (e.g., which of two unfamiliar tennis players will win a match; Ellsberg, 1988). Still another situation arises in which strong evidence is brought forward both for a hypothesis and against it (e.g., two well-supported theories which give different predictions for a particular experimental outcome). It seems plausible to suppose that the first case involves a different type of uncertainty (*risk*) than the second case (*ambiguity* or *ignorance*) or the third (*conflicting* evidence).

In real-world decision making, people do not reduce all types of belief and uncertainty to a single measure. Yet traditional probabilistic models have been dominated by single measures (such as probability as a measure of belief, and entropy as a measure of uncertainty[7]). Shruti provides a richer vocabulary, and a capability for more naturally representing a variety of uncertainty-handling strategies. The basis for this flexibility is Shruti's independent registration of evidence for and against a hypothesis and, analogously, the reasons for and against performing an action. This enables the metacognitive system to discriminate four qualitatively distinct uncertainty patterns that may exist at a single node at a given time:

- *conclusion or decision*: significant activation either for or against a hypothesis but not both;
- *incompleteness* of information: little or no activation either for or against a hypothesis (corresponding to ignorance or ambiguity)
- *lack of resolution*: moderate amounts of activation both for and against a hypotheses, in which the sum of + and - activation is one or less (e.g., the coin toss)

---

[6] Another factors that promotes co-existence is that the object discriminations that have been made reflexively are embedded implicitly in the evidence-conclusion relationships at the reflective level. These object discriminations are frozen rather than dynamically tracked, however.

[7] Entropy is, roughly, the degree to which probabilities of each event approximate *1/n* for *n* possible events.

- *conflicting* information: strong activation both for and against a hypothesis, in which the sum of + and - activation is greater than one

In addition, by adding the temporal dimension, we get a fourth pattern:

- *dependence on assumption*: activation that is subject to change as more information is considered (implicit assumption), or as different choices are made (explicit assumption).

Attending to a belief reflectively results in recognition of uncertainty patterns pertaining to that belief. In addition, the decision maker may employ more proactive strategies for uncovering uncertainty. Such strategies involve clamping truth values in order to consider hypothetical situations and plans, i.e., *what if* reasoning. These hypothetical beliefs and actions may activate relevant information that could not otherwise be considered. This information in turn may lead to recognition of additional gaps, conflicts, and/or assumptions.

The three major types of uncertainty (incompleteness, conflict, and assumptions) all appear in the Sanna's Post example.

<u>Incompleteness.</u>

In Figure 3, movement of armored vehicles on the road has been predicted. Nevertheless, this node may represent a gap, since its strength of activation may be quite weak until the movement has actually been observed. Other gaps are represented by other components of the *intent* mental model that have not participated in the arguments for or against enemy intent. In particular, although enemy capability relative to friendly forces has received some activation, no specific information has been retrieved. Such information could have important consequences. For example, if the enemy expected us to be stronger in this sector, they might choose to reinforce the battalion fight by some other route.

A more proactive method for uncovering gaps is to shift attention from evidence (the presence of trucks) to conclusions (e.g., the enemy intends to use Sanna's Post as a logistics base), and clamp the conclusion true – that is, ask, *what if the conclusion is true*? This query results in the reflexive propagation of activation that is equivalent to two questions:

(i) How can the conclusion be explained? The attention shift may draw attention to additional gaps in the mental model representing causes of the intent. For example, the commander has considered the element of opportunity as it affects the enemy (e.g., the road), but not very thoroughly. He has not thought about other possible routes the enemy might use, and whether they offer any advantages, e.g., are they less vulnerable to ambush?.

(ii) What does the conclusion predict? The attention shift may help generate new predictions regarding expected enemy actions that can be used to test the assessment of intent.

Conflict.

In Figure 4, conflicting arguments are activated about whether or not Sanna's Post is vulnerable to an attack. On the one hand, logistics bases are likely to be weakly defended. On the other hand, enemy armor is present at Sanna's Post.

Strategies for filling gaps can always lead to conflict, and so they double as strategies for finding conflict. More specific and more proactive strategies for discovering conflict are also available. To discover potential conflicting evidence, the commander can clamp *the current conclusion as false*, or *the current plan as failing to achieve its objectives.* For example, to find out if there is any evidence for an alternative explanation of the trucks at Sanna's Post, the commander can imagine that he knows that the enemy's intent is *not* to use Sanna's Post as a logistics base. This attention shift is equivalent to asking two questions of the reflexive system:

> (i) How or why is the conclusion could be false? This may activate knowledge of alternative possible explanations for the trucks. For example, the enemy might wish to deceive us in order to fix the company and prevent us from joining the battalion fight ourselves.
>
> (ii) What does the falsity of the conclusion predict? The alternative explanation may lead to other predictions (e.g., regarding overall enemy strength and recent losses), which may be verified by further information collection.

Assumptions.

We have already seen that a particularly important kind of assumption concerns beliefs at the edge of the currently active network: Does the historically aggregated information actually fit the current situation, or will conclusions change as more specific information about the present situation is considered? Figure 3 includes four different assumptions: regarding the prior probabilities of enemy purpose and opportunity, the expected utility of the enemy's moving reinforcements along the road through Sanna's Post, and the feasibility of destroying the enemy in Sanna's Post. In each of these cases, it is possible that past experience was non-representative, or, conversely, that the present situation is unusual.

A simple way to discover assumptions begins by recognizing the heightened *possibility* of dependence on assumptions for beliefs at the edge of the current reasoning process. These are beliefs which have simply been accepted as true, and are not embedded in a web of reasoning or observation. The decision maker can shift attention to such nodes, and determine whether the degree of belief in the event represented by the node changes as a result of newly activated information. Such newly activated information may represent plausible causes of the event, or testable predictions implied by the event.

A context in which assumptions are especially important is in the effort to resolve conflict. Since a proposition and its contradiction cannot both be true, conflict itself is a strong cue that assumptions somewhere in the system of belief are false (Cohen, 1986). The stronger the conflict, the more likely it is that one or more of the beliefs responsible for the conflict must depend on assumptions that are false in the present situation. Such assumptions may be uncovered by shifting attention to beliefs which contribute to the conflict.

A more proactive strategy for ferreting out hidden assumptions is similar to the strategy for finding conflict. This involves clamping a conclusion as false and the evidence for the conclusion as true. In response, the system reflexively searches for alternative explanations of the evidence. If these alternatives turn out to represent gaps, i.e., there is little evidence for or against them, then the original conclusion depended on the implicit *assumption* that these alternative explanations were false.

The metacognitive system increases its effectiveness by learning strategies that are specifically tailored to each type of uncertainty.

12. *The metacognitive system learns to combine a set of simple operations: inhibiting recognitional responding, activation of new information internally by shifting focal attention, and clamping truth values.* These operations are simple and both psychologically and biologically plausible. The metacognitive system *learns* to combine these operations in response to different patterns of uncertainty by reinforcement and associative learning processes. Through such learning, the metacognitive system acquires a rich repertoire of uncertainty handling strategies based both on knowledge acquired about the specific domain, and on more generalizable principles.

13. *Strategies utilized by the metacognitive system contain both domain-specific and general-purpose elements, to varying degrees as a function of experience.* After extensive experience in a domain, decision makers learn which concepts are likely to be the major determinants of different kinds of uncertainty for other concepts. In other words, they learn to identify likely *culprits*, i.e., nodes in the active network that are likely to be responsible for the gap, conflict, or unreliable assumption that has been identified. This might include, for example, knowledge of what the typical gaps in understanding are likely to be in estimating enemy intent, what the likely causes of conflicting conclusions about enemy intent are, and where the hidden assumptions lie. These learned associations guide reflective processing. They bring with them both an increase in efficiency and the risk that novel sources of uncertainty will be overlooked

In filling gaps and in resolving conflict, the metacognitive system utilizes measures of culpability that reflect the sensitivity of belief in one node (e.g., the node where there is a gap or conflict) to changes of belief in other nodes. These are closely related to the parameters that are used to tune weights in the reflexive system to environmental correlations. The result is mastery of domain-specific strategies for metacognitive critiquing and attention-shifting.

More general strategies for reflective reasoning may also be learned, by abstracting from experience or by explicit instruction. These general strategies identify likely causes and cures for different types of uncertainty by using argument relationships. For example, decision makers may learn than when a conflict in beliefs is discovered, they should shift attention to *evidence* for the conflicting *conclusions*.

Our discussion of Figure 3 illustrates a more general strategy. The reflective system recognized that the decision to destroy the enemy at Sanna's Post might well depend on assumptions, since the chain of reasoning that led to that response had not been thought through very deeply. As a result of this uncertainty and the high stakes of the decision, the reflective system inhibited the initial recognitional response. Instead of acting reflexively on that response, the reflective system identified it as the conclusion and shifted attention to it in order to evaluate the assumptions upon which it depended,. The aim was to explore the chain of reasoning more deeply and expose information beyond the edge of the currently active network. Exposure of this information resulted in the discovery in Figure 4 of a conflict between evidence that Sanna's Post will be vulnerable to attack and evidence that it will not.

The grounds of both the optimistic and the pessimistic argument are at the edge of the current network, and thus likely to be dependent on implicit assumptions about the representativeness of historically aggregated information. To resolve the conflict, in Figure 5 the commander shifts attention to the conclusion of one of the two competing arguments (i.e., that we are outgunned by tanks at Sanna's Post). The purpose of this attention shift is to activate assumptions that may turn out not to fit this situation, that is, to expose incorrect generalizations about the strength of this defense. Activation of new information, if it results in revision of assumptions, may eliminate the conflict. If this fails, the decision maker might shift attention to the conclusion of the other competing argument.

Metacognition helps guide the dynamic retrieval and collection of new information, and facilitates learning. It thus helps *create, maintain, and improve the belief network.*

14. *Value tradeoffs contribute to the control of reflective processing.* Recognition-based and reflective responding fall along a spectrum that varies the number of attentional cycles enlisted to arrive at a conclusion or response. Decisions about whether to reflect more (i.e., engage in additional attentional cycles) or act at once on the current best response are determined by the current uncertainty of recognitional conclusions, the costs of delay, and the potential costs of errors.

Concepts for Uncertainty Visualization

A graphical user interface for the Reflexive-Reflective System might have the following features:

The GUI provides a graphical display of the inferential relationships in the *network* of beliefs and actions that represents the situation.

The GUI dynamically displays the current uncertainty status of each active node in working memory, in terms of *qualitatively different types of uncertainty*.

The GUI supports *domain-specific* critical thinking strategies in real time, by providing information about the degree of historical *culpability* of each node in the network for uncertainty at other nodes.

The GUI supports *general-purpose* critical thinking strategies in real time, by providing information about the generic roles that nodes are currently playing within *arguments* that bear on crucial uncertainties in the current situation.

Summary

In sum, there are inherent, and dynamic, limits on the scope of LTM information that can be brought to bear in interpreting any evidence. The key interaction between the reflexive and reflective systems is the adaptive direction of focused attention within the reflexive memory by means of learned metacognitive behaviors. Recency effects are used to assemble such intermediate results into composite assessments. The model suggests that the development of executive attention functions (metacognitive strategies) may be necessary for, and integral to, the development of working memory, or dynamic access to LTM.

Work has been done in the present project to implement or enhance many of the above features, including:

- In the reflexive system: Causal reasoning, utility propagation and decision making, priming, and supervised learning capability.
- In the reflective system: Development of mathematical algorithms for measuring different types of uncertainty and for measuring domain-specific degrees of culpability.
- For the graphical user interface (GUI): Partial implementation of displays for the network of beliefs, different types of uncertainty, and diagnosing problems in arguments.

An Application Programmer Interface (API) for the Shruti simulator was also developed to enable the integration of the Shruti simulator with the Reflective System, and several modifications were made to the Shruti simulator code to facilitate the interaction between Shruti and the Reflective System.

CHAPTER 12
THE REFLEXIVE SYSTEM

Converging Design Rationales for Shruti

Shruti's design is supported by (i) its neural plausibility, and (ii) its success in accurately simulating rapid recognitional behavior. In the system architecture that we will describe, however, Shruti plays another role, as the object of monitoring and control by a higher-order reflective layer. The evidence for this reflective layer also is both behavioral (see Part 1) and neurophysiological. For the monitoring function to be carried out (as described in Chapter 13), the reflexive system embodied in Shruti must posses certain properties, to be described below. The plausibility of the overall model is enhanced if those properties are independently motivated by (a) the requirement to support reflective monitoring and control, and (b) the original design goals of the reflexive system (neural plausibility and accurate simulation of rapid recognitional performance).

The attempt to integrate the reflexive system with a metacognitive component has informed the design of Shruti in a number of ways, and in particular has led to enhancements of Shruti within this project. These include:

- the treatment of causal reasoning,
- integration of information across attentional shifts by means of priming,
- the representation and propagation of utility as a part of planning, and
- the tuning of network weights via learning.

At the same time, the inferential and representational characteristics of Shruti confirm the basic hypotheses underlying the Recognition / Metacognition model, and have supported the detailed computational specification of metacognitive strategies and metacognitive learning (Chapter 13).

Situation Assessment in Shruti

Among the cognitive components in Shruti are (a) *relational instances*, (b) their links to other relational instances via *rules* (c) *facts* of several different kinds which allow rules to become activated, and (d) *objects* whose identities must be mapped onto the appropriate roles in relational instances and tracked through rule-based inferences. Each of these will be discussed in the context of relatively simple examples, intended to show how Shruti can be used to represent Army battlefield knowledge. We will focus first on situation assessment and then move on to mechanisms that are specific to decision making processes.

The Need for Relational Reasoning

Upon learning that an attack by Company A occurred at a particular time, 0400, against a position where an enemy unit had positioned a logistics post, and that the enemy was not expecting an attack at that time and place, we quickly answer questions such as, "Did any friendly company attack an enemy position?" and even, if we have enough other contextual information, "Is Company A likely to be successful in destroying the logistics post?" There are at least two important implications of this ability:

(1) We instantly encode more about this situation than the mere *co-occurrence of features* representing the enemy unit, the logistics post, Company A, 0400, not expecting, and attacking. Our situation awareness includes the *roles* associated with **attacking** (i.e., *event*, *attacker*, *attacked)*, the roles associated with **occurring** (*event*, *time*, *place*), and the roles associated with **not expecting** (*event*, *time*, *place, participant*), and we instantly encode and use information regarding which *entities* fill which of these roles in this particular case. For example, it is Company A that has attacked the enemy at a particular time and place. It is the enemy that did not expect the attack at that time and place.

(2) The long-term knowledge underlying the prediction is not simply a rule that associates **attacking**, **not expecting**, and **occurring** with **destroying**. We also know how the *roles* associated with these relations must map onto each other for the inference to go through. For example, the *attacked* role must be played by the same entity that plays the role of the *agent* that does not expect [the attack]. The *event* that is not expected must be the same type as the attack *event*. The *time* and *place* of the attack must be the same as the *time* and *place* that the *participant* does not expect the *event*. If these consistency constraints are satisfied, then we might conclude that the entity playing the role of *attacker* will destroy the entity playing the role of *attacked*. If we lacked the ability to dynamically bind objects to roles, we could just as easily predict that any unit was likely to destroy any other at any particular time and place, as long as those units, times, and places are co-active in our attention.

Our ability to understand language (Kintsch, 1988; Just & Carpenter, 1977) and recognize situational relationships in real-time suggests that we are capable of performing a wide range of inferences rapidly, spontaneously and without conscious effort, as though they are a *reflex* response of our cognitive apparatus. In view of this, such reasoning may be described as *reflexive* reasoning (Shastri & Ajjanagadde, 1993). This remarkable human ability poses a challenge for cognitive science and computational neuroscience. The inability to enforce consistency among roles in relations has been regarded by some critics as a fatal flaw in current parallel architectures. Workers in this area have been forced to trade off representational power for computational efficiency. Systems that are able represent relational information (e.g., Anderson & Lebiere, 1998) depend on relatively slow, serial processing, while faster, parallel systems have typically been unable to reason quickly about relationships among specific objects.

Shruti demonstrates how a network of neuron-like elements can encode a large body of semantic and causal knowledge, and yet perform a wide range of relational inferences within a few hundred milliseconds (Shastri & Ajjanagadde, 1993; Shastri & Grannes, 1996; Shastri, 1999(a)(b); Shastri & Wendelken, 1998). In Shruti, the encoding of relational knowledge is mediated by networks of focal-clusters representing instantiations of n-ary predicates. Rules correspond to links between focal-clusters, and inference corresponds to a transient propagation of rhythmic activity across focal-clusters. A dynamic binding between a conceptual role and an entity filling that role is expressed within this rhythmic activity by the *synchronous firing* of nodes representing the entity and the roles it fills on that occasion.

Relational Instances

A relational instance is a possible instantiation of an n-ary predicate, that is, a relation (such as **attacks**) that applies to one or more entities.[8] The representation of a relational instance in Shruti is accomplished by a *focal cluster* that has several important component nodes. Figure 1 shows the anatomy of the focal cluster corresponding to the relation **attacks**. (plus a fact that potentially instantiates it). The main components of the cluster are: (a) the *collector* nodes indicated by "+" and "-", (b) the *enabler* node indicated by "?", (c) the *role* nodes corresponding to *event*, *attacker*, *attacked*, and (d) *utility* nodes "+w" and "-w" representing the degree of preference associated with the truth ("+w") and falsity ("-w") of an instantiation of the relation, respectively.

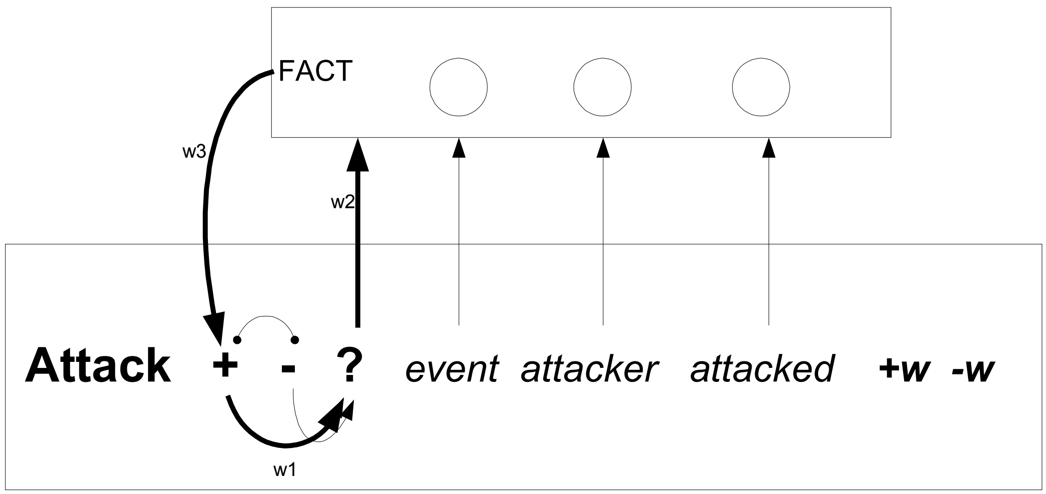

Figure 6. Structure of a cluster representing an illustrative relation. The relation **attack** has three role fillers (*event*, *attacker*, and *attacked*), a positive collector ("+"), a negative collector ("-"), an enabler ("?"), and nodes indicating expected utility for the truth ("+w") or falsity ("-w") of an instantiation of the relation. The *Fact* cluster checks to see if specific objects fill the roles in the relation.

Positive and negative collectors.

The positive ("+") and negative ("-") *collectors* for each relation receive inputs that confirm or disconfirm the relation, respectively. Stable activation of the positive collector ("+") means that the system is affirming that the relation **attacks** is true of a set of specific objects. Stable activation of the negative collector ("-") means that the system is denying the that **attacks** is true of a specific set of objects. Activations into "+" and "-" collectors mutually inhibit one another, so that a large amount of activation in one collector suppresses a small amount of activation in the other collector.

Enabler.

Relational clusters also contain a component called an *enabler*, which we represent as a *query*, or question mark ("?"). Activation of the enabler means that the system is

---

[8] It is a relational *instance* because more than one instance of the same relation (e.g., **attacks**) could be referred to in the same line of reasoning.

seeking an explanation for, or other information about, the currently active instance of **attacks**. Activation automatically flows from the positive ("+") or negative ("-") collectors of a relation to the enabler ("?") for that relation. Shruti thus takes any support for or against a relation as an occasion to search for other information bearing on that relation. The weight parameter w1 (which varies from 0 to 1000) measures the tendency of the system to seek explanations for instances of this relation.

Relational roles.

When objects are recognized or referred to in reasoning, *object clusters* (similar to relational focal clusters) are activated. Each separately identifiable object is assigned a phase in the rhythmic activity that is propagated across the belief network. At the same time, the relational cluster contains a node for each role associated with that relation. When a specific object is recognized either perceptually or through linguistic decoding to fill a specific role, the role node and object cluster are assigned the same phase of rhythmic activity.

Facts about the world.

Facts involved in reflexive inference can be of several kinds: dynamic (i.e., ongoing observation that specific objects satisfy a relation), episodic (recollection that specific objects satisfy a relation), and taxon (statistical information that certain *types* of objects have a certain likelihood of satisfying a relation).[9] The role nodes in a fact (see Figure 6) function as temporal pattern matchers to determine whether there are any active objects that consistently fill the roles in the relation. If so, the fact allows activation to flow from the enabler to the collector of the relation. The weight parameter w2 reflects the likelihood, based on experience, that an explanation will be found along the pathway leading to a dynamic fact, while the parameter w3 reflects the maximum strength of belief that a dynamic fact can provide. In Figure 7, three objects have been identified that fill the three roles in the **attacks** relation.

Getting the (dynamic) facts.

Reports from higher headquarters indicate that *Company A is attacking the enemy logistics post*. Simplifying by necessity, these external inputs to the system simultaneously activate (i) the positive collector of **attacks**, (ii) its three roles (*event*, *attacker*, and *attacked*), and (iii) positive collectors for three entities (the attack itself, Company A, and the enemy logistics post).[10] So far, this activation merely means that the system recognizes (i) that *there is an attack by someone against someone* (but there are likely to be many other recognized relations as well), and (ii) that the three entities are present (but there are likely to be many other recognized entities also). [11] These unconnected recognitions are not sufficient for the specific assertion that *it is Company A that attacks the enemy logistics post*. Shruti accomplishes the correct matching of entities to roles in the relation by a process of temporal pattern discrimination.

---

[9] When we discuss reflexive decision making, we will introduce other kinds of facts.

[10] Although in this case activation occurred through linguistic decoding of a report, it might also have occurred through perceptual recognition.

[11] This is as far as most pattern matching or feature recognition architectures go.

Each individual entity in a particular situation interpretation is assigned a distinct phase of neural firing. Thus, if the same real-world entity simultaneously activates two or more nodes in the belief network, all of these nodes will be firing synchronously. Detection of temporal synchrony is therefore a sufficient basis for equating the object references of two nodes. In this example, the dynamic instantiation of the relation **attacks** requires (i) that the node for Company A and the node for the role of *attacker* fire in synchrony; (ii) that the node for the enemy logistics post and the role of *attacked* fire in synchrony; and (iii) that the node for attack37 and the role of *event* in the **attack** relation fire in synchrony. Detection of these rhythmic correspondences is accomplished by the *Fact* cluster in Figure 7 (where detection of synchronous firing has been represented by circles).

Activation automatically flows from the positive ("+") collectors of the **attack** relation to its enabler ("?"). And since appropriate objects have been found to fill the roles in the **attacks** relation, the flow of activation from the enabler to the positive collector of **attacks** is allowed to go forward, and a stable circuit of activation comes into being: from collector to enabler and, via the fact, back to collector. Subsidiary cycles of activation involving the specific object representations and their respective roles are also set in motion by temporal synchrony. (When an object such as *attack17* is matched to a role in an **attacks** relation, such as *event*, activation is allowed to flow from the enabler ("?") to the collector ("+") of the object cluster.) The totality of this stable activation pattern constitutes the dynamic fact that *Company A attacks the enemy logistics post*, as shown in Figure 7.

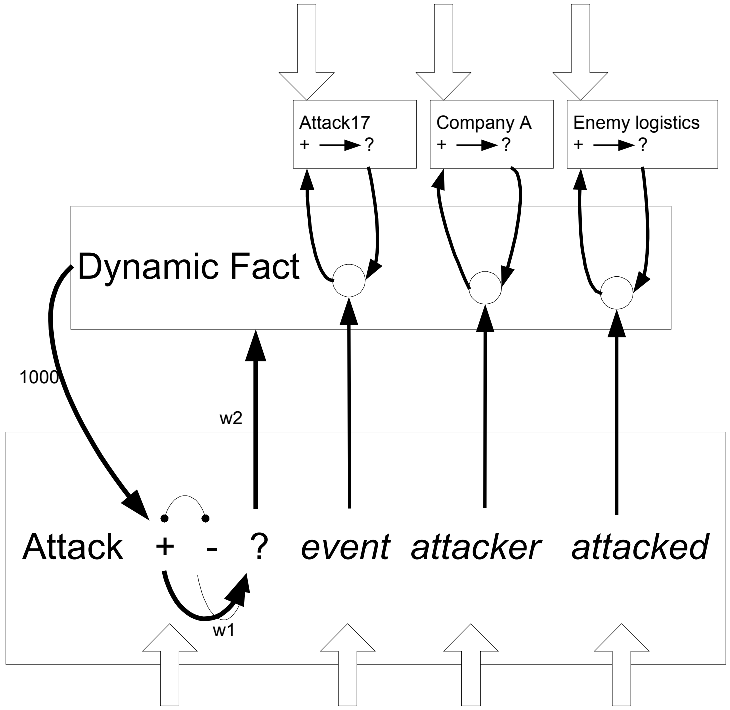

Figure 7. A dynamic fact, representing the instantiation of the attack relation by specific entities. Large empty arrows represent external inputs that initiate activation of specific nodes in the system. Narrow solid arrows represent stable cycles of activation that result from these inputs. Circles represent matching of objects to roles by means of temporal synchrony of neural firing.

Of course, perceptual and/or linguistic inputs will typically activate a great number of dynamic facts. Figure 8 shows three relational instances, two of which (**attacks** and **occurs**) have been recognized as true. The same entity, *attack17*, is involved in both dynamic facts: one fact specifies who the attacker and attacked are and the other fact specifies the time and place of the attack. We do not yet know whether this attack was **unexpected**.

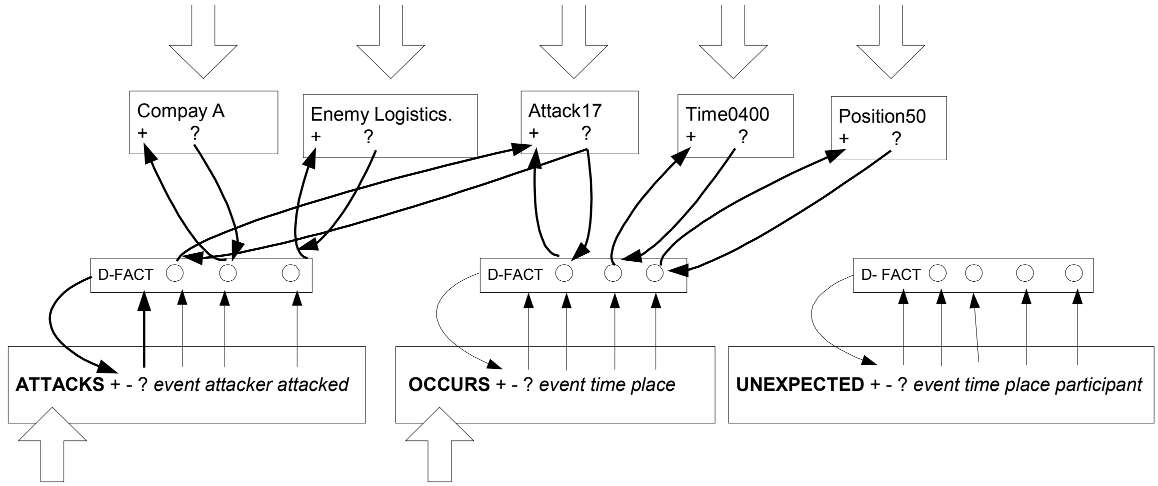

Figure 8. Closure of enabler-to-collector links for two relations (**attacks** and **occurs**) by dynamic facts (D-facts). An additional relevant fact, not shown here or in the following figure, is that EnemyLogistics **is-located-at** Position50.

Rules: Predicting Consequences of the Facts

Inferential connections among relations are encoded in Shruti by means of *rules*. Figure 8 represents part of a rule that might contained in the long-term memory of a military officer. Rules of this sort are be associated with maneuver warfare tactics (see the discussion of the Critical Thinking Training in Chapter 7 and the Appendix of Volume II). In ordinary terms, this rule says: *If an attack occurs at a time and place that is unexpected by the force under attack, the attacking force is likely to defeat the attacked force*.

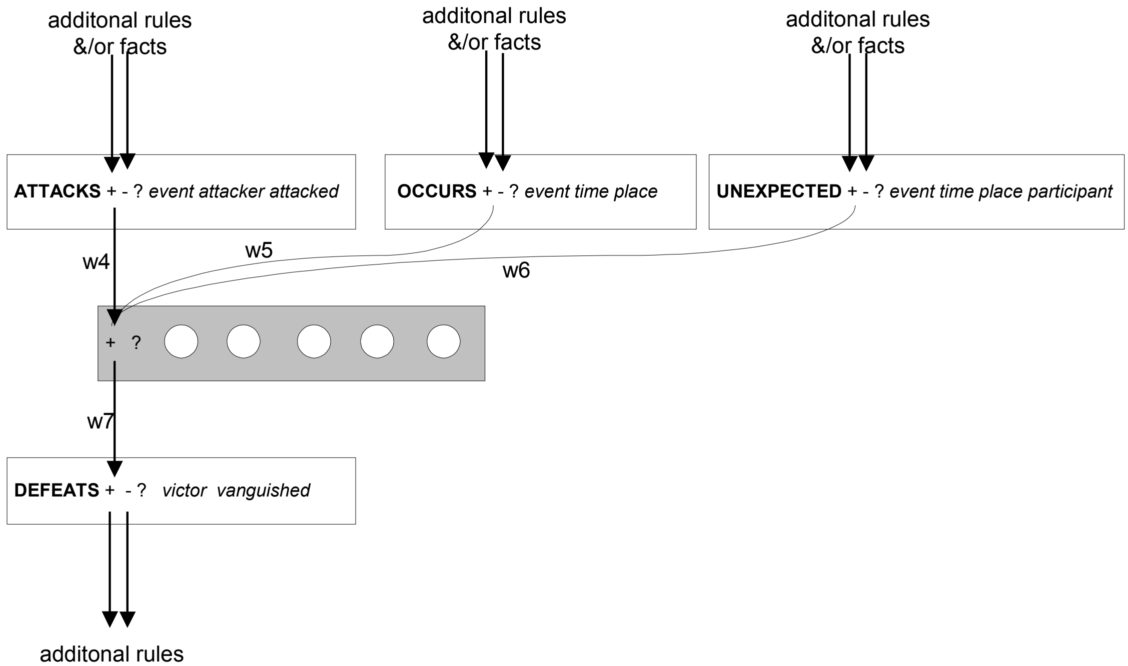

Figure 9. Part of a belief network corresponding to a rule in maneuver warfare, concerning a likely relation between surprise and success. Only collector-to-collector connections are shown. Rule mediator cluster is shaded.

For a rule to become active in Shruti, and for a conclusion to be established, a stable cycle of activation must flow through the rule. In the case of a predictive inference (as in Figure 9), activation will flow from facts associated with the antecedents to the consequent, and back. In the case of an abductive inference, which we will illustrate later, activation will flow from facts associated with consequents through statistical facts regarding antecedents to specific conclusions regarding antecedents. A stable, self-reinforcing activation cycle of this kind across an interlinked set of relations constitutes an *interpretation* and/or *plan* for the current situation. We will review each kind of link within a rule that is required to support such cycles of activation.

Collector-to-collector links.

In the rule shown in Figure 9, the positive collectors of **attacks**, **occurs**, and **unexpected** are linked first to the positive collector of a *rule mediator* cluster. The strength of connection for each antecedent (weights w4, w5, and w6 in Figure 9) measures the importance of each antecedent to the rule. Inputs from the two dynamic facts (**attacks** and **occurs**) are combined at the positive collector of the rule mediator cluster by a *soft-and* function that does not require complete activation of all three rule antecedents. Activation then flows from the rule mediator node to the positive collector of the consequent, **defeats**. The weight for this link (w7) is a measure of the support offered by the rule (with all antecedents satisfied) to the consequent, i.e., the chance that a surprise attack will cause a defeat.

Role to role links.

In parallel with collector-to-collector activation, activation flows along the connections between *roles* of the antecedent and consequent relations (Figure 10). For activation to flow from antecedents to consequents and back, the objects that fill relational roles must satisfy the pattern of identities demanded by the rule.[12] These identity constraints are enforced by temporal synchrony-matching at the rule mediator node. The required identities are represented in Figure 10 by the convergence of the arrows from antecedent roles to rule mediator roles. In our example, with two active antecedents, there is only one identity constraint: The *event* that satisfies the **attack** relation and the event whose place and time are described by **occurs**, must be the same event. (If and when **unexpected** is activated, the participant in the **unexpected** relation would have to be the attacked in the **attacks** relation; and the time and place of that same attack would have to be the same time and place at which the attack is **unexpected** by the participant.)[13]

The rhythm of firing in object nodes is matched to roles in antecedent relations that are filled by those objects. The rhythm of firing in the antecedent roles is in turn transmitted through the rule mediator node to the roles that are filled by the same entities in the consequent. In our example, the *victor* role of **defeats** fires in synchrony with the object cluster for Company A; and the *vanquished* role of **defeats** fires in synchrony with the object cluster for enemy logistics post.

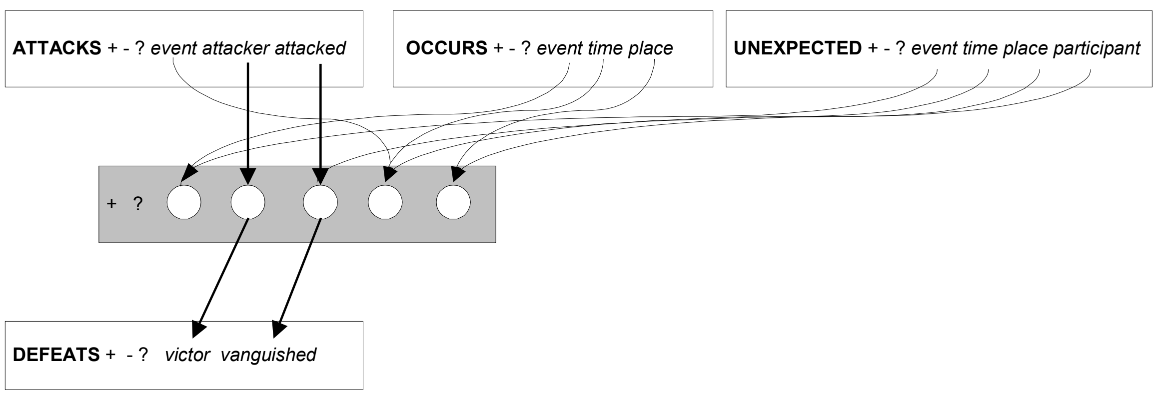

Figure 10. Links from roles that must be filled by the same object converge on the same temporal matching site (represented by a small circle) at the role mediator node. This node verifies that the pattern of role-filler identities required by the rule is satisfied.

---

[12] In other words, from a logical point of view, many rules underlying rapid reflexive reasoning cannot be expressed as *If…then* relations among *propositions*. They must be expressed using resources from the predicate calculus, i.e., as quantified statements containing bound variables. Most attempts to model recognitional performance fall short in this regard. They are unable to account for how people keep track of object identities across relations.

[13] Links from the *event* role in **attacks**, the *event* role in **occurs**, and the *event* role in **unexpected** converge on the same temporal synchrony site in the rule mediator cluster. Similarly, links from the *attacked* role in **attacks** and links from the *participant* role in **unexpected** converge on the same site. Activation from antecedent relations is not allowed to flow to the consequent unless the relevant role-filler links are firing in synchrony.

Collector-to-enabler and enabler-to-enabler connections.

As a result of the collector-to-enabler link within a relation, and the enabler-to-enabler links in a rule, Shruti immediately converts the *affirmation* or *negation* of any claim (activation of "+" and "-" collectors) into a *query* about that claim (activation of enablers ("?"). The query is sent to other relations that might support or negate the claim via rules. If this activation backchains (via rules) to a relation that is associated with a *fact* of some kind, that fact serves as a bridge for the activation to flow from the enabler to the "+" or "-" collector of that relation. This closes the circuit and the activation flows back from the discovered relation, along collector-to-collector links of the rule to the original claim, supporting or negating it.

From the positive collector of **defeats** activation automatically flows to the enabler of **defeats**. The system thus seeks to prove or disprove the relevant instantiation of the **defeats** relation. This process queries any facts directly connected to **defeats**,[14] and queries other relations linked to **defeats** via rules. As shown in Figure 11, the enabler of **defeats** queries the rule mediator node. The weight w8 reflects the likelihood that the explanation of the consequent (i.e., one force **defeats** another) will be found through *this* rule as opposed to other rules that represent other possible explanations for **defeats**. Then the rule mediator node queries each of the rule antecedents, **attacks**, **occurs**, and **unexpected**, with weights w9, w10, and w11, representing the experienced importance of each antecedent for such an explanation.

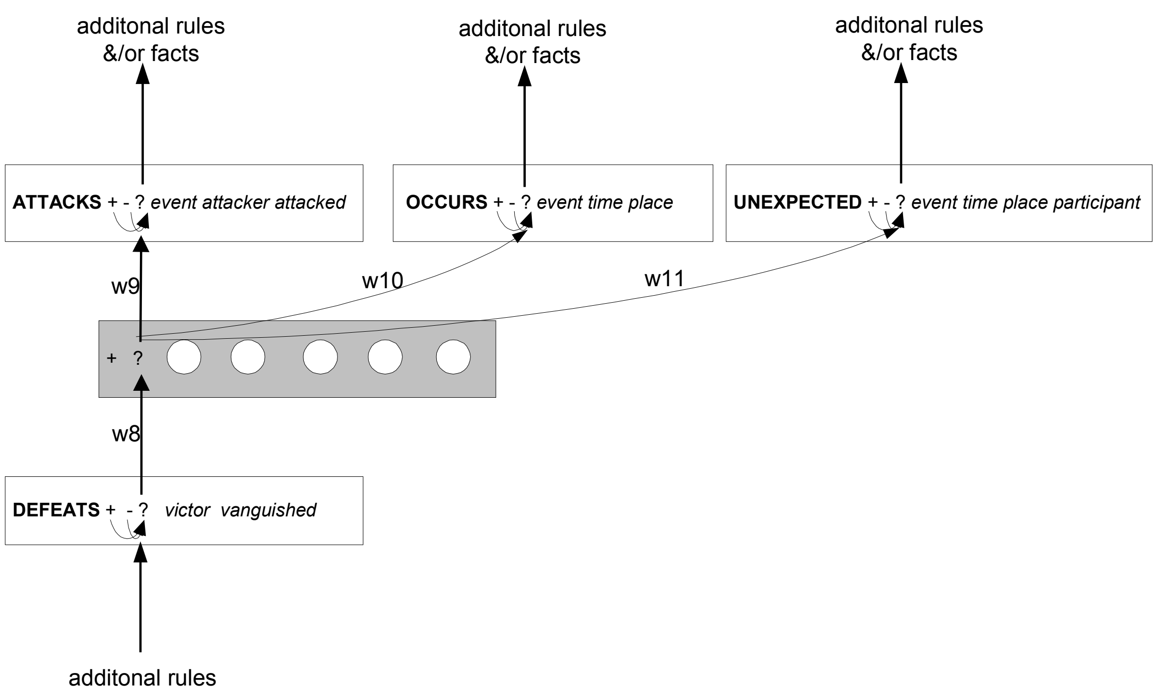

Figure 11. Collector-to-enabler and enabler-to-enabler connections

[14] As we shall see, each instantiated fact can be associated with a prior probability even in the absence of direct evidence.

Rules Interacting with Facts

Activation can be sustained, and conclusions arrived at, only if there is a completed cycle of activation.[15] For such a return loop to be created, enabler-to-collector links must be closed by means instantiated *facts* in the chain of reasoning.[16] In our example, we have knowledge of facts regarding **attacks** and **occurs**. Figure 12 shows how instantiation of two of the relations in the antecedent closes the enabler to collector links for those relations. Figure 12 summarizes all the links described so far, and shows how they create a stable, self-reinforcing cycle of activation.

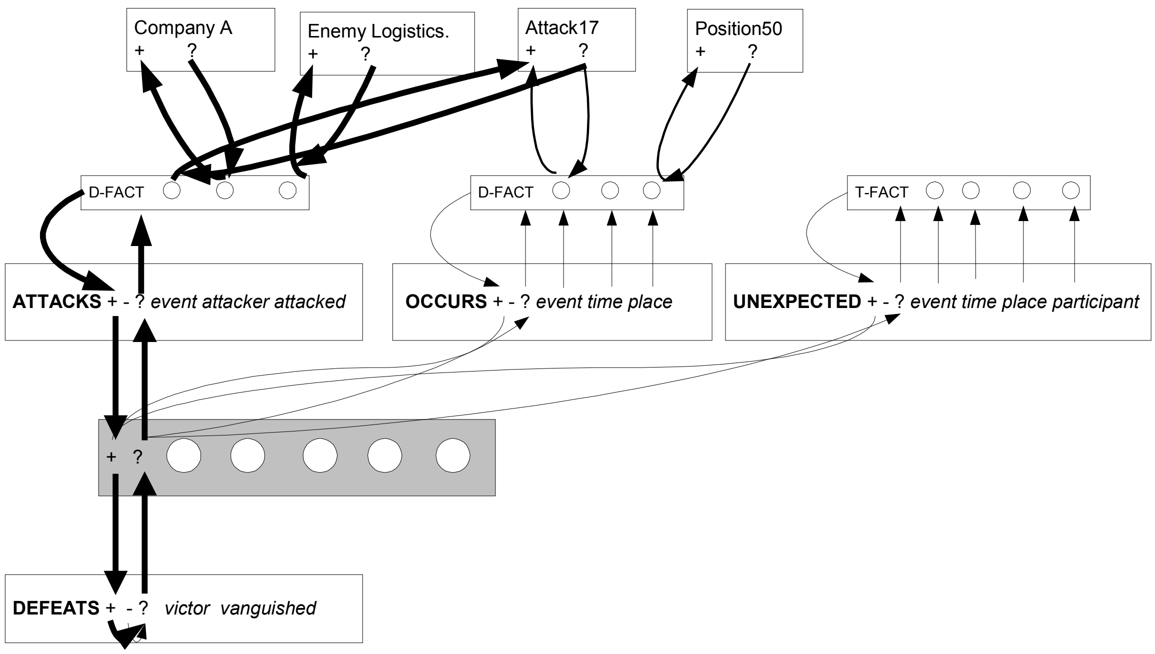

Figure 12. Combination of links in previous figures, showing a complete cycle of activation. Included are pre-existing collector-to-collector, collector-to-enabler, and enabler-to-enabler links, as well as the enabler-to-collector links closed by instantiation of two relations (**attacks** and **occurs**). Heavy lines indicate the cycle of activation involving the relation **attacks**. A reinforcing, partially overlapping cycle of activation exists involving the relation **occurs**. However, although **unexpected** is queried, the cycle involving **unexpected** is not yet closed.

---

[15] Strength of activation in a simple, single antecedent cycle is roughly equal to the product of the input activation times the weights of all links in the cycle. For the multiple antecedent rule in Figure 12, the strength of activation in the cycle depends on the combination rule at the rule mediator node. Activation of the consequent by only two of the three antecedents in this example depends on use of a soft-AND combination rule.

[16] Enabler-to-enabler querying continues up a backward chain of rules until the inferential horizon is reached. Any facts encountered along the way close a loop of activation, and allow support to return down the corresponding collector-to-collector chain.

Limits on Computational Resources

Shruti identifies a number of constraints on the representation and processing of relational knowledge and predicts the capacity of the active (working) memory underlying reflexive reasoning (Shastri, 1992; Shastri & Ajjanagadde, 1993). First, on the basis of neurophysiological data pertaining to the occurrence of synchronous activity in the γ band, Shruti leads to the prediction that a large number of facts (relational instances) can be active simultaneously and a large number of rules can fire in parallel during an episode of reflexive reasoning. However, the number of distinct entities participating as role-fillers in these active facts and rules must remain very small -- no more than [π_m /Ω], where π_m is the maximum period of oscillation, and Ω is the window of synchrony. Biologically plausible values of π_m and Ω suggest a limit of about 7. Second, if we assume that the quality of synchronization degrades as activity propagates along a chain of cell clusters, it follows that systematic reasoning involving dynamic bindings can only be sustained over a limited depth. As the depth of inference increases, the jitter in the propagation builds up until eventually, inference reduces to a mere (associative) spreading of activation. Thus Shruti assumes that reflexive reasoning has a limited inferential horizon. The depth of horizon, however, can be modulated by attentional mechanisms. Third, Shruti predicts that only a small number of *instances* of any given relation can be active simultaneously, and this also limits the depth of recursion.

The need for summary information.

Because of limits on computational resources, information that is relevant to an inference may be quiescent, or beyond the current inferential horizon. *Taxon-facts* provide a local estimate, or average, of the impact of information that is currently not active, but which has been experienced in the past. Taxon-facts are formed and tuned by a history of exploring (in mind as well as action) possible explanations of a relation, and every relation that has been explored in this way is associated with a taxon-fact. As attention shifts, however, taxon-facts are, in effect, dynamically decomposed into the specific inferences and facts that apply in the current situation.

Taxon-Facts as Prior Probabilities.

There was no initial activation (by observation or report) of the third antecedent, **unexpected**. However, that relation would be *queried* immediately after the activation of **defeats** via enabler-to-enabler links. Shruti asks the question, Will attack17 at Position50 and time 0400 be a surprise? Even though specific information about the objects in a relation may not available (i.e., in the form of dynamic or episodic facts), Shruti still draws on *general* information about the *types* of objects in the relation. For example, we may have information regarding the overall likelihood that an attack at this *type* of place and/or at this *type* of time will be a surprise. We might know that Position50 is a mountainous area, and that attacks in mountainous areas have a high probability of being unexpected. We might also know that time 0400 is at night, and that attacks at night have a better than average chance of being unexpected. In the absence of specific information about Position50 and time 0400 (in the light of this particular enemy), this *general* information is used as a prior probability that those *specific* objects satisfy the relation. Prior probabilities in Shruti are known as *taxon-facts* because they are based on semantic (i.e., taxonomic) information about the *types* of objects that appear in instantiated bindings.

Unexpected will now return activation to the positive collector of the rule mediator cluster because there is a *fact* in which the appropriate objects satisfy the **unexpected** relation. The taxon-fact closes the loop between enablers and collectors in backwards, or abductive, reasoning, but its strength of activation corresponds only to the prior probability, or historical average. The activation of **attacks**, **occurs**, and **unexpected** then determines the strength of the activation that ultimately flows through the rule mediator cluster to the conclusion, **defeats**.

Causal Explanation

The example just described involved a prediction of the consequences (**defeat**) of an observed or reported event (**attack**). We can step back, however, and look at some of the reasoning that might have led Company A to decide to attack Position50 in the first place. Decision makers infer the cause or causes of observed events. For example, observed actions may suggest the intent of the enemy, and the intent may lead to expectations of specific enemy actions in the future.

Figure 13 starts with two dynamic facts, the observation that *trucks and other supply vehicles are **near** a town called Sanna's Post,* and determination that *they **belong to** the enemy's 37th Division*.[17] These dynamic facts supply activation to the positive collectors of **near** and **belong to***,* which then flows through their enablers. Enabler-to-enabler connections allow the system to query for any information that might explain these observations. In this example, two rules, offering two different explanations, are found: Sanna's Post may be a **logistics** facility for the enemy 37th Division, or it may be intended as a **deception** to convince us that Sanna's Post is being used as a logistics facility.

[17] The example in this section is drawn from a Tactical Decision Game called "Battle of Sanna's Post" and published in the *Marine Quarterly*. We used it as part of the training evaluation described in Chapter 9, Volume II.

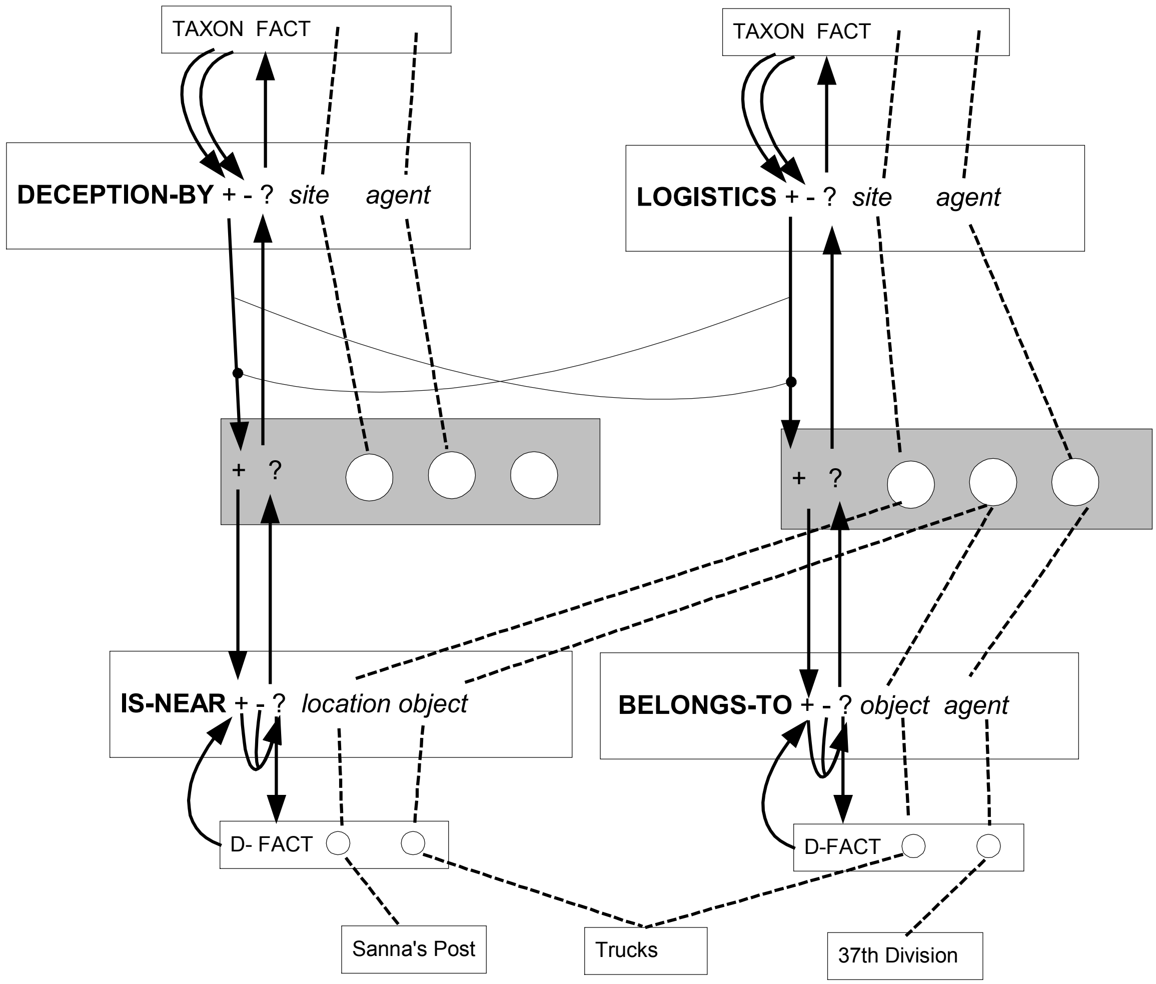

Figure 13. Two dynamic facts (trucks are **near** Sanna's Post, and the trucks **belong to** the enemy 37th Division) with two competing explanations: Sanna's Post serves as a **logistics** facility for the enemy 37th Division or it is intended as a **deception**. Object identity constraints are shown in full only for the rule supporting explanation as a **logistics** post.

Competing explanations.

Each of the two explanations is supported by taxon-facts, or prior probabilities. These may be based on the similarity of Sanna's Post to other sites used for logistics by divisions of this enemy army, and on the frequency with which this enemy has used deception in a situation of this kind. Each rule supports a completed circuit of activation, whose strength is determined by these taxon-facts and by weights on the collector-to-collector links. If the activation in one circuit is significantly weaker than the other, it will be suppressed by inhibition from the other circuit, and a single explanation will emerge reflexively. On the other hand, if both explanations receive significant, and comparable, support, they will both remain activated. It is, of course, not impossible that a single set of observations might have more than one cause. However, if the network contains other information implying that the two explanations are mutually exclusive, they should not both be active. Activation will flow to both the positive and negative collectors of some of the relations (e.g., Sanna's Post will appear both to be and not to be a **logistics** base), and the system will be in *conflict*. As we shall see in the next chapter, reflective processes are utilized both to diagnose and to resolve this type of uncertainty.

Taxon-facts at the edge of the active network.

Within Shruti, various biologically plausible constraints interact with the mechanisms of inference with the result that the propagation of activity throughout the causal network is limited. As noted, there are several sources of these limits, including allocation of temporal phases for dynamic binding, patterns of independence in the causal network (that limit the degree to which activation is self-reinforcing), limits on multiple instantiation of variables and rules, and attenuation of belief signals below plausible representational thresholds. In Shruti, the vast majority of all variables and rules within the network will be *quiescent* when reasoning in a problem domain of any plausible size. This quiescence represents a lack of information, not an assertion that the variable is false. (Negation is handled by explicit support for the falsity of the variable.)

The limits on the parallel spread of activation result in the dynamic determination of network boundaries. In our example, evidence regarding **defeats** grounds out at its possible causes, **attacks**, **occurs**, and **unexpected**. These relations, and the rules that link them, are meant to express only a portion of a vastly larger causal model. In general, the backward chaining of inference from effects to causes will have to ground out at some point on relations in the knowledge base. In order to determine the likelihood of consequences of those relations, we need to know the prior probability of those for which no dynamic or episodic facts exist (e.g., **unexpected** in Figure 12). Similarly, if we wish to infer the probability of different possible causes given information about their consequences (as in Figure 13), we must know the prior probabilities of the causes. Since backward chaining inference can ground out at any relation in the knowledge base, we need to have priors for *all* relations within that knowledge base. These priors are automatically adapted by Shruti using patterns of co-occurrence. In order to make use of these ubiquitous priors, we need to have them be active whenever they are encountered during backward chaining, and we need them to automatically *decompose*, or dis-aggregate, whenever the backward flow of activation passes through the relation to causes that are prior to that relation. By deactivating the prior associated with the relation being decomposed, and by activating the priors associated with *its* causes, Shruti automatically disaggregates the (statistically valid) prior probability into its dependency relationships. In turn, these relationships account for specific circumstances directly, i.e., do not rely on the statistical aggregation.

Quiescent knowledge.

Psychological research suggests that not all knowledge in human long-term memory is active, that is, readily available to the decision maker, at all times. In this example, there might have been additional knowledge in long-term memory that, had it been active, could have supported a more extended chain of reasoning about the instantiation of **unexpected**. Results in Shruti depend more heavily on prior probabilities (taxon-facts) when more specific knowledge cannot be activated due to limits on computational resources and, in particular, the limited inferential horizon. Figure 14 provides examples of two chains of reasoning (forward from cause to effect, and backward from effect to cause, respectively) that might have shed more light on whether the attack will be **unexpected** by the enemy:

(1) Forward reasoning from cause to effect. Imagine that the command staff of Company A have not yet conducted a personal reconnaissance of the terrain at Position50. However, reports have been received by intelligence officers from local inhabitants that imply that concealment near that position is *below average for mountainous terrain*. This episodic fact, in conjunction with rules linking **concealment** to **unexpected**, implies that the chance of surprise is less than predicted based on the prior probability, which is simply an average.

(2) Backward reasoning from effect to cause. Suppose there are also reports that the enemy may be preparing for a night attack at Position50. Since, preparation is a typical effect of expecting an attack, it suggests that the enemy does expect the attack at that time and place.

In our example, these two kinds of information are not taken into account because of cognitive resource limitations. Activation beginning with the dynamic instantiations of **occurs** and **attacks** did spread (via collector-to-collector links) to **defeats**, and from there (via enabler-to-enabler links) to **unexpected**. If there were no computational limitations, activation would then continue to flow from **unexpected** via enabler-to-enabler links to **concealment**, where the episodic fact (absence of good concealment) would close the loop and send support to the negative collector of **unexpected**, ultimately decreasing confidence in the prediction that the enemy will be **defeated**. Similarly, if there were no resource limitations, collector-to-collector links would take activation from the negative collector of **unexpected** to the predicted effect, the positive collector of **preparing-for-attack**. The confirmation of the prediction by an episodic fact (that preparation is taking place) would amplify that loop, returning more activation to the negative collector of **unexpected,** disconfirming the claim the attack is likely to be **unexpected**, and reducing confidence in **defeating** the enemy.

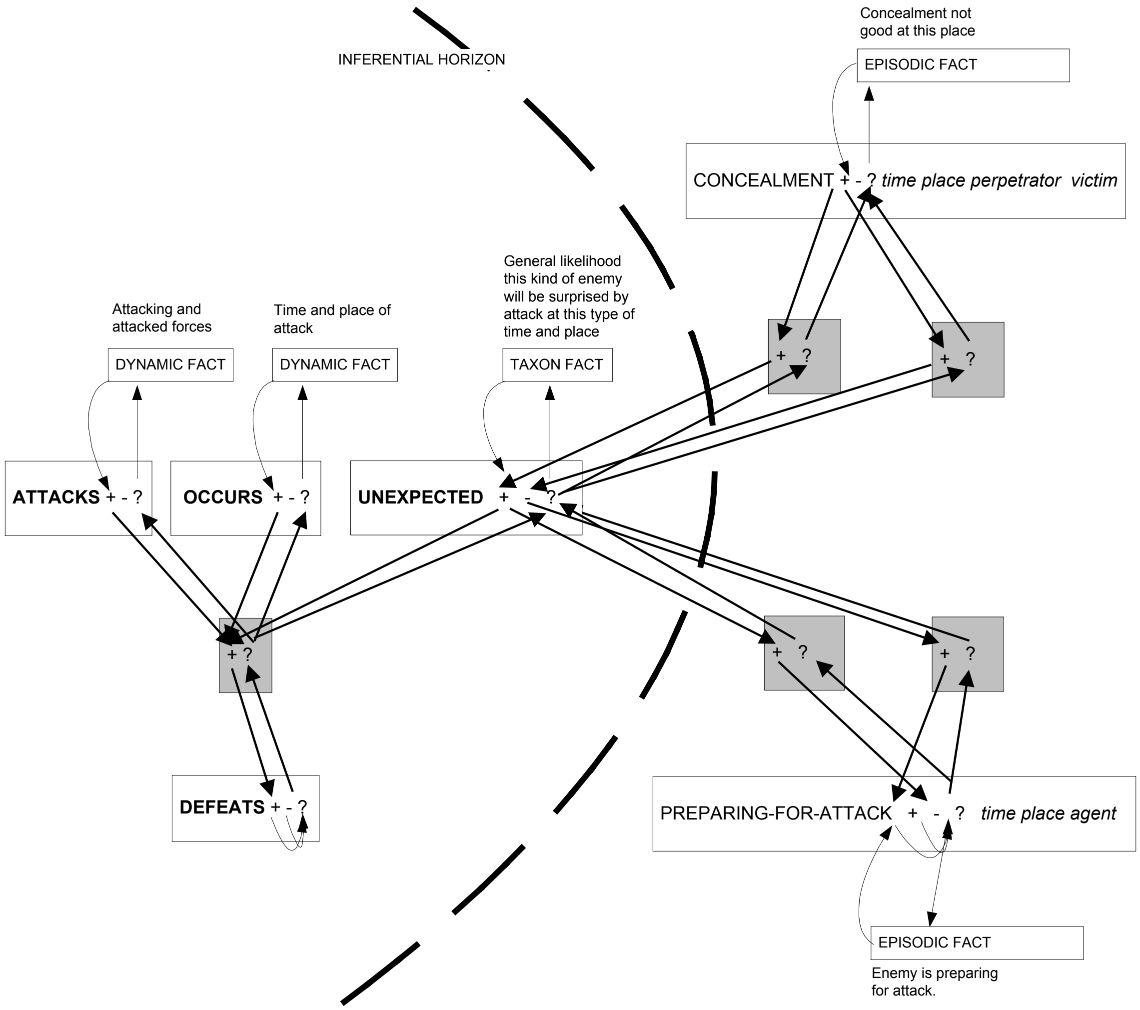

Figure 14. Inferential network in which certain rules and episodic facts lie outside the inferential horizon. As a result, the taxon-fact regarding **unexpected** comes into play. For clarity, many links and role fillers are omitted.

Unfortunately, as Figure 14 shows, both of these lines of reasoning do lie beyond the current inferential horizon (which centers on the original points of activation at **occurs** and **attacks**). In other words, the commander simply does not "make the connection" to this information. When the spread of activation stops at the **unexpected** relation, the prior probability of this instantiation of **unexpected** is the sole determinant of the activation that travels back to **defeats**.

As noted, priors are automatically adapted by Shruti to co-occurrence patterns in the decision maker's past experience. In this example, the strength of the taxon-fact is an average of all past information regarding **unexpected**. Suppose that in the decision maker's past experience, there has been good concealment in mountainous terrain 70% of the time, and no concealment 30% of the time. As a result, the *average* activation returning to the positive collector of **unexpected** from **concealment** over past incidents (when the decision maker did actively consider the information about concealment) has been approximately 700 out of a thousand, and the average returning to the negative collector of **unexpected** from **concealment** has been 300. The current situation, however, is far from the average. A more accurate estimate of the chance of concealment in this specific situation might be 10%. Prior probabilities allow a summary of currently unattended information to play a role in reasoning, but the solution is imperfect when the specific situation deviates significantly from the average.[18] The lack of specific information forces Shruti to operate reflexively with statistical averages, and in this case to produce an excessively optimistic prediction regarding success of the attack.

One of the key contributions of reflective processes is their ability to shift attention within the belief network, to activate areas of knowledge that are likely to significantly change the current reflexive conclusion.

Decision Making and Action

In real-world tasks, situation assessment and decision making tend to be tightly interwoven. Observations lead to predictions of events and those predictions trigger goals, which may lead rapidly and reflexively to actions or decisions. Shruti performs reflexive decision making through the very same processes that it uses for reflexive situation assessment. Shruti not only seeks *explanations* of activated events, but seeks actions to *produce* desirable events and *prevent* undesirable ones.[19] Moreover, the need to act guides the process of seeking explanations, and explanations guide the need to act.

Utility and Feasibility

The basic ingredients of reflexive decision making in Shruti are the following nodes associated with relations:

1. Weights on *J-facts* represent the expected utility of an instantiated relation (or its negation). J-facts are important only for relations that are on the edge of the currently active network, and they are an average based on past experience of the consequences of the relation that are currently not active in the network. The J-fact associated with a relation does not include its intrinsic utility, only the utility it inherits from its consequences.
2. Weights on *utility facts* represent the *intrinsic* desirability of an instantiated relation's being true or false. The J-fact of a relation aggregates over the utility facts associated with the currently inactive consequences of the relation.
3. Weights on *feasibility facts* represent the expected availability of a feasible action sequence that could be used to make a particular relation true. Feasibility facts aggregate over the causal precursors (rather than consequences) of a relation.

[18] Priors also average in the effects of information the decision maker that is not actually in long-term memory, but which the decision maker has collected in the past by further investigation of the external environment.

[19] Some of the design concepts to be described in this section, pertaining to utility and decision, have not yet been fully implemented.

4. Activation of the "+w" or "-w" node of a relation represents the current *preference* for acting to make that relation true or false. This activation is a function of cumulative activation across a closed loop that includes relevant J-facts, utility facts, and feasibility facts (as well as enabler to enabler and w-to-w connections).

J-facts, utility facts, and feasibility facts together create closed, self-reinforcing loops of activation through +w and -w nodes and enablers. Strong activation in a +w (or –w) node corresponds to *wanting* to act to cause the truth (or falsity) of the associated relation, because (i) there is a feasible action to do so, and (ii) the expected utility of doing so is high. Stable cycles of activation through the +w or –w node of a relation makes that relation a goal or subgoal, and is the precursor to forming an intention or plan and acting on it.

Formation of Intent

Figure 15 shows how values and available actions interact to generate reflexive decisions. In this example, Company A's assigned mission is to **destroy** a logistics post belonging to the enemy, and the commander needs to decide how to do so. The relation, *Company A **destroys** logistics* is attended and hence queried as a result of linguistic decoding of the order to **destroy** the enemy logistics post. This querying produces activation in the enabler ("?") of **destroys.**

Whenever the enabler ("?") node of a relation is activated, for whatever reason, Shruti searches for other relations that might causally lead to (or explain) the relation in question. This search involves the flow of activation along the enabler-to-enabler links that lead from **destroys** to **attacks** (since **attacking** is a means of **destroying**). Two elements must be present if a complete cycle of activation is to return to the "?" node of the original relation, **destroys**:

(i) A feasible action must be represented in long term memory, by means of which the decision maker could bring about one of the relations in the causal chain. In Figure 15, activation flows from the "?" node of **destroys** through enabler-to-enabler links until a relation is found that is associated with an appropriate action. It turns out that **attacking** the logistics post is a feasible action. The presence of this feasibility-fact provides a switch that allows activation to return from the enabler of **attacks** to the +w node of **attacks** and from there to the +w node of **destroys**.

(ii) A high positive expected utility must be associated with the original relation.[20] In Figure 15, the J-fact associated with **destroys** reflects the influence of the assigned mission, and the potential contribution of destroying the logistics post to accomplishment of all relevant purposes. This J-fact provides a switch that allows

[20] The "expected utility" of a relation roughly comprises (i) weights on the J-facts of any of its active consequences that happen to be on the edge of the currently active network, and (ii) utility facts (if any) of both the relation and its active consequences. There will be no utility facts in the currently active network if the utility of active relations is not intrinsic but inherited from currently inactive consequences. For this reason and for simplicity, we have only represented J-facts in the diagrams.

activation to flow from the "+w" node of **destroys** to the enabler ("?") node of **destroys**, where the cycle began.

In some instances, more than one feasible action is available that could produce a highly valued consequence. In this case, Shruti can reflexively settle on the "best" of the competing options, in the same way that it settles on the "best" explanation of an observation when more than one is available (see Figure 13). In problematic cases, where there exists significant conflict, gaps in information, or hidden assumptions, reflective intervention may result in a deeper exploration of the sources of utility and feasibility of the competing options.

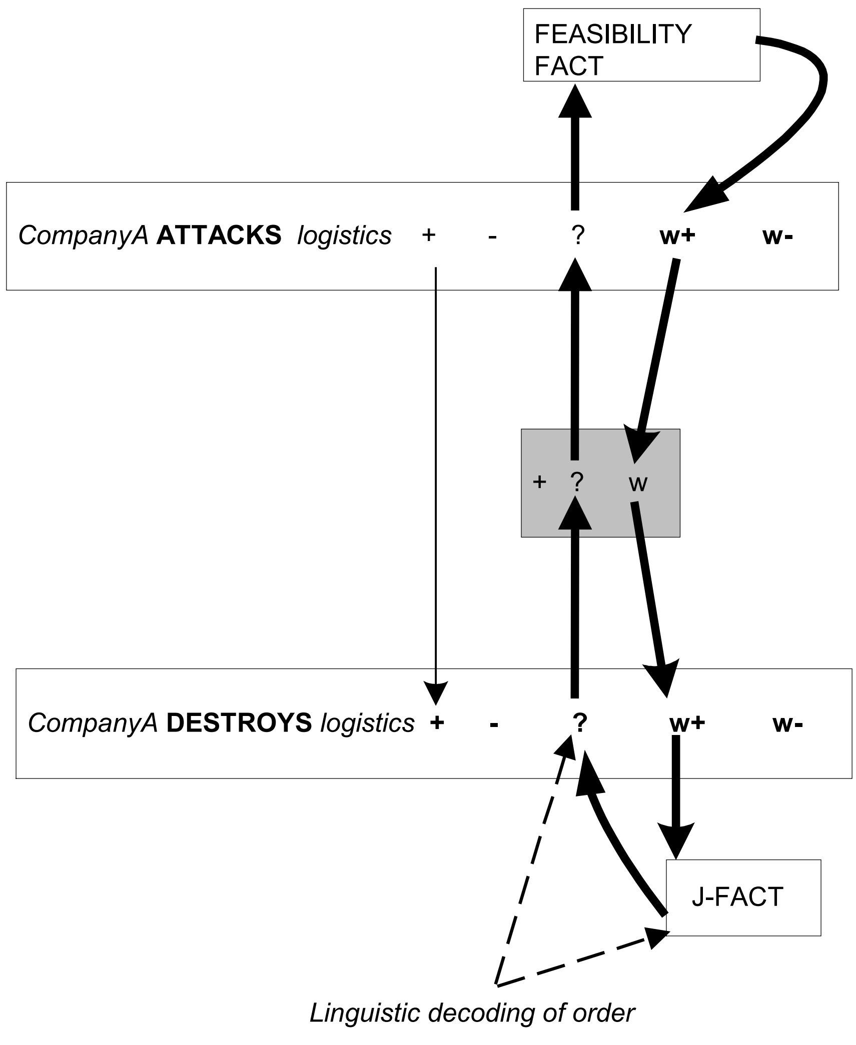

Figure 15. A simple network in which an intent is generated to perform an action because it has the consequence of bringing about a desired state of affairs.

Prediction and intent.

In Figure 15, a decision to **attack** was prompted by an order to **destroy**. This order had two effects: an increase in the expected utility of **destroying** and a query for potential actions that could cause the logistics post to be **destroyed**. Another way that a reflexive decision process could get started involves a *belief* rather than a specific order. As shown in Figure 16, the instigating event might be new predictions based on observation or analysis that a positively valued event will not occur. In this example, an instantiated relation (**destroys**) has preexisting positive expected utility, and the decision maker concludes for whatever reason that the relation is *not* likely to occur in the absence of action. This belief would generate a query that might lead to discovery of a feasible action (**attack**) to bring the relation about. This is represented in Figure 16 by activation flowing from elsewhere in the network to the negative ("-") collector of **destroys** and from there to the enabler of **destroys** and on to the enabler of **attacks**.

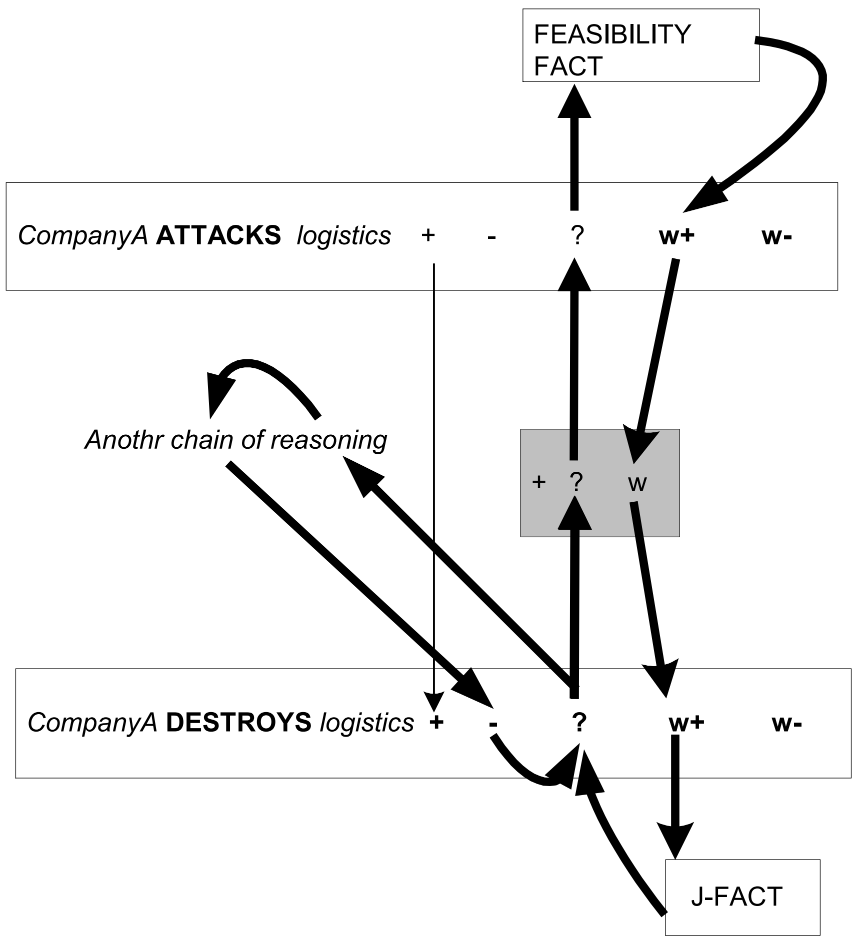

Figure 16. Initiation of a decision process by a prediction that a positively valued outcome will not occur in the absence of action to bring it about.

Another possibility is represented in Figure 17. Here, the prediction that a *negatively valued* outcome *will* occur (the enemy **reinforces** its forces, using the logistics post) leads to a decision process to find a way to prevent it (e.g., **destroy** the logistics post). This example illustrates several other points as well. First, negative expected utility is represented by a J-fact (or utility fact) that links the -w node, rather than the +w node, to the enabler of **reinforces**. Second, **destroys** is a potential cause of the enemy's *failing* to **reinforce**. This negative causal relationship is reflected by the link between the "+" node of destroys and the "-" node of reinforces. Third, to the extent that **reinforces** is negatively valued, anything that prevents it (such as **destroys**) is positively valued. So, because of the negative causal relationship, there is also a negative relationship between preferences for the two relations. Thus, the -w node of **reinforces** is linked to the +w node of **destroys**.

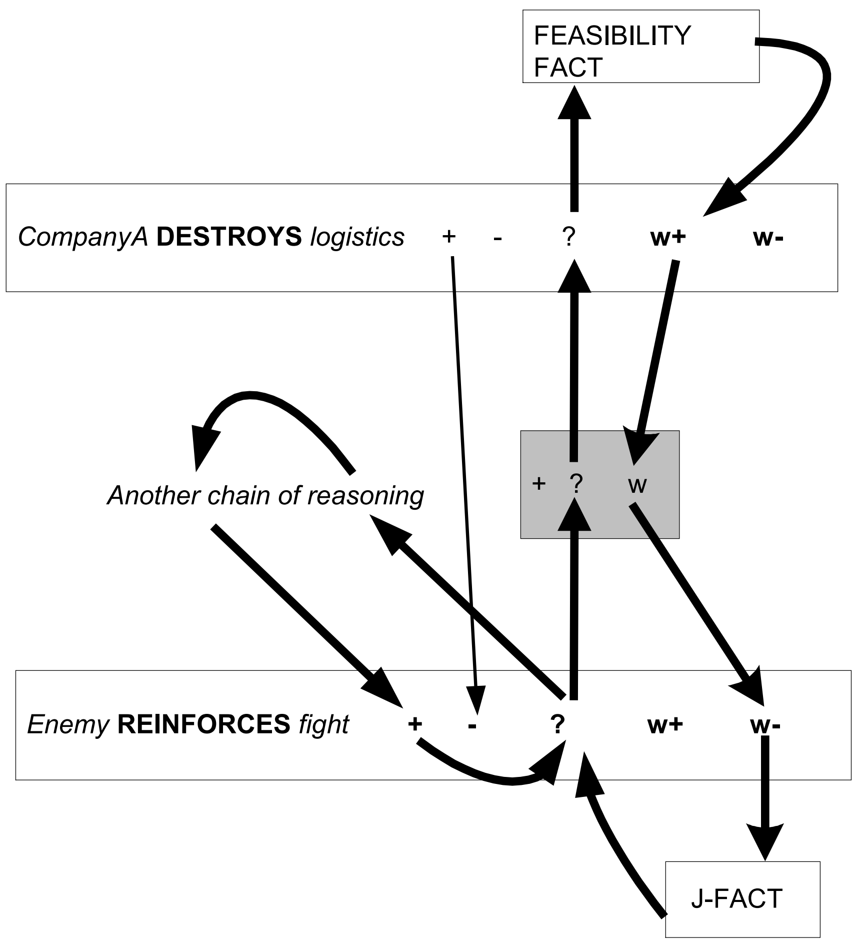

Figure 17. Instigation of decision process by prediction that a negatively valued event is likely to occur in the absence of action to prevent it.

More generally, w-to-w links run parallel to collector-to-collector links, since both represent causal links between relations. Because of this, prediction of a negatively valued event leads to a stable cycle of activation only if there is a feasible action that prevents that negatively valued event. Thus, belief that the enemy will **reinforce** leads to querying for an action, such as **destroying**, that will prevent it. **Destroying** becomes desirable (activation of +w) to the degree that the predicted event (**reinforces**) was undesirable (activation of –w). These connection thus allow the return flow of activation through the +w node of **destroys** to the –w node of **reinforces**.

It is important to understand that the feasibility of **destroying** the logistics post is an *average*, based on the historical feasibility of **attacking**. Since **attacking** and the causally prior steps involved in implementing it are outside the active network for the time being, the estimated feasibility of **attacking** does not take into account many of the particulars of this situation. Reflective processing, in which additional cycles of attention are directed to these issues, will be necessary to verify this decision.

Finding a feasible action to cause or prevent an event is analogous to the discovery of an explanation for an observed event or belief.[21] Feasibility facts can be uncovered by this flow of activation in the same way as dynamic or taxon-facts are uncovered in the quest for an explanation. Rather than showing *why* the relation is true or false, the feasibility fact shows *how* it can be *made* true or false.

What if a positively valued event is expected to occur, or a negatively valued event is expected not to occur, *without any action* on the part of the decision maker? For example, suppose that situation assessment led to the belief that the logistics post will be **destroyed** independently of any action by Company A to **attack** it. In this case, the cycle of activation leading to a decision will *not* occur. As shown in Figure 15, the *positive* collector of **destroys** inhibits activation flowing from the *positive* J-fact of **destroys**. There is no point in intending to make a relation true, no matter how desirable it is, if it is already true or predicted true. Similarly, although not shown, the *negative* collector of a negatively valued relation will inhibit activation from the *negative* J-fact. Notice that the activation flowing through w nodes represents the *preference for an action*, and when stable, the formation of *intent*. It does not represent the experience of *desirability*. A relation may be just as desirable or undesirable whether or not action by the decision maker is either sufficient or necessary to bring it about or prevent it.

In precisely the same way, *after* a successful **attack** is executed, dynamic facts are instantiated corresponding to Company A **attacks** logistics and logistics is **destroyed**. This leads to a self-reinforcing cycle of activation through the positive collectors ("+") of **destroys** and **attacks** This in turn shuts off the desire to *make* **destroys** true, since it is now known to be true. Thus, goals are released as actions are executed and purposes achieved.

---

[21] When options are mutually exclusive, each suppresses the other in a winner take all competition. However, independent (i.e., non-mutually exclusive) options are combined using abductive metrics similar to Bayes rule for inferring the probability of different mutually compatible explanations.

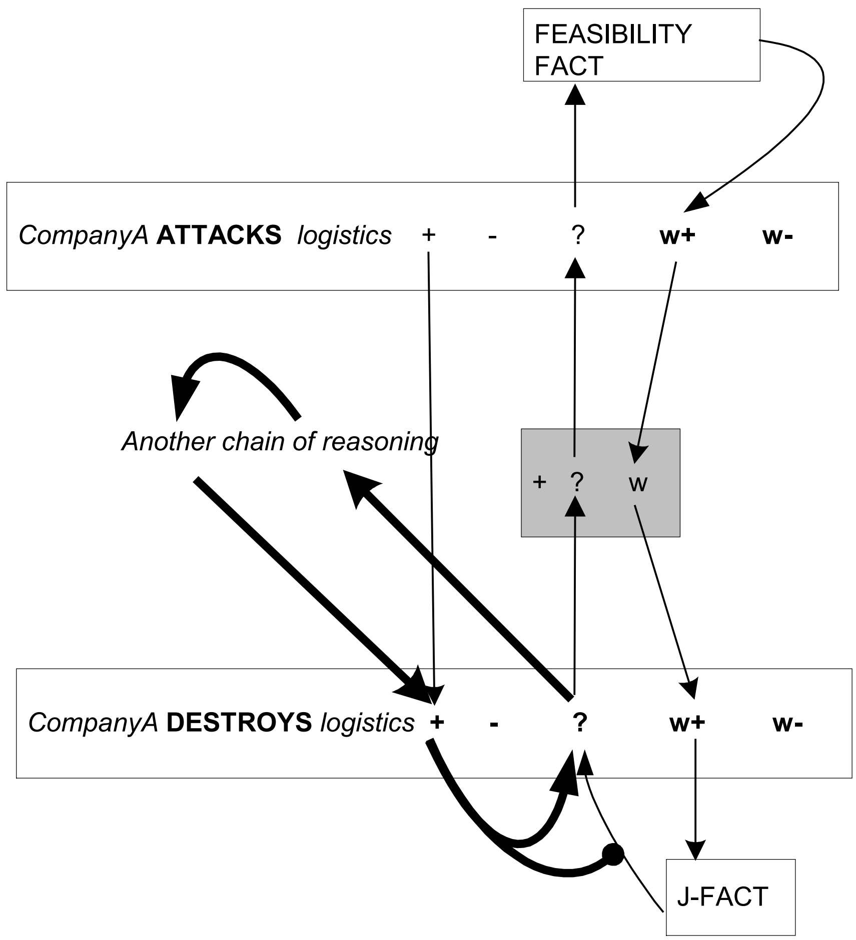

Figure 18. A prediction that the logistics post will be **destroyed** (e.g., by other friendly units) suppresses the cycle of activation that leads to Company A's decision to attack.

Inferencing and decision making.

Inferencing and decision making work hand in hand. The value of an event can amplify the force of the inferencing process concerning that event. Any prediction of the occurrence or non-occurrence of an event leads to queries (via enabler to enabler links) regarding other possible causes of its occurrence or non-occurrence. However, this process is more intense when the system predicts that a *negatively* valued event *will* occur or that a *positively* valued event *will not* occur. This same intensified querying process also searches for feasible actions to prevent the negatively valued event or bring about the positively valued event.

Conversely, the formation of intent is guided by causal inferencing through the discovery of positively and negatively valued consequences of a relation and feasible actions to bring it about. Inference also plays an important role in showing when action is unnecessary. For example, queries triggered by the prediction that a desirable state of affairs will not occur may ultimately show that the event *will* occur after all without the decision maker's intervention. Similarly, the prediction that a negatively valued event will occur may be discovered to be mistaken. In these cases, the inference process, although running parallel to the search for an action, eventually brings that search to a halt.

In reflexive decision making, the system rapidly settles on an *intent* to act in the same manner, and as part of the very same process, in which it settles on a situation interpretation (i.e., a self-reinforcing set of observations, predictions, and explanations). Actual performance of the action may sometimes be immediate; in other instances, there may be a delay of a second or two (e.g., responding to a subordinate's question with an order), or longer (e.g., waiting until night to attack the enemy logistics post). According to recent research in experimental psychology (Brandimonte, Einstein, & McDaniel, 1996), knowledge of a future action that an agent intends to perform is stored in *prospective* memory. This is analogous to episodic memory of *past* events, but it is the memory of actions that the agent has *decided* to perform in the future. An effective prospective memory is as important in highly skilled performance as an effective episodic memory. In the special case where the action is appropriate immediately, the prospective fact is like a zero latency timer, and executes as soon as the activation loop is closed. When the delay is longer, of course, there is the possibility that intervening events might destabilize the solution and change the intent before the action is executed.

<u>Conflicting Goals and Actions</u>

Reflexive decisions sometimes involve uncertainties that it pays to examine more closely through more reflective processes. A simple example is the case in which an action has both good and bad consequences, i.e., there is a *conflict* among goals. Figure 19 illustrates such a conflict. Attacking the logistics post will degrade the enemy's capability to support reinforcements, but it will also expose the presence of Company A to the enemy, diminishing its ability to ambush any reinforcements that do attempt to go through. One goal (preventing the **reinforcement** of the enemy by use of the logistics post) favors **attacking** the logistics post, while another goal (keeping the presence of the company **unknown**) supports not **attacking** the logistics post. In this case, there are competing cycles of activation supporting both the performance and the withholding of the action (**attacking**).

If the strength of the activation through one cycle is significantly greater than the strength of the activation through the other. the system will resolve the conflict reflexively, quickly settling on one or the other of these options. The difference in activation strengths will depend on the respective utility and J-facts associated with the outcomes of the two options, and the causal, w-to-w connections between the actions and those outcomes. If each option has significant strength, reflective strategies may be called into play to identify and resolve the conflict. As we shall see in Chapter 13, such reflective processes may dig in long-term for relevant information that is not currently active memory but which can increase understanding of the decision.

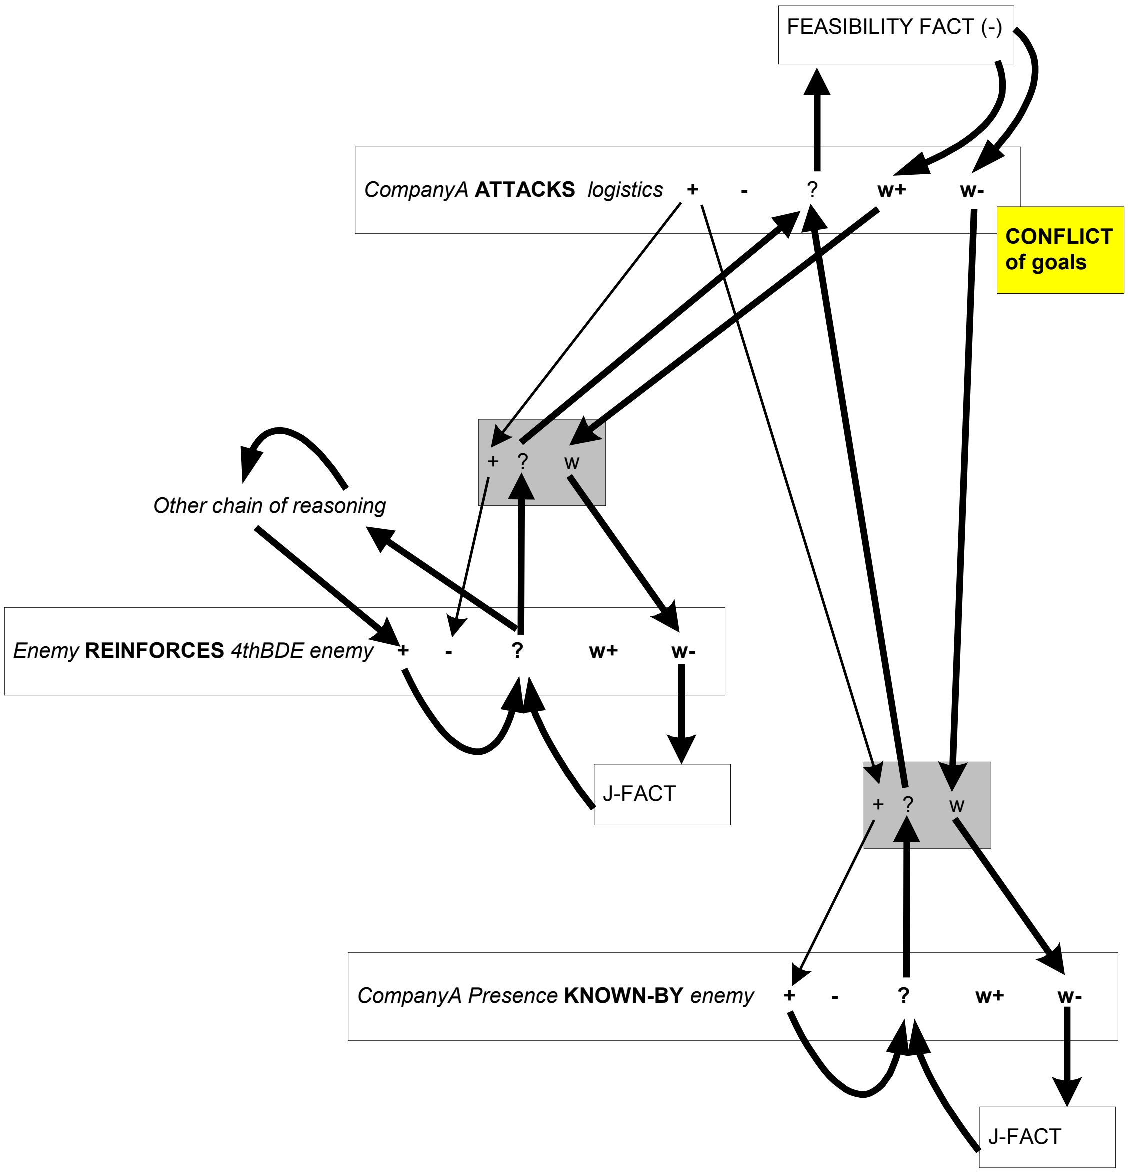

Figure 19. An action (**attacking** the logistics post) (i) prevents one undesirable state of affairs (enemy **reinforces** enemy 4th Brigade) but (ii) causes another undesirable state of affairs (presence of Company A is **known** by enemy). Both performing the action and withholding the action have inherited value, leading to conflict.

Statistical Averages, Attention, and Novelty

Limits on computational resources typically cause activation to stop before reaching "terminal nodes" of a potentially relevant web of reasoning. The focus of decision making is often somewhere in the middle of such a web, starting with goals that tend to be *signals* of positive utility rather than desired in themselves, and settling on an intention whose feasibility requires more fleshing out before an action can actually be implemented. Similarly, in seeking a causal explanation, Shruti typically reasons a few steps back to precursor events, but stops well before it has explored all the causal conditions under which the precursor events would or would not occur.

To handle limitations on the scope and complexity of reasoning, Shruti associates decomposable associative memories with all variables in the network, i.e., taxon facts, utility facts, and feasibility facts. When variables are at the edge of the current reasoning horizon, Shruti completes cycles of activation by utilizing associative memories containing statistical aggregations of past experiences.

In particular, weights on j-facts are tuned to remember the average cumulative expected value at the variable to which they are attached. This is accomplished by an update rule similar to the one used for J* in adaptive critic models, called *heuristic dynamic programming* (Lukes, Thompson, & Werbos, 1990).[22] The estimates of J*(*t*) distributed throughout the network (in the form of weights on j-facts) are updated by reference to the activation flowing through the w nodes. Suppose that variable A (representing a simple or complex event at time t) is a causal influence on variable B (representing a simple or complex event at time t+1), and that A is not at the edge of the current network (B may or may not be). The weight on the j-fact at A (corresponding to J*(*t*)) is adjusted to reduce an error term defined as the difference between: (1) the activation of w nodes at A (analogous to an estimate of J*(*t*)) and (2) the activation of w nodes at B (analogous to an estimate of J*(*t+1*)) plus the activation of any intrinsic utility-fact at A.

Links of some kind at the edge of the active network are necessary in Shruti because of (i) the limited spread of activation, and (ii) the requirement that activation return to its source and settle into stable cycles. There are at least two arguments for more specifically equating these links with statistically aggregated facts: (iii) The existence of a plausible learning procedure as described in the previous paragraph, and (iv) the resulting enhanced accuracy of conclusions, given the limits on computational resources. In particular, weights on j-facts and activation through w nodes are expected to be *well-calibrated* with respect to the actively represented evidence and context in which reasoners find themselves. That is, given sufficient tuning by experience, these quantities should closely approximate actual probabilities for the *conditions that are discriminated by the active network* (see Cohen, Parasuraman, Serfaty, & Andes, 1997). The problem with these estimates is not calibration, but *resolution*, i.e., the fineness with which

---

22 Heuristic dynamic programming (HDP) refers to systems that implement an incremental approximation of dynamic programming (DP) where events and feedback are experienced over time and the behavior of the system converges to results expected for DP. HDP was first formulated as a neural-network approximation to dynamic programming in Werbos (1977). We start with the following simplification of the Howard (1960) version of the Bellman equation for an absorbing Markov chain, similar to the simplification found in Werbos (1989) and Lukes, Thompson, Werbos) (1990):

$$J^*(R(t)) = \mathrm{Max}_{u(t)} \langle U(R(t),u(t)) + J^*(R(t+1))\rangle,$$

where u(t) is a vector representing the choice of action at time t, angle brackets denote the expected value, R is a vector describing the current state of the process being controlled, and R(t+1) depends -- of course -- on the choice of action u(t). HDP uses this function to compute the targets for a supervised learning module that learns, through experience, to predict J* given R.

However, the update rule for *j*-facts is constrained by the patterns of dependence among variables so that it accounts only for the variable to which the *j*-fact is attached and its effects, rather than an entire system state at time t. Also, like taxon-facts, *j*-facts will be sensitive to type information when present.

different conditions are discriminated by the active network. By taking more information into account, finer discriminations (i.e., higher resolution estimates) can always be made, and the reasoner will be able to generate more adaptive plans for *novel* experiences.

Another problem is *temporal* resolution, that is, the time required for the system to respond to changes in the environment. Temporal resolution is affected in two ways: adjustments of the situation model may be delayed, and changes in behavior as a result of changes in the situation model may be delayed. Only nodes that are active at the time can adapt their weights to surprising events. As a result, opportunities to learn the specific character of some changes will be missed due to inattention. And even when beliefs do change, the changes will not affect behavior on a given occasion if they are not within the currently active region of the network. Changes that have long-distance effects (in terms of the length of the inferential paths required to move from evidence to conclusion) will be incorporated more slowly as the statistical aggregates converge to new values.

Reflective thinking contributes both to the quality of decisions and the speed of adaptation by identifying situations in which low resolution (represented by statistical averages) are inadequate. Attention shifting under reflective control causes the reasoning horizon to change, incorporating more information into the solution. At each moment, as activation in the network reaches a new edge, the statistical estimate associated with the new edge - a taxon fact for a cause, a feasibility fact for an action, or a j-fact for a consequence - is utilized for computations of belief and expected value. As the activation passes that edge and reaches a new set of edges, the previously used statistical estimate is decomposed. That is, it is inhibited and replaced by (i) statistical estimates from the new edges, combined with (ii) specific information associated with newly active variables.

We turn to such reflective processes in the next chapter.

CHAPTER 13
THE REFLECTIVE SYSTEM

Overview of the Reflective System

A reflective architecture has been designed and partially implemented that sits on top of Shruti (as described above) and influences its performance. While Shruti supports inference based on the user's mental model of the situation and plan, this reflective layer supports critical thinking about the mental model and plan, in light of its inferential implications. The mental model, Shruti, and the reflective layer together provide an integrated ensemble of capabilities in support of the decision making task.

A technical goal in this development was to keep the reflective layer as simple as possible, and to base each reflective function on plausible psychological or neurological findings in regard to human reasoning. The hypothesis that we wish, ultimately, to test is that an extremely simple and small set of reflective functions can have a very significant impact on the performance of the reflexive system, especially for novel and uncertain decision making tasks. The human user is the source of the mental model representing the relevant research issue. Because of the computational limitations on human reflexive reasoning, reflective control over attention can significantly leverage the power of the reflexive system (Cohen et al, 1996). Reflection can help decision makers bring to bear relevant knowledge that they already possess, but do not have available for use on a particular occasion. The critical thinking training system is designed to provide support for this kind of reflection.

Recognition, as implemented by the reflexive reasoning of Shruti, and metacognition, as implemented by the reflective layer to be described in this section, interact continuously. Metacognitive processes assess the belief network to find high-leverage points of uncertainty. Shifting attention to these points results in the activation of new long-term knowledge, lying beyond the edge of the currently active belief network. Metacognitive decisions about *where* in the argument network to shift attention lead to recognitional retrieval of further information from long-term memory, which is integrated with previously activated information, leading to new reflexive conclusions. These new conclusions affect the further course of reflection, and so on. Strategies for shifting attention in response to reflexive uncertainty are learned through processes of associative and reinforcement learning, and are ultimately shaped by their results, successful decisions.

We will describe the reflective layer in several stages. First, we briefly outline the fixed structural features of the metacognitive model. These serve as the building blocks out of which a variety of different reflective strategies can be constructed. Second, we will describe how domain-specific strategies for critical thinking might be acquired as proficient decision makers learn to assemble the basic elements into more and more effective tools for critical thinking. Third, we will describe certain general features that might be shared across a variety of domain specific strategies. Truly general-purpose critical thinking strategies might be abstracted from the domain-specific strategies with the support of appropriate training, and in turn guide the development of more refined domain-specific strategies.

Building Blocks: The Basic Metacognitive Model

The fixed structural features of the metacognitive system are the following:[23]

1. All hypotheses are in long-term memory. Some (but not all) hypotheses in long-term memory are part of stable activation cycles, that is, in working memory. Some (but not all) hypotheses in working memory are in focal attention.

2. Shifting focal attention to a currently active hypothesis has the following three effects:

- Automatic awareness of the degree of activation in "+" and "-" (belief) collectors, and in "+w" and "-w" (utility) nodes of the attended hypothesis. That is, attention generates awareness of both support and preference.
- Automatic querying of the attended cluster. Attention supplies activation to the "?" node, causing the reflexive system to seek other information that bears on the attended hypothesis.
- Optionally, clamping of "+", "-", "+w" , or "-w" nodes of the attended cluster at 1.0 or 0.0. Persistent attention to the *positive* collector of a cluster creates a "what-if" assumption that the hypothesis is true. Persistent attention to the *negative* collector creates a "what-if" assumption that the hypothesis is false. Persistent attention to the "+w" or "-w" node creates an assumption that *making* the attended relation true or false, respectively, is preferred or intended. The reflexive system generates the implications of these explicit assumptions.

3. Shifts in attention can change the outcome of reflexive processing. This capability depends on the following characteristics of the system:

- As noted in 2, attention may be associated with changes in activation at the attended hypothesis through automatic querying, and optionally, through clamping of truth values or desirability. Returning activation due to such clamping may create new, stable cycles of activation.
- Shifting attention moves the center of the inferential horizon to the attended node. This brings a different set of long-term memory hypotheses into working memory, permitting new stable cycles of activation to occur, and potentially changing the activation levels at hypotheses that were already in working memory.
- What-if assumptions (persistent attention to the "+", "-", "+w" , or "-w" nodes) may lead to the activation of hypotheses that would have been suppressed by other, competing hypotheses in the absence of the assumption.

---

[23] In the following we will use the term *hypothesis* very broadly, to refer to any semantically meaningful cluster of nodes in Shruti. Shruti uses clusters of nodes to represent a relation, an episodic fact, a dynamic fact, a taxon fact, a utility-fact, a J-fact, or a decision. A hypothesis may therefore be a general or specific proposition about facts or values, or an action option.

- Attentional shifts influence subsequent reflexive processing even after attention moves on. During the initial attentional cycle, activation causes priming of the affected hypotheses. During subsequent attentional cycles, priming maintains the hypotheses in working memory even if they are no longer within the inferential horizon. Conclusions reached after multiple attentional cycles therefore will integrate a greater span of long-term knowledge than conclusions reached after a single cycle of reflexive processing.

4. Through associative and reinforcement learning mechanisms, the metacognitive model has the ability to learn the following:

- Local patterns of activation. The metacognitive system can learn to recognize recurrent, meaningful patters of activation of "+" and "-" collectors, and "+w" and "-w" nodes, at a single attended hypothesis. Such local patterns might include, for example: both "+" and "-" collectors of a cluster have low activation; both have high activation; there is a large difference between "+w" and "-w" activation; and so on. These meaningful local patterns of activation correspond to *different types of uncertainty*.
- Influence relationships, i.e., correlations between changes in local patterns of activation at specific hypotheses and changes in activation at specific other hypotheses. With experience in a domain, the system learns which relations, facts, or decisions tend to be responsible for, or have the most impact on, uncertainty at other relations, facts, or decisions.
- Argument roles. Abstractions of these learned correlations can lead to the ability to recognize different functional roles that hypotheses play in arguments across all domains: i.e., conclusions versus evidence versus rebuttals.
- Domain-specific attentional behavior. Through experience, the system can learn to shift attention to, and/or clamp activation at, hypotheses in response to (i) local patterns of uncertainty, (ii) correlations of those patterns with activation in other hypotheses, and (iii) associations between attentional behavior under conditions defined by (i) and (ii) and changes in utility (w). Second-order reflective behavior is shaped by its consequences, like other recognition-based behavior.
- General reflective behavior. Through experience and education, the system can learn to shift attention and/or clamp activation at, hypotheses in response to (i) local patterns of uncertainty, (ii) the pattern according to which other hypotheses fill functional roles (e.g., evidence, rebuttals, conclusions) in arguments regarding the hypothesis of interest, and (iii) strategies that specify which functional roles are likely to be most relevant to resolving different types of uncertainty.

This small set of fixed elements can lead, through learning, to a rich repertoire of reflective, or critical thinking, skills. We will discuss some of the potential results of such learning in the remainder of this chapter.

Qualitatively Patterns of Local Uncertainty

Given that attention to a hypothesis brings awareness of the activation levels of its component nodes, what sense will the metacognitive system make of these activation levels? Qualitatively different *types of uncertaint*y are definable in terms of different local patterns of activation at attended hypotheses. Because these different types of uncertainty arise from consistently different conditions elsewhere in the belief network, they call for different responses if they are to be most efficiently resolved. It is plausible, therefore, to suppose that the metacognitive system takes advantage of these readily available local clues, learns to recognize the qualitatively different types of uncertainty that they signal, uses them to infer problems elsewhere in the active belief network, and leverages such inferences to guide attentional behavior.

Identification of qualitatively different types of *local* uncertainty must, in Shruti, be based on the possible activation states of the positive ("+") and negative ("-") collectors, and the "+w" and "-w" nodes of each hypothesis. We will focus here on the "+"/"-" collectors, since patterns for "+w" and "-w" nodes are analogous. Activation at each collector varies from 0 to 1 independently of the other collector. It follows that the unit square shown in Figure 20 includes, and is sufficient to recognize, all possible types of local uncertainty. The four most extreme (and qualitatively distinct) possibilities are at the corners of the square, and two of these represent different categories of uncertainty:

(a) neither + nor – activated (*incompletenes*s, or insufficient information)

(b) + activated but not - (tendency to decide for hypothesis)

(c) - activated but not + (tendency to decide against hypothesis)

(d) both + and - activated (*conflict*, or contradictory information)

Figure 20 illustrates the two types of uncertainty. (a) and (d) in Figure 20 represent total incompleteness (0,0) and total conflict (1,1), respectively, concerning the relevant hypothesis. The solid diagonal in Figure 20, which runs from (0,0) to (1,1), represents an axis along which incompleteness and conflict trade off with one another, as activation increases or decreases at "+" and "-" collectors by comparable amounts.

A third kind of uncertainty is identifiable in the center of the unit square. Here, as for incompleteness and conflict, activation at the "+" and "-" collectors is approximately the same. However, there is *too much* information to characterize the uncertainty as incompleteness, which is the failure to develop relevant arguments at all. And there is *too little* support for this to be characterized as a conflict; the contradiction is not serious enough to imply the presence of errors in reasoning[24]. In this region, the problem is likely

---

[24] Heuristically, conflict occurs when activation at "+" and "-" collectors is similar and their sum is greater than 1.0. Since this is impossible in a probability framework, it implies an error in reasoning (Cohen, 1986). Lack of resolution is the case where activation at "+" and "-" collectors is similar, but the sum is approximately equal to 1.0. This corresponds to the situation where there is lack of detailed causal knowledge underlying an empirically observed frequency, but not a flaw in reasoning.

to be of a different kind: There is plausibly consistent information, but because it averages over relevant conditions, it does not sufficiently resolve the truth or falsity of the hypothesis. We will refer to this type of uncertainty as *lack of resolution*.[25]

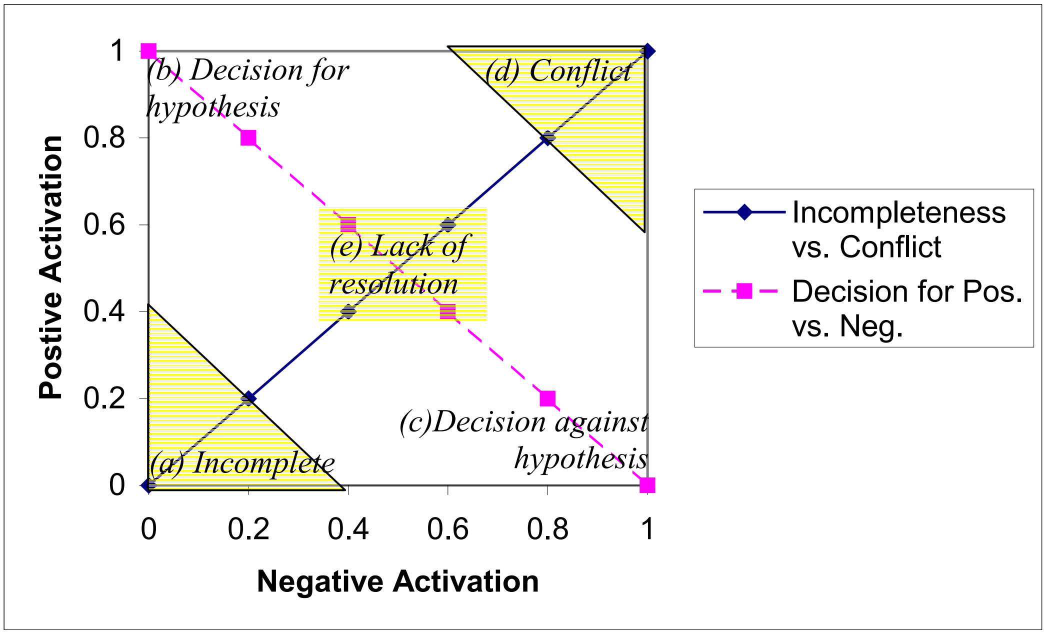

Figure 20. Two qualitatively different types of uncertainty, incompleteness (a) and conflict (d), defined in terms of relative activation levels at the positive and negative collectors of a hypothesis.

A fourth and final kind of local uncertainty can occur at the two remaining corners of the unit square (b or c). The truth or falsity of the hypothesis may be supported as the result of an explicit *assumption*. Assumptions can also be implicit, but only explicit assumptions, due to clamping, are local. The assumption is *explicit* (and local) when the activation levels of "+" and/or "-" collectors are being *clamped* (or persistently attended) by the metacognitive system at 1.0 or 0. The assumption is *implicit* when (a) the truth or falsity of the attended hypothesis actually depends on the truth or falsity of other hypotheses, but (b) the system fails to explicitly represent those other hypotheses, and (c) the system reasons *as if* the other hypotheses had specific truth values. When reasoning depends on statistical aggregations of causes, utilities, or action feasibility, as discussed in Chapter 12, implicit assumptions are involved. Implicit assumptions are not a purely local kind of uncertainty, since examination of local activation levels alone cannot distinguish them from reliable conclusions.

While incompleteness, conflict, and lack of resolution refer to the inability to discriminate different conclusions *now*, dependence on assumption represents temporal instability. The current activation levels of the "+" and "-" collectors reflect a definitive conclusion that might, despite the appearance of decisiveness, *change* in the light of

[25] Other ways of characterizing it would be as *high entropy*, or as *classical Bayesian uncertainty*.

further thinking or discovery of further evidence. Although implicit assumptions cannot be readily identified from local activation patterns, the very *possibility* of such hidden weaknesses in an argument is important. Limitations on attention mean that at any given time in an argument, there are some hypotheses that have thus far been accepted without question. The *appearance* of a decisive conclusion – either (0,1) or (1,0) in the unit square – should not necessarily be the end of critical thinking.

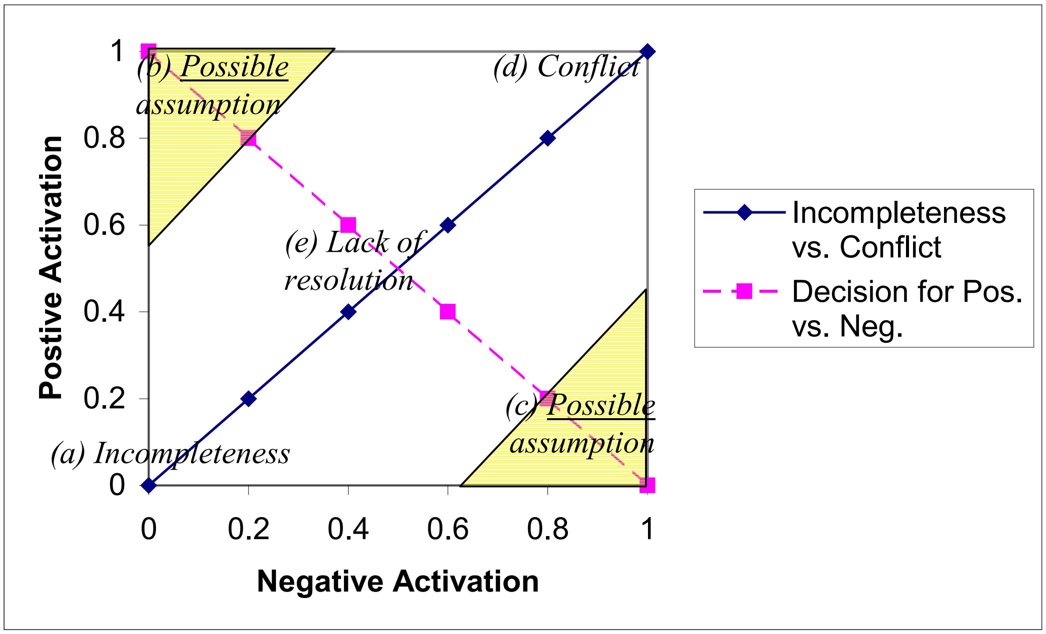

Figure 21. Apparently decisive conclusions may represent explicit or implicit assumptions

In sum, four local types of uncertainty can be identified (or, in the case of implicit assumptions, suspected) at an attended hypothesis:

1. *incompleteness* = insufficient information, i.e., support for *neither* the hypothesis nor its negation
2. *conflict* = contradictory information, i.e., support for *both* the hypothesis and its negation that is sufficiently high to suggest something is wrong with the reasoning
3. *lack of resolution* = aggregated information, which fails to discriminate conditions in which the hypothesis tends to be true from conditions in which the hypothesis tends to be false
4. *dependence on assumption*: (i) explicit local assumptions = intentional clamping of the truth value of a hypothesis as true or false (e.g., by persistent attention); (ii) implicit non-local assumptions = dependence of a conclusion on the truth or falsity of conditions that are not explicitly represented in the active belief network (e.g., reliance on statistical aggregations for variables at the edge of the currently active network).

In this section, we pursue the hypothesis that experienced decision makers may learn to recognize and use patterns of uncertainty like those just discussed. In this section we show that simple numerical measures can capture the degree to which the belief state at any hypothesis belongs to each of these patterns. We then show, by means of these measures, that these concepts as a set provide a plausibly exhaustive carving up of all possible types of local uncertainty.

<u>Incompleteness</u>

A gap in information at a node exists to the degree that both the “+” and - components of that node have zero activation. We refer to activation at the positive (“+”) collector of an attended hypothesis *a* as a^+ and activation at the negative (“-”) collector as a^-. Then,

$$Incompleteness = (1 - a^+)(1 - a^-)$$

Figure 22 shows how incompleteness varies as a joint function of a^+ (positive activation) and a^- (negative activation). Incompleteness is at a maximum when both kinds of activation are at 0, and at a minimum when at least one of them is at 1.0. Decision makers seeking to reduce incompleteness of information about a hypothesis would attend to parts of the inference network likely to help move them lower in this landscape.[26]

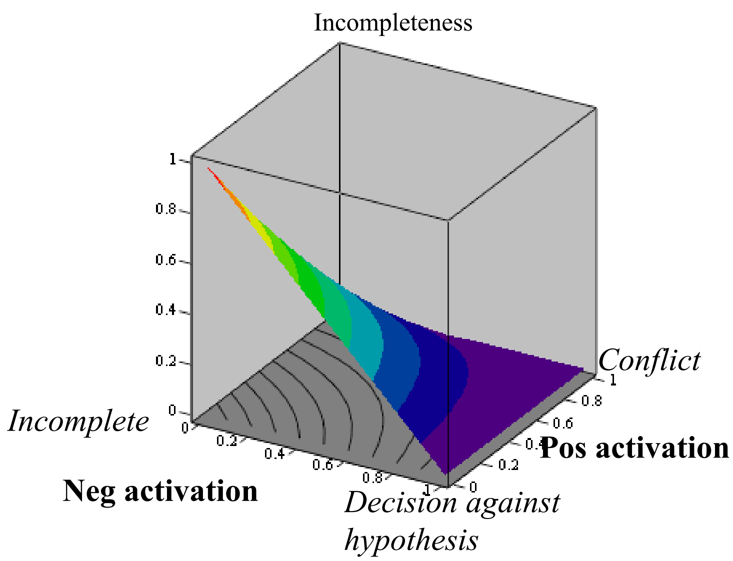

Figure 22. Local incompleteness (on the vertical axis) as a function of positive and negative activation. The unit square shown in Figure 20 is represented on the two horizontal (x,y) dimensions. The point labeled *incomplete* involves no activation for either positive or negative components.

---

[26] As can be seen from Figure 22, the slope of the descent is steepest at its maximum, near the point of total incompleteness (0,0), and tapers off as it gets closer to conflict (1,1). When one component (+ or -) has zero activation, incompleteness is a relatively steep, linear decreasing function of activation at the other component, until a decision point ((0,1) or (1,0)) is reached.

Conflict

A conflict in the evidence for hypothesis *a* exists to the degree that both the "+" and "-" collectors of *a* have maximum activation. Thus,

$$Conflict = (a^{+})\,(a^{-})$$

Figure 23 shows how conflict varies as a joint function of a^{+} and a^{-}. This surface is a mirror image of the surface representing incompleteness in Figure 22. Conflict is at a maximum when both kinds of activation are at 1, and at a minimum when at least one of them is at 0.[27]

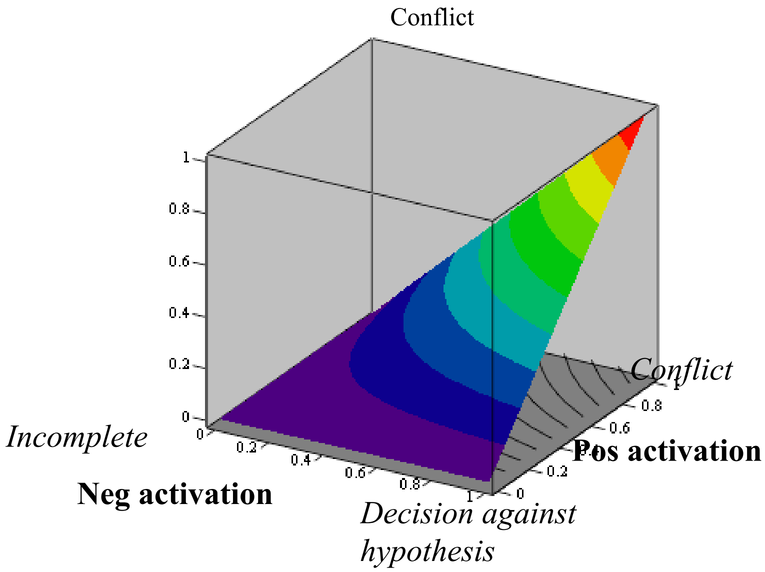

Figure 23. Local conflict (on the vertical dimension) as a function of negative and positive activation. The point labeled *conflict* involves full activation for both components.

Lack of Resolution

Lack of resolution for hypothesis *a* is at a maximum when the activation for both "+" and "-" collectors is .5. Thus,

$$Lack\ of\ resolution = (a^{+})\,(1 - a^{+}) + (1 - a^{-})\,(a^{-})$$

Figure 24 shows how lack of resolution varies as a joint function of a^{+} and a^{-}. Lack of resolution is at a maximum when both kinds of activation are at .5, and at a minimum when at least one of them is at 0 or 1.0.[28]

---

[27] As can be seen from Figure 23, the slope of the descent is steepest at its maximum, near the point of total conflict (1,1), and tapers off as it gets closer to incompleteness (0,0). When one component (+ or -) has activation of 1.0, conflict decreases steeply and linearly with decreasing activation at the other component, until a decision point ((0,1) or (1,0)) is reached.

[28] Unlike the previous two surfaces, the slope of the descent is relatively flat near its maximum (.5,.5) and is steepest at its four minima in the corners of the square. This suggests that lack of resolution gains in

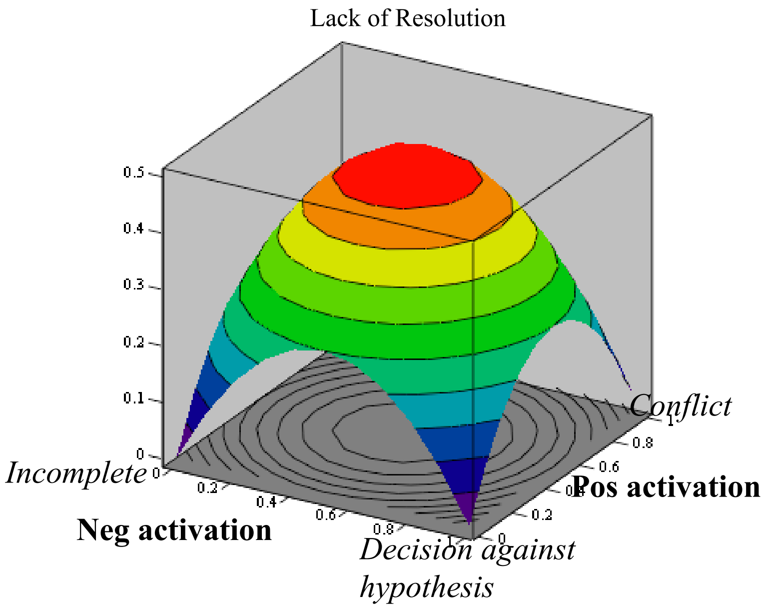

Figure 24. Lack of resolution (on the vertical dimension) as a function of negative and positive activation.

<u>Indecisiveness: Total Local Uncertainty</u>

A plausible overall measure of local uncertainty, which unifies the three measures, is based on the *difference* between the activation levels of the "+" and "-" components:

$$Decisiveness = (positive\ activation - negative\ activation)^2$$

The evidence is indecisive to the extent that there is no difference between activation of the "+" and "-" components. Thus,

$$Indecisiveness = 1 - (a^+ - a^-)^2$$

Figure 25 shows the level of indecisiveness as a function of activation of the "+" and "-" collectors, respectively. Indecisiveness is at a maximum when both positive and negative activation are the same (no matter at what level). Indecisiveness, therefore, imposes its maximum penalty wherever *any one* of the other three kinds of uncertainty (incompleteness, conflict, and lack of resolution) would impose a maximum penalty. Indecisiveness is at a minimum only when all the other three kinds of uncertainty are at a minimum: i.e., at the two corners where activation of one conclusion is 1.0 and the other zero. A decision maker using a unidimensional strategy of uncertainty reduction would aim to travel lower in this landscape, attempting to move *directly* to a conclusion in favor of either truth or falsity of the hypothesis.

---

importance as the decision maker approaches a conclusion, as we discuss later. When one component (+ or -) has activation of 0 or 1.0, lack of resolution is an inverted-u shaped function of activation in the other collector.

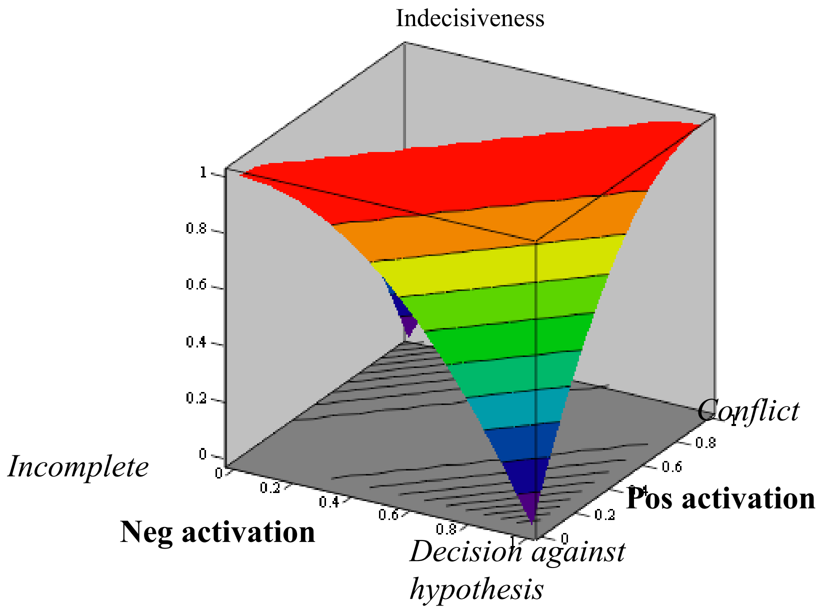

Figure 25. Indecisiveness (on the vertical dimension) as a function of positive and negative activation.

It turns out that indecisiveness can be decomposed precisely into the sum of incompleteness, conflict, and lack of resolution:

$$Indecisiveness = 1- (a^{+} - a^{-})^{2}$$

$$=1 - (a^{+})^{2} - (a^{-})^{2} + 2\,(a^{+})\,(a^{-})$$

$$= (1- a^{+})\,(1- a^{-}) + (a^{+})\,(a^{-}) + a^{+}\,(1 - a^{+}) + a^{-}\,(1 - a^{-})$$

$$= Incompleteness + Conflict + Lack\ of\ resolution$$

We will argue shortly that uncertainty handling strategies that address multiple dimensions of uncertainty are generally more likely to succeed than strategies based on a single aggregated measure like indecisiveness (or entropy). The reason for this advantage is that uncertainty handling in the real world does not calculate a conclusion (e.g., with Bayes rule) from a *static* belief network. Real-world uncertainty handling involves finding *new information* and introducing it to the argument: elaborating and modifying the active model of the situation, including the significance and diagnostic strength of cues, by retrieving, collecting, and reorganizing knowledge. As a result of the need to add information, the path to a firm conclusion is often indirect. For example, correcting conflict may lead to incompleteness; reducing the resulting incompleteness may increase lack of resolution; lack of resolution may be resolved by adopting explicit assumptions; and so on. A zigzag path of this kind may be the *only* way to activate all the knowledge that is necessary to arrive eventually at a consistent, stable answer.

Nevertheless, it is of interest that there exists such a simple mathematical relationship between the three component measures and the single overall measure of indecisiveness. This relationship provides support for the component measures *as a set* in addition to the arguments for each of them individually. Whether or not indecisiveness itself is used to *guide* the process of uncertainty handling, minimizing indecisiveness –

increasing the differentiation between the truth and the falsity of a hypothesis – is an intuitively positive measure of success.[29] Because of this face validity, the fact that incompleteness, conflict, and lack of resolution sum to indecisiveness supports the exhaustiveness of the three measures taken as a set. Strategies guided by these three measures would be predicted to counterbalance one another in such a way that, in the end, they minimize indecisiveness. As guides for uncertainty handling, however, they should do better: Because they explore the problem space more thoroughly, they are more likely to uncover implicit assumptions, to bring new knowledge to bear on a problem, and thus lead to decisions that are more stable in time.

Comparison with Bayesian Modeling

In a Bayesian model, the probability of a hypothesis is constrained to be one minus the probability of its negation. Thus, if we were to equate activation levels to Bayesian probabilities, we would get the following:

$$Incompleteness = (1 - a^{+}) (1 - a^{-}) = (1\text{-}Prob_a) Prob_a$$

$$Conflict = (a^{+}) (a^{-}) = Prob_a (1\text{-}Prob_a)$$

$$Lack\ of\ resolution = a^{+} (1 - a^{+}) + a^{-} (1 - a^{-}) = 2\ Prob_a (1\text{-}Prob_a)$$

$$Indecisiveness = 1\text{-} (a^{+} - a^{-})^2 = 1 - (Prob_a - (1 - Prob_a))^2 = 4\ Prob_a (1\text{-}Prob_a)$$

It is apparent that incompleteness, conflict, lack of resolution, and indecisiveness cannot be distinguished within a classical Bayesian framework! All four are reducible to a single measure: the chance of the event times the chance of its complement. Figure 26 shows what this single Bayesian measure (scaled so that its maximum equals 1.0) looks like as a function of the activation levels for "+" and "-" components. It can be seen that for the permissible combinations of activation levels (i.e., whose sum is 1.0), the resolution measure in Figure 26 is the same as the indecisiveness measure in Figure 25.

Bayesian models are unable to accommodate separate measures of incompleteness, conflict, and lack of resolution, and the rich repertoire of uncertainty handling strategies that this differentiation supports (for more discussion of these issues, see Cohen, 1986; Cohen, Schum, Freeling, & Chinnis, 1984). To support a more effective array of decision making strategies, an inferential syntax is required in which activation of "+" and "-" components are free to vary independently between 0 and 1.0.

---

[29] In addition, as we shall see in the next section, indecisiveness corresponds to the *only* Bayesian measure of local uncertainty.

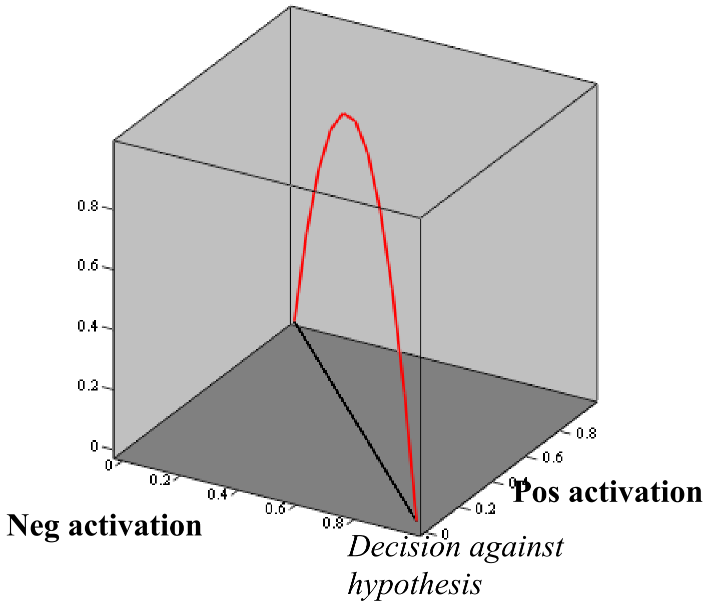

Figure 26. Bayesian measure of uncertainty: *1 - resolution*. Negative activation always equals one minus positive activation. (Vertical axis values are multiplied by 4 to make scale comparable to other figures.)

Dependence on Assumptions

Assumptions are introduced because they generate useful effects in other parts of a belief network, e.g., assuming a worst case scenario helps in evaluating an option; assuming that the enemy has a particular intent helps in predicting actions the enemy might take and how we might respond. The downside of assumptions is, of course, that they may turn out to be wrong. We may sometimes learn whether or not an assumption is false by directly confirming or disconfirming it. But many times the only method available is to look again elsewhere in the network, at the conclusions the assumption helps us to arrive a. If evidence is found that invalidates any of those conclusions, we may have to reconsider the assumption. As a result, the burden in evaluating an assumption is not local to a particular hypothesis, but will involve looking at its effects on uncertainty (especially conflict) in as large a portion of the belief model as possible.

From the local point of view, an explicit assumption simply amplifies or reduces the *decisiveness* of a node (without changing the *direction* of support for a^+ versus a^-). A hypothesis depends on an *explicit* assumption if the decision maker has reflectively clamped activation levels of positive and/or negative collectors so as to change the *difference* in activation between them. If a_b^+ and a_b^- are the baseline reflexive activation levels that the "+" and "-" collectors would have had without clamping, and a^+ and a are the new levels of the "+" and "-" collectors after the assumption is made, then

$$\textit{Dependence on explicit assumption} = (a^+ - a^-)^2 - (a_b^+ - a_b^-)^2$$

as long as the sign of the two differences $(a^+ - a^-)$ and $(a_b^+ - a_b^-)$ is the same. Recalling that we defined *decisiveness = (positive activation –negative activation)^2*, it follows that this measure of dependence on an explicit assumption is equivalent to the *change in decisiveness* brought about by clamping or persistent attention.

Dependence on an assumption is a measures of *epistemic risk* if the decision maker has increased the difference in activation between positive and negative collectors, or *epistemic caution* if the decision maker has reduced the difference in activation. A simple two-dimensional representation will illustrate some key points, represented by numbers 1 through 3 in Figure 27:

(1) In this example, the initial belief state before any assumptions are made is represented by the point at (.85,.45). If an assumption were to move the belief state along the diagonal drawn though this point, it would not change decisiveness, since the difference between positive and negative activation remains constant on any line of slope = 1. Therefore, assumptions that merely trade off incompleteness, conflict, and lack of resolution, are *epistemically neutral*, and do not have any effect on decisiveness or on our measure of *dependence* on assumptions.

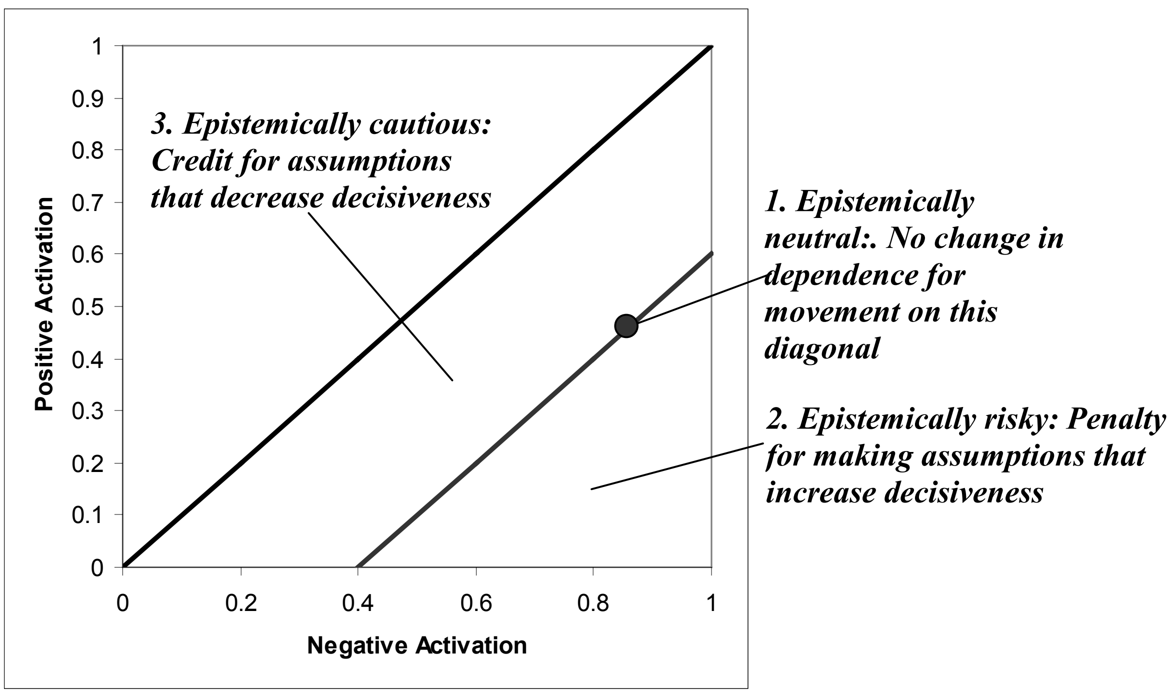

Figure 27. Numbers correspond to different effects on the measure of dependence on assumptions as a function of the starting point (black dot t .85,.45) and the of assumptions of different kinds relative the prior state of belief (black dot a).

(2) There is a penalty for adopting an assumption to the extent that it goes beyond the current evidence and *adds* to decisiveness, that is, to the extent that it is *epistemically risky*. Therefore, any assumption that moves the belief state closer to a conclusion, gains decisiveness at the cost of increased dependence on assumptions. The maximum cost

always occurs at the two corners, (0,1) and (1,0), but the largest possible penalty depends on how far it is necessary to move from the starting point, $a_b^+ - a_b$.[30]

(3) If the assumption is used to *reduce* decisiveness, that is to discount or dilute existing evidence, the assumption is *epistemically cautious*, and there is a local *credit* rather than a debit. Assumptions that move belief toward neutrality (the two collectors' having equal values), have a negative dependence score. The maximum dependence credit occurs for assumptions on the (0,0) to (1,1) diagonal, but the amount of the credit depends on how far it is necessary to move from the starting point, $a_b^+ - a_b$.[31]

At the local level, all an assumption does is transform uncertainty from one form (indecisiveness) to another (dependence on assumption) with no change in the total. Any increase (decrease) in dependence on assumptions is exactly counterbalanced by an opposite decrease (increase) in indecisiveness, i.e., in the sum of incompleteness, conflict, and/or lack of resolution. This invariance can be described as the *conservation of total local uncertainty*, given the current evidence and beliefs.

$$
\begin{aligned}
&= \textit{Indecisiveness before assumption} \\
&= 1 - (a_b^+ - a_b^{--})^2 \\
&= 1 - (a^+ - a^{--})^2 + (a^+ - a^-)^2 - (a_b^+ - a_b^{--})^2 \\
&= \textit{Indecisiveness after assumption} + \textit{Dependence on assumptions}
\end{aligned}
$$

Figure 28 shows how dependence on an assumption varies as a joint function of a^+ and a^- when $a_b^+ - a_b^- = -.4$, that is, for the example of Figure 27, in which a decision maker is leaning slightly toward the falsity of the hypothesis before adopting an assumption. For contrast, Figure 29 shows dependence on an assumption when $a_b^+ - a_b^- = -.8$, that is, the evidence points more strongly toward the falsity of the hypothesis. These situations differ both in the maximum penalty for epistemic risk, and in the maximum gain for epistemic caution. In both cases, however, the total range of the measure (difference between maximum and minimum) is 1.0.

The measure of dependence on explicit assumptions requires that a decision maker compare the local +/- activation levels before and after clamping. Hence, while the measure is indeed locally available at the attended cluster (it does not require direct examination of any other parts of the belief network), it requires more time and/or memory to assess than the other local measures we have considered. As a result, we might expect greater difficulty in gauging and keeping track of assumptions, even explicit ones, than in the case of other types of uncertainty. Of course, ferreting out *implicit* assumptions is a far more challenging task.

---

[30] When the initial difference between collectors is only .4 (as in Figure 27), clamping one collector at 1.0 and the other at 0 has a cost of $1 - .4^2 = .84$. When the initial difference is .8, there is less need for boldness, and a smaller cost is incurred, $1 - .8^2 = .36$.

[31] When the initial difference is.4, there is not much conviction to surrender, and clamping both collectors at the same value brings a credit of $0 - .4^2 = .16$. When the initial difference is .8, there is more belief to give up, and a larger credit is gained by discounting the evidence, $0 - .8^2 = .64$.

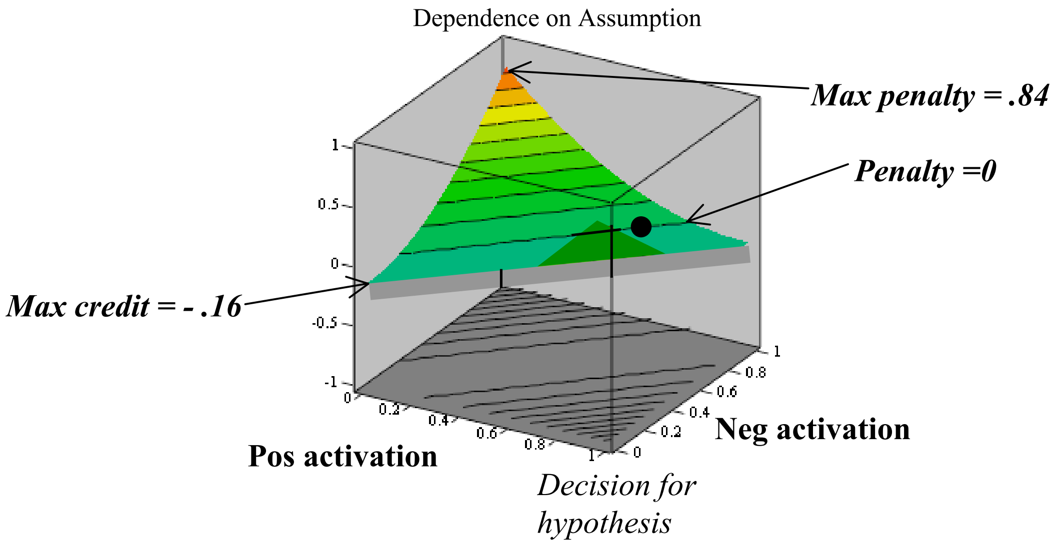

Figure 28. Dependence on assumptions, with the original difference between activation at collectors = .4.

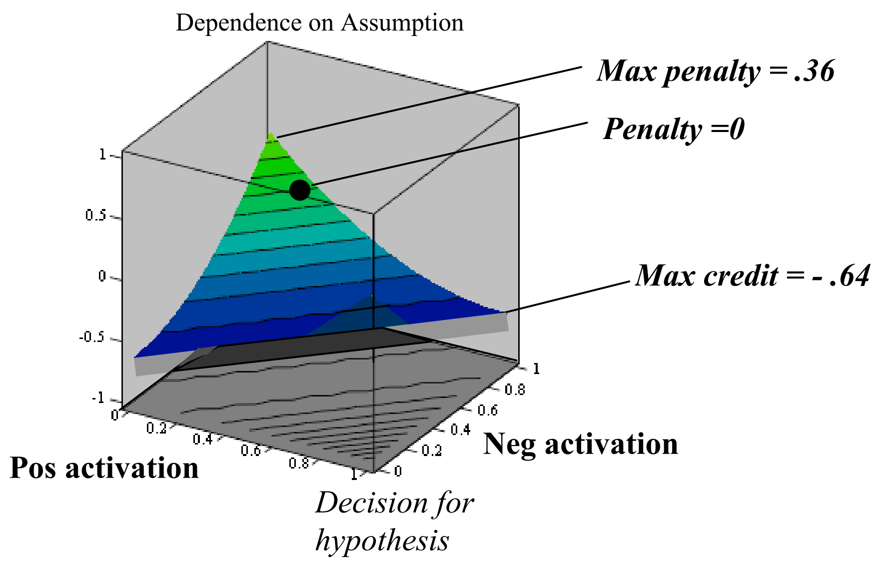

Figure 29. Dependence on assumptions, with the original difference between activation at collectors = .8.

Visualization of Local Uncertainty

Local measures of uncertainty have a dual role: First, they reflect qualitatively different decision subproblems that can be most effectively addressed by different uncertainty-handling tactics. Secondly, they are substitutable components of a single overall measure of local uncertainty, the indecisiveness of the conclusion regarding that hypothesis. It would be useful to have a relatively simple graphical representation of local uncertainty that captures both of these aspects. Such a representation should vividly convey the *total uncertainty*, but also allow an immediate determination of its *qualitative* nature: whether it is determined most by incompleteness, lack of resolution, or conflict, and whether any assumptions have been adopted. In this section, we explore the possibility of such a visualization, and show that the relationships among the various measures make possible a relatively simple solution.

We will first consider a representation of total uncertainty, then the relative contribution of different component types.

Total Uncertainty

The belief state represented by the point at (.85,.45) in Figure 30 has total uncertainty equal to $1 - .4^2 = .84$. Every other point on the same diagonal has the same total uncertainty, because the difference between positive and negative collectors is constant along that diagonal. As a result, the size of the shaded triangle in Figure 30 can be used to gauge the total uncertainty of this set of belief states. Moreover, every belief state in the unit square determines a similar triangle, whose size reflects its own total uncertainty.[32]

In Figure 31, triangles have been generated corresponding to different levels of indecisiveness. Note that the topmost triangle is the same as the one shown in Figure 30 (both have a difference in activation of .4, hence, an indecisiveness of .84). The two remaining triangles represent approximately equal decreases in indecisiveness, to .51 and .19, respectively. As the difference in activation increases, the size of the triangles shrinks. Figure 32 is a top-down view, which shows the contour lines of Figure 31 projected onto the x-y plane. Each of the contour lines represents a different level of indecisiveness, and each of them determines a different sized triangle. Figure 33 represents indecisiveness as a function of activation levels in the negative and positive collectors.

[32] Each triangle is formed by drawing a diagonal with slope = 1 through the point representing the belief state (negative and positive activation levels). The desired triangle encloses the area lying between the diagonal and the nearest conclusion, (0,1) or (1,0).

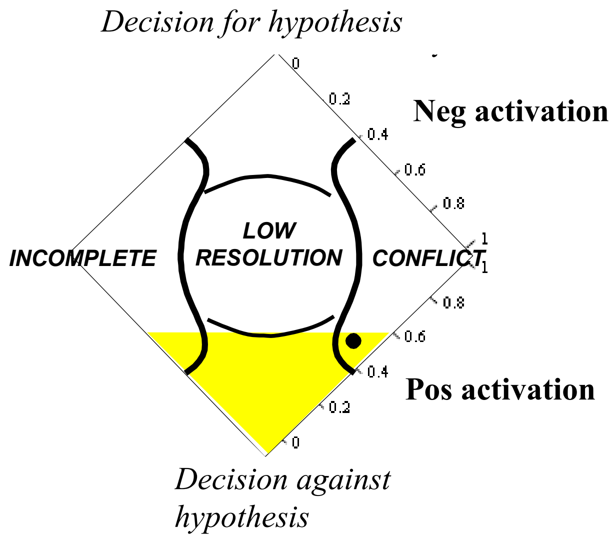

Figure 30. Graphical representation of zones in the unit square where different types of uncertainty predominate. The illustrative point represents .85 negative and.45 positive activation, and the size of the shaded triangle corresponds to the total amount of uncertainty at that point. The shaded area is the triangle determined by the belief state corresponding to the indicated point. Its size is related to the total uncertainty of that belief state.

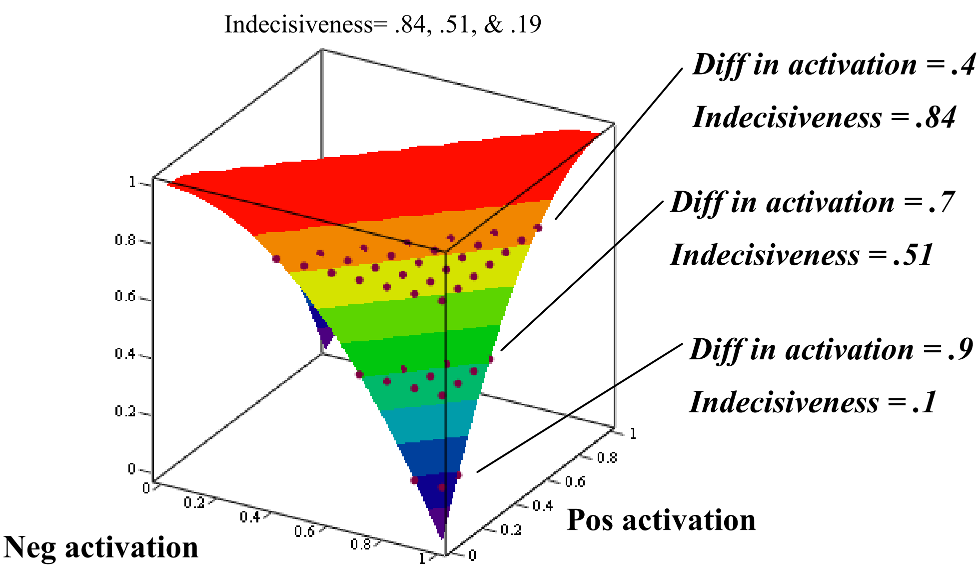

Figure 31. Indecisiveness as a function of positive and negative activation. Triangles are located at three different levels of indecisiveness, corresponding to three differences in negative and positive activation.

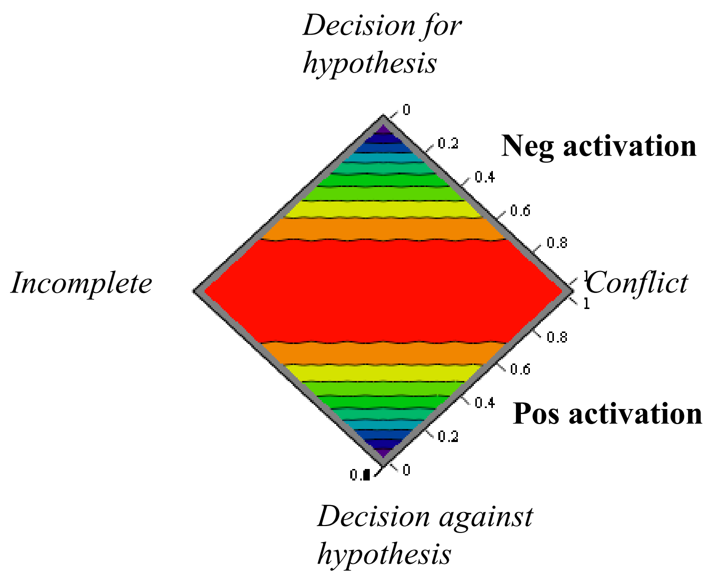

Figure 32. A top-down view of Figure 31, showing contour lines projected onto the x-y plane.

As can be seen from the spacing of the contour lines in Figure 32, the area of the triangle is not a linear function of indecisiveness. Each contour line represents an equal change in indecisiveness. The areas between contour lines represent the areas added to the triangle by equal increments of indecisiveness. It is clear that the area of the triangles increases more rapidly with indecisiveness at high levels of indecisiveness (toward the top of the surface in Figure 31 and the corresponding central diagonal of Figure 32). The triangle area metric emphasize differences among more uncertain states, where the need to reduce uncertainty is greatest. On the other hand, the length of the base of the triangle provides an untransformed linear representation of the *difference* in activation levels between the two collectors.

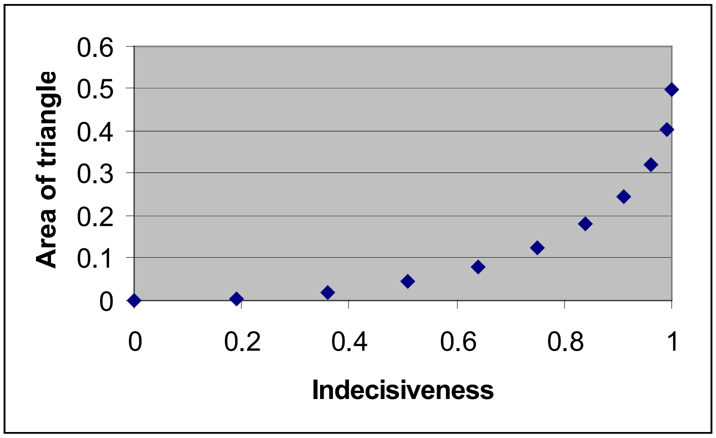

Figure 33. Area of triangle as a function of total uncertainty, or indecisiveness.

Types of Uncertainty

The same visual context can convey the qualitative *type* of uncertainty that is most crucial in any particular belief state. As a starting point, Figure 34 plots the level of the specific type of uncertainty that predominates at each combination of positive and negative activation levels. Figure 35 shows the resulting contour levels projected onto the x-y plane, and shows the regions in which each type of uncertainty predominates. Figure 30 uses these results to carve up the unit square into three zones. Conflict and incompleteness are most important in the respective corners, (0,0) and (1,1). Lack of resolution prevails in the center of the square, along the (0,1) to (1,) axis. It is immediately apparent that the predominant type of uncertainty at the point used in this example (.85,.45) is conflict.

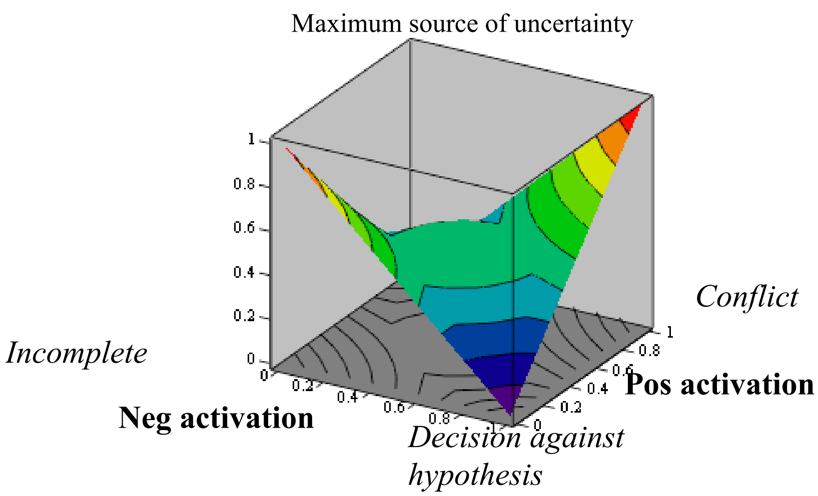

Figure 34. Plot of the maximum of conflict, incompleteness, and low resolution measures at each combination of positive and negative activation.

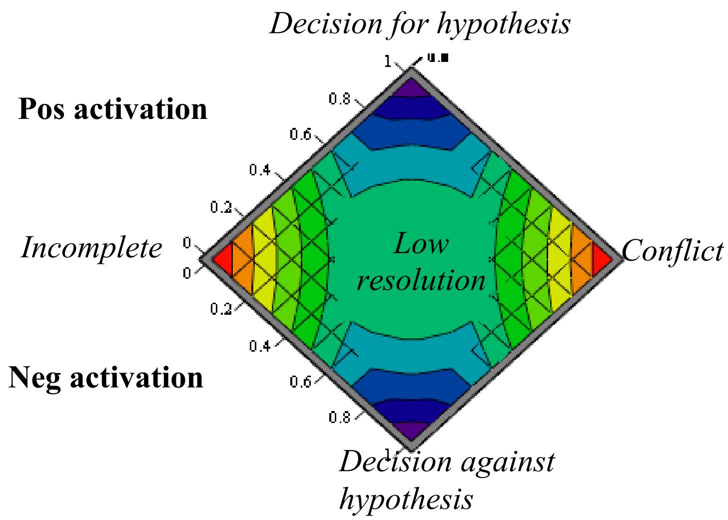

Figure 35. Contour lines of maximum type of uncertainty in Figure 34. Cross-hatched areas show where incompleteness (on the left) and conflict (on the right) contribute the most to total uncertainty. In center, low resolution is the predominate uncertainty type.

The effect of assumptions on local uncertainty can be represented in the same context. Figure 36 shows two types of assumptions a decision maker might adopt from the same original point (.85,.45). In Case 1, the decision maker has clamped activation levels at (.9,0). This increases decisiveness in the direction of the original evidence (primarily by reducing conflict). It brings with it an equivalent dependence on assumptions, as indicated by the red arrow. The length of the arrow is the distance from the point representing the original belief state (.85,.45) to the new point (.9,0) *measured along the decisiveness axis*. This length reflects the degree of dependence that the assumption introduces as well as the corresponding increase in decisiveness.[33] The size of the triangle, however, remains the same, since improved decisiveness is exactly offset by dependence on the new assumption.

The other kind of assumption represented in Figure 36 reduces decisiveness. Case 2 reflects a more cautious stance toward the evidence, and thus receive a credit for negative dependence on assumptions. Arrows for assumptions that decrease decisiveness are aimed inward toward the center diagonal rather than outward toward the conclusions.

[33] More precisely, the length of the arrow equals $[(a^+ - a^-) - (a_b^+ - a_b^-)] / \sqrt{2}$, which is the square root of one half of our measure of dependence on assumptions. This is the perpendicular distance from the diagonal on which the original point sat (the boundary of the triangle) to the diagonal on which the new point sits.

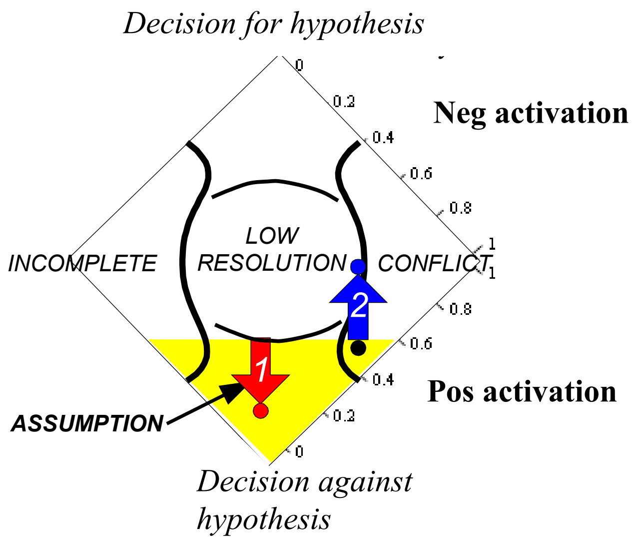

Figure 36. Two different types of assumptions starting from the same original belief state, represented by the black dot. Assumption that incurs dependence and increases decisiveness is shown in red; assumption that decreases decisiveness and generates a negative dependence is shown in blue. The length of the arrow represents the magnitude of the positive or negative dependence and the corresponding change in decisiveness.

Uncertainty Handling Strategies

Given that a decision maker notices a particular type of uncertainty at an attended node, what does she do about it? In this section, we will describe two broad strategies: (1) A domain specific strategy (which nevertheless has some very important general components), based on learned correlations between uncertainty at specific nodes and activation at other nodes. (2) A more general strategy based on distinctions between evidence and conclusions in a particular situation.

Domain-Specific Strategies: Culpability

Through experience in a domain, decision makers may learn not only to recognize different types of uncertainty at a single hypothesis, but how the rest of the active belief network influences that uncertainty. We discuss measures of the responsibility, or *culpability*, borne by other hypotheses in working memory for the local uncertainty in an attended node. These measures capture the potential for new information at the "culpable" node to change the local uncertainty at the attended node. They show where a decision maker may usually find new information that fills a gap, reduces conflict, or confirms or denies an assumption at the hypotheses of interest. The measures of culpability thus give the metacognitive system promising leads about where to shift its attention next. They prioritize attention in terms of where more digging might have the biggest payoff.

The culpability of some hypothesis a_j for local uncertainty at the hypothesis of interest, a_i ($j \neq i$) is based on the answers to three questions:

1. Domain-specific learning: How much influence can changes in the activation at a_j have on activation of positive and negative collectors at a_i ?
2. General learning: What is the effect of activation of positive and negative collectors at any node (in this case, a_i) on the relevant type of local uncertainty at that node?
3. Domain-specific learning: How likely are changes in activation of a_j to occur – that is, what is the expected total amount of local uncertainty of a_j itself?

The first and third of these components can only be learned through specific experience in a domain, while the third (the recognition of patterns of uncertainty) is general across domains.

The following formulas represent the first two components of culpability, showing how the domain-specific and general elements are combined mathematically. Each formula captures the *impact* of the positive collector of a_j on a particular kind of uncertainty at hypothesis a_i. This is broken down (via the chain rule) into elements that correspond to (1) and (2) above. Formulas for the impact of the negative collector of a_j can be derived in the same way.[34]

Impact of a_j^+ *on incompleteness of* a_i

$= \partial$ incompleteness of $a_i / \partial a_j^+$

$= [\partial (1 - a_i^+) (1 - a_i^-) / \partial a_i^+] (\partial a_i^+ / \partial a_j^+) + [\partial (1 - a_i^+) (1 - a_i^-) / \partial a_i^-] (\partial a_i^- / \partial a_j^+)$

$= - (1-a_i^-) (\partial a_i^+ / \partial a_j^+) - (1-a_i^+) (\partial a_i^- / \partial a_j^+)$

Impact of a_j^+ *on conflict of* a_i

$= \partial$ conflict at $a_i / \partial a_j^+$

$= \partial (a_i^+ a_i^-) / \partial a_i^+) (\partial a_i^+ / \partial a_j^+) + (\partial (a_i^+ a_i^-) / \partial a_i^-) (\partial a_i^- / \partial a_j^+)$

$= a_i^- (\partial a_i^+ / \partial a_j^+) + a_i^+ (\partial a_i^- / \partial a_j^+)$

Impact of a_j^+ *on low resolution of* a_i

$= \partial$ lack of resolution of $a_i / \partial a_j^+$

$= \partial ((a_i^+)(1-a_i^+) + (1-a_i^-)(a_i)) / \partial a_i^+) (\partial a_i^+ / \partial a_j^+) + \partial ((a_i^+)(1-a_i^+) + (1-a_i)(a_i)) / \partial a_i) (\partial a_i^- / \partial a_j^+)$

$= (1-2a_i^+) (\partial a_i^+ / \partial a_j^+) + (1-2a_i^-) (\partial a_i^- / \partial a_j^+)$

---

[34] In each case, the relevant formula is: (Impact of a_j^+ on the uncertainty of a_i) = (Impact of a_j^+ specifically on a_i^+) (Impact of positive activation on uncertainty at any node)+ (Impact of a_j^+ specifically on of a_i^-) (Impact of negative activation on uncertainty at any node)

These formulae require the acquisition of two kinds of information by associative learning: (1) Decision makers must learn the general impact of positive and negative activation on each of the relevant types of uncertainty of a node. (2) Decision makers must learn the domain-specific correlations between (a) changes in "+" and "-" activation at some nodes and (b) changes in "+" and "-" activation at other nodes in the network. The latter, domain-specific derivatives are precisely the same as measures already utilized in associative learning through backpropagation.[35]

A reasonable culpability measure combines these measures in the following way:

> *Culpability of a_j for local uncertainty of a_i*
>
> *= (Average Indecisiveness of a_j) Max {(Impact of a_j^+ on uncertainty of a_i), (Impact of a_j^- on uncertainty of a_i)}*

If the decision maker is trying to resolve uncertainty at more than one attended hypothesis, e.g., $a_1 \ldots a_n$, the overall culpability attributed to a_j is simply the sum of its culpability for each of the attended hypotheses (other than itself):

> *Culpability of a_j = $\Sigma_{i \neq j}$ Culpability of a_j for local uncertainty of a_i*

The summation increases the chance that high leverage hypotheses will be selected for attention, so that problems of uncertainty at more than one critical hypothesis might be resolved all at once. An example would be an assumption shared by more than one problematic argument.

We conclude that this type of metacognitive skill appears to be in part general, and in part domain-specific. Unlike logic or decision theory, it cannot be practiced in the absence of some level of prior knowledge in a domain. Nevertheless, learning in one application may well hasten the acquisition of comparable skill in another through transfer of the general components.

General Strategies: Argument Roles

Mental models represent persisting, general knowledge about the relationships (which are often causal) among events in the domain. We saw in Chapter 12 how Shruti can be used to represent and reason reflexively with such knowledge, and in this Chapter how simple metacognitive learning processes can lead to improvements in reasoning. However, a limitation of the reflective skills discussed in the previous section is their domain-specificity. Patterns of uncertainty are general, but the culpability relationships that enable decision makers to search efficiently for solutions must be learned, to some degree at least, anew in each domain.

---

[35] See, for example, Rumelhart, Hinton, and Williams (1986). Backpropagation is a way of incrementally adjusting the weights on links in a belief network across many trials so that the network learns to match its output to the "correct" targets. Weights are adjusted in proportion to the size of the errors at the output node and their responsibility for the errors. Errors occur when the output node's activation is too high or too low relative to a teaching input, or target. Degree of responsibility is measured by the derivative of the error with respect to changes in the weight. In a simple chain of connections, this is simply the product of the weights on the links.

We now explore the hypothesis that there are more purely general reflective skills. In particular, we hypothesize that decision makers can learn to identify general patterns in a activated belief network that are diagnostic of strong versus weak culpability. In sum, decision makers may learn generalizable patterns that indicate which hypotheses are likely to be having the most impact on uncertainty in other hypotheses.

Identifying argument roles.

The key step in the acquisition of general reflective skill is an awareness of *argument* relationships. An argument, to borrow a part of Toulmin's (1958) definition (see Chapter 4), consists of:

1. a conclusion
2. grounds for that conclusion
3. rebuttals, i.e., conditions under which the link between grounds and conclusion would not hold

Argument networks are built by *reflection* on recognitionally activated mental models. Argument networks represent temporary, situation-specific awareness of the evidence-conclusion relationships among nodes in the activated part of long-term memory. That is, they represent decision makers' understanding of the epistemic priority of their beliefs on a given occasion, i.e., what beliefs function as evidence and lead to what other beliefs that function as conclusions *in the current situation*.[36] An argument network can be thought of parsimoniously as the set of attended beliefs *labeled* according to their current function as evidence, conclusion, or rebuttals. Figure 37 presents part of an argument network. This argument is based on (but is not the same as) a causal mental model of *enemy intent*.

Once decision makers become familiar with the roles that hypotheses play in arguments,[37] they can learn to identify the roles that hypotheses play in specific arguments by reflection on the activated part of long-term memory.[38] We do not assume that decision makers are able to directly recognize causal relationships among their own beliefs (i.e., that they came to believe X because they learned Y). We do assume (i) that they can identify the hypothesis or set of hypotheses that is currently of primary interest (e.g., will the enemy attack in the south?), and (ii) that they can recognize and remember the sources of their beliefs. In particular, they can identify *dynamic* and *episodic* facts: information about specific events that is received through the senses or stored in episodic

---

[36] Arguments thus provide causal explanations of why decision makers hold the beliefs that they hold. They reflect somewhat idealized causal relations among mental events of believing, not necessarily causal relations among the external events that those beliefs represent.
[37] There is evidence in the cognitive development literature (e.g., Kuhn, Amsel, & O'Loughlin, 1988; King & Kitchener, 1994) that the ability to discriminate evidence from conclusions evolves over time.
[38] This task is quite different from the one typically described as "argument analysis" in the literature on critical thinking (e.g., Ennis, 1996). In "argument analysis," the issue is to identify the argument of an author in a written or spoken passage, and cues such as words like "because" or "therefore" are helpful in distinguishing grounds from conclusions. In the case of meta-recognition, decision makers must be able to identify the reasons for their own conclusions on the fly.

memory. We also assume the basic set of metacognitive operations: that decision makers can query and clamp hypotheses.

1. The first step in building an argument network is identification of a set of nodes in the mental model as the *primary hypotheses*, or issues of main concern to the decision maker. These are automatically labeled as *conclusions* for the arguments in which they participate. For example, in Figure 37, the primary hypothesis might be P001, that the enemy's intent is to attack in the south.
2. The next (or concurrent) step in building an argument network is identification of events that are known or observed on this occasion (i.e., *dynamic or episodic facts*), and which influence the strength of the primary hypotheses. These are labeled as *grounds* for the arguments in which they participate.
3. The third step is to refine the understanding of argument relationships, if possible, by identifying *intermediate* grounds and conclusions among the remaining activated hypotheses. This could happen in a variety of ways. One is by recalling temporal relations among beliefs, e.g., that I came to believe X after I saw Y, or after Tom told me X was the case. In addition, or when temporal cues are not available, the decision maker draws on explicit causal knowledge that specifies which types of events generally cause what other types of events. Another, more active strategy is a combination of querying and clamping conclusions in order to discover the causal relationships that lead from conclusions to grounds.[39]
4. The fourth step (which may happen concurrently with the third) is to identify nodes that are on the *edge* of the activated part of the mental model, but are neither dynamic/episodic facts nor conclusions. These are hypotheses that are evidentially relevant to the conclusions of interest, but have not been considered deeply. The decision maker has simply accepted them as given. The strength of these hypotheses is determined either by statistical aggregation of past experience (*taxon facts, J-facts, feasibility facts*) or by explicit assumptions by the decision maker. These nodes are labeled as

---

[39] Decision makers with sophisticated metacognitive skills should be excellent at eliciting their own *implicit* knowledge of chains of argument, starting with a known conclusion and known grounds. For an underlying causal mental model, either the grounds must be a cause and the conclusion a predicted consequence, or the grounds must be an effect from which the conclusion is inferred as its cause. (More complicated cases, such as predicting one effect from another effect of the same cause, can be understood as combinations of these two possibilities.) To identify the intermediate grounds and conclusions of a causal argument chain, decision makers could start by querying the conclusion (i.e., asking, *What could have caused this to happen?*), querying the results of that query, and so on. The strategy is successful if it eventually links up with the grounds (which in this case turns out to be a cause, and the conclusion a prediction). If this fails, decision makers might try clamping the conclusion as true (asking, *What does this event lead to?*), clamping the results as true, and so on, until linking up with the grounds. (In this case, the grounds turns out to be an effect and the conclusion a cause.) If both of these strategies work, some of the grounds are causes (see P002 in Figure 37) and some are effects (see P003) of the conclusion. If neither of these strategies works, then the conclusion may not be argumentatively linked to the alleged grounds at all. There is a serious gap in the argument!

assumptions (or their negations are labeled as rebuttals) for the arguments in which they participate. Activation of additional parts of long-term memory may prove them wrong, and neutralize the current arguments for the primary hypotheses.

Figure 37 illustrates the result of using such reflective strategies to identify argument relationships in a causal mental model. The argument links run parallel to the causal relationships, but not necessarily in the same direction. In particular, one grounds for believing in the conclusion about enemy intent (P001) refers to a *cause* (P002), and the grounds of the other argument for that conclusion (P001) refers to an *effect* (P003). The negation of the rebuttal (P004) is part of the conditions, along with P001, that would lead to movement of engineers (P003). Since there is as yet no direct support for or against it, aside from general probabilities, it is shown as a rebuttal to the argument for P001 based on P003.

Use of argument structure to find culpable hypotheses.

The matching of hypotheses to roles in an argument network helps decision makers handle uncertainty reflectively in domains where they do not have expert knowledge. First, it breaks the problem down into component parts, or subproblems. These subproblems involve decisions regarding the truth or falsity of one or more hypotheses (conclusions) by reference to the truth or falsity of other propositions (grounds), conditional on assumptions about other factors (rebuttals). More importantly, when evaluating the truth of a conclusion, the argument structure contains important clues about where to look to resolve problems with different types of local uncertainty. In fact, the most likely causes of each type of local uncertainty in a hypothesis of interest can be readily identified in terms of argument structure. As a result, decision makers can leverage learning about argument structure into more general strategies for resolving uncertainty.

General strategies for incompleteness.

Figure 38 illustrates three different features of the argument *context* that can cause a hypothesis to be *incomplete*:

1. There is no activation for +P001 from argument A001 because there is no evidence for the grounds of that argument, i.e., no reason yet to believe P002.
2. There is no activation for +P001 from argument A002 even though there is evidence for the grounds of that argument (+P003). But there is also evidence for the rebuttal (+P004), and this rebuttal neutralizes the support from argument A002.
3. There is no activation for -P001 because there is no argument that could support the conclusion that P001 is false.
4. There are no additional arguments for +P001.

To resolve the incompleteness at P001, the decision maker needs to solve at least one of these problems:

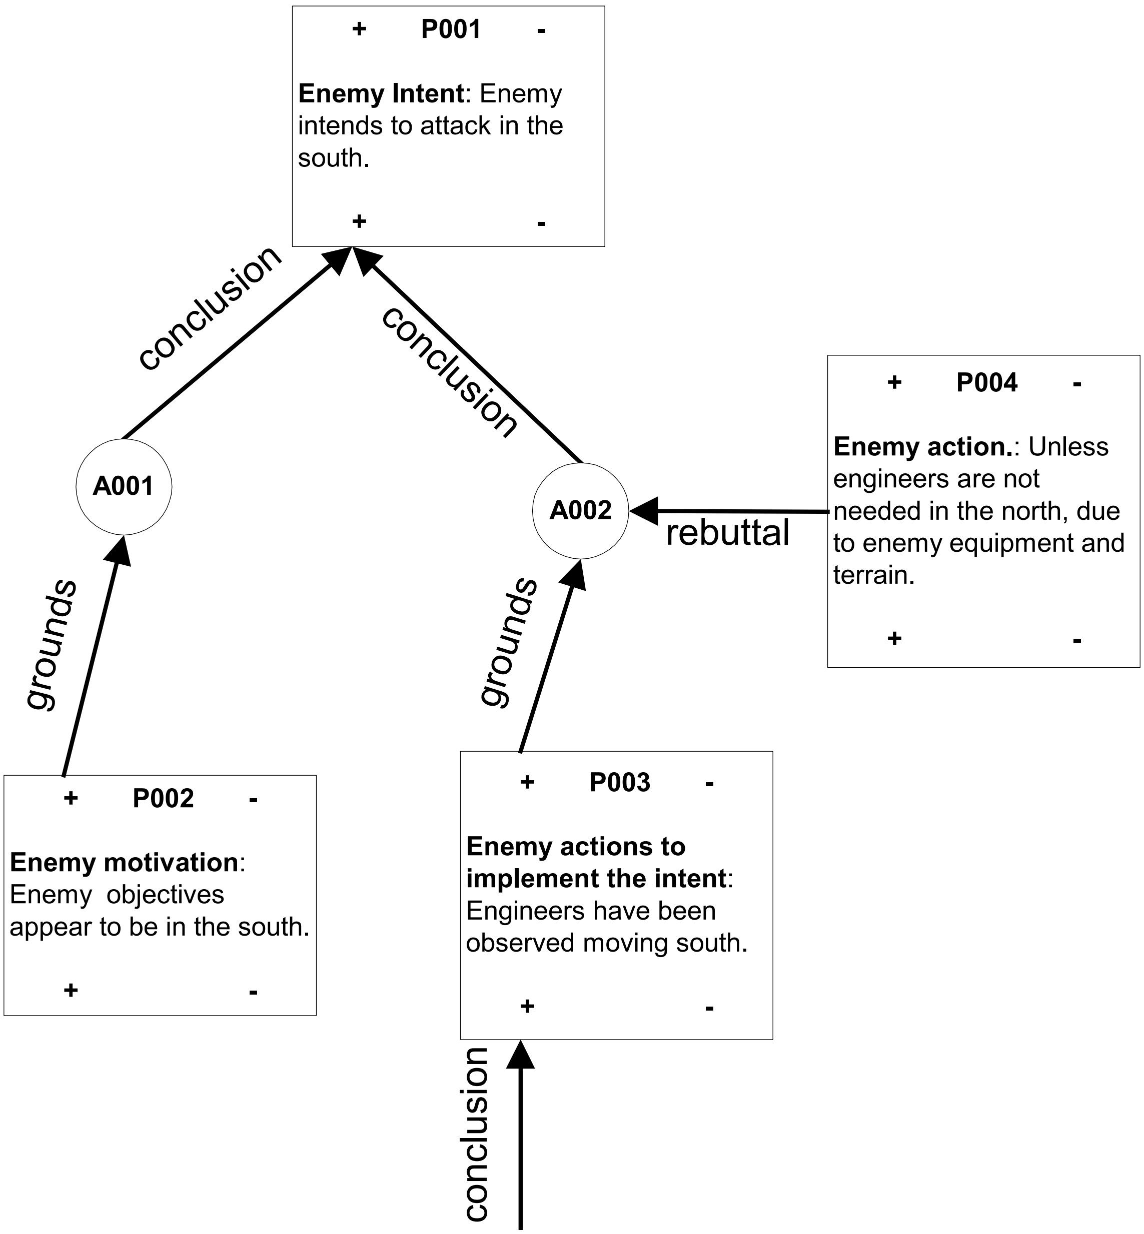

Figure 37. An argument network based on a causal mental model. P002 is a cause of P001, while P003 is and effect, yet both may be grounds for believing P001.

(1) In order to activate information in long-term memory that might support P002, the decision maker can shift attention there. The result is a query for possible antecedents of P002. In the case of a causal mental model, the decision maker can clamp –P002 true, and see if this leads to conflict with activation from some event of which P002 is the best explanation. If these tactics work, P001 will receive activation from P002, and incompleteness will be resolved.

(2) The decision maker can shift attention to the rebuttal P004, and query for possible antecedents. This may result in activation of information that supports the negation of the rebuttal, releasing support from this argument for P001. (Of course, it may also result in confirmation of the rebuttal.) In the case of a causal mental model, the decision maker can clamp the rebuttal true; this may lead to conflict if the rebuttal is the most plausible explanation of some other hypothesis which is thought to be true.

(3 & 4) If these do not work, the decision maker can focus on P001 itself. Querying to retrieve possible antecedents of either +P001 or –P001 probably won't work: P001 was the hypothesis of primary interest so it has already been attended. But clamping the + or "-" collector may lead to the activation of some hypothesis that conflicts with it.

This example illustrates three general strategies for dealing with incompleteness in a hypothesis of interest:

1. find evidence for existing arguments
2. rebut rebuttals
3. create new arguments

In all these cases, the goal is not simply to make an *assessment* of the truth of a hypothesis (e.g., P001, P002, or P004). The goal is to generate *reasons* for an assessment. The decision maker tries to activate previously implicit information in long-term memory, use it to elaborate the active belief network, and thereby find new *arguments* to resolve the incompleteness in P001. Decision makers must practice knowledge elicitation on themselves, and metacognitively skilled decision makers learn where to look.

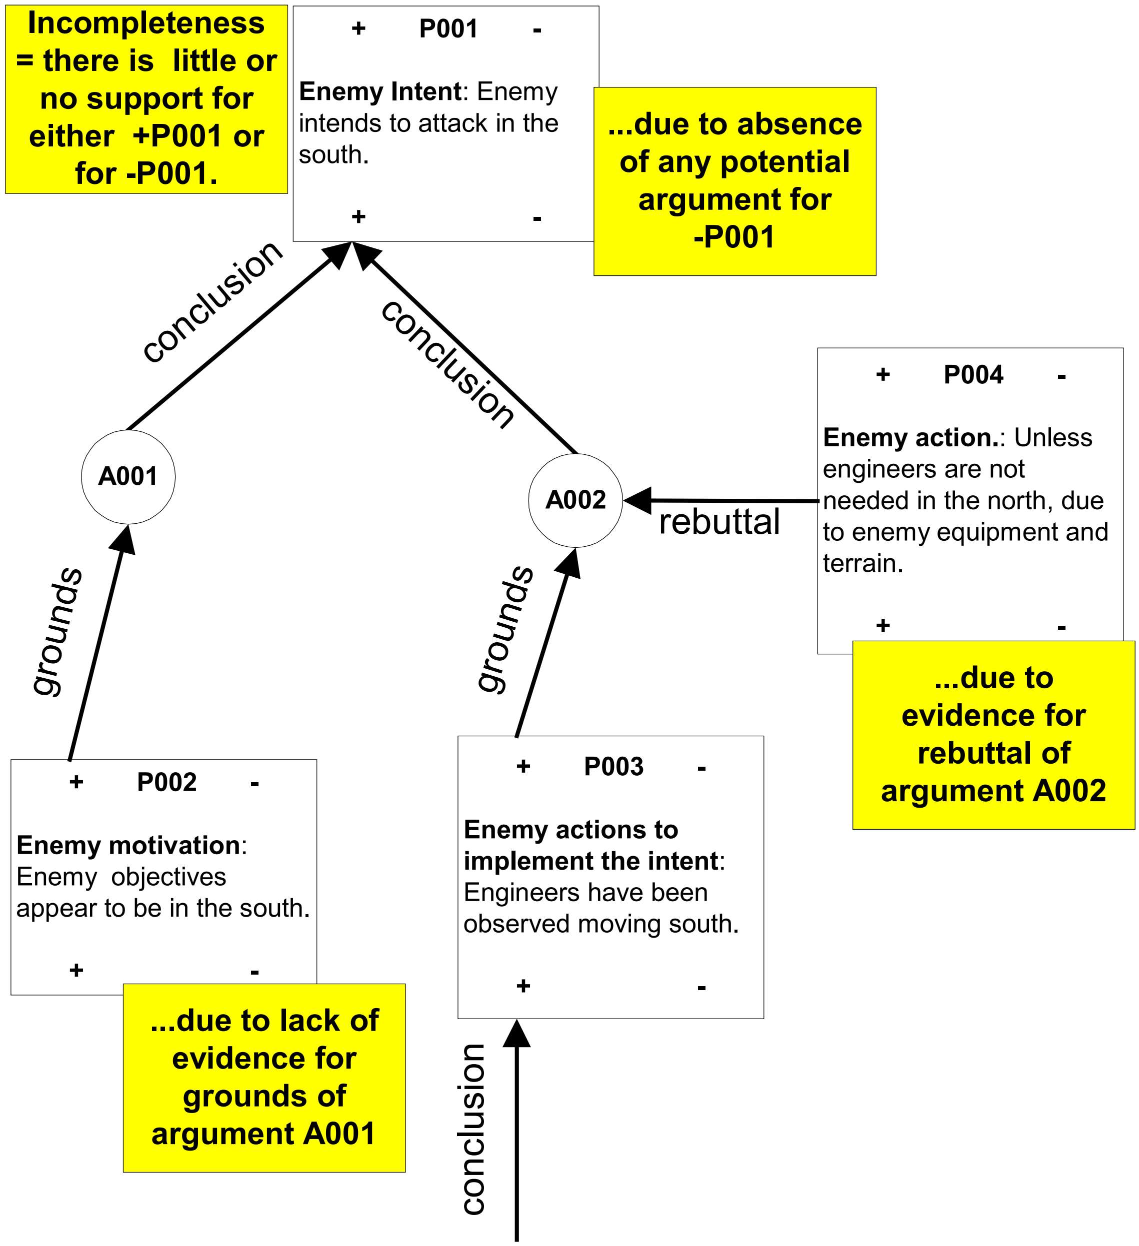

Figure 38. Contextual reasons for incompleteness of information at node P001. Nether the + or "-" collector of P001 is activated because of the pattern of argument elsewhere in the network.

General strategies for conflict.

Like incompleteness, conflict is a local measure that brought about by certain features of the argument context. Suppose, for example, that Figure 39 was the result of resolving the incompleteness of P001 in Figure 38. A more thorough analysis has uncovered convincing reasons for believing that the enemy's objectives' are in the south, supporting +P001; but a conflicting argument has also been developed for -P001, based on a well-supported intel report that enemy assets are concentrating in the north. The result is *conflict* between +P001 and -P001. Moreover, the decision maker is not currently aware of rebuttals of argument A001, and has no evidence to support the rebuttal of argument A003.

As in the case of incompleteness, there are several steps the decision maker can take to deal with this conflict:

(1) Shift attention to the rebuttal of one of the arguments, P006. This may result in retrieval of information that would support it. If not, the decision maker can clamp it false, and see if that elicits any conflict. If this works, one of the two conflicting arguments will be neutralized by the rebuttal, and the conflict at P001 will be resolved.

(2) If this doesn't work, the decision maker might turn to the other argument A003. No rebuttal has yet been thought of for A003. In order to elicit possible rebuttals, the decision maker can assume that the rule is false, and look for an explanation (i.e., a condition that could neutralize the rule). To do this, the decision maker clamps P001 false and P002 true. This is equivalent to querying for an explanation of the failure of the rule. If new rebuttals are generated, the decision maker may then attempt to evaluate them, by looking for grounds for their truth or falsity.

(3 & 4) The ultimate cause of the conflict at P001 is the existence of grounds for both lines of reasoning, i.e., P002 and P005. The decision maker may now look for ways to discredit those premises. They may be queried to look for disconfirming evidence, or clamped true to look for conflict. If these succeed, the conflict at P001 will be resolved but at the cost of new conflict at P002 or PP005, which the decision maker may wish to address. If these fail, the decision maker should at least have succeeded in retrieving grounds *in favor* of P002 and P005. The next step could be to look for rebuttals for these arguments, as in (2) or if rebuttals have been retrieved, for grounds to believe the rebuttals, as in (2). Alternatively, the decision maker might conclude that there were no convincing grounds for P002 or P005 in the first place, i.e., they depended on an implicit assumption.

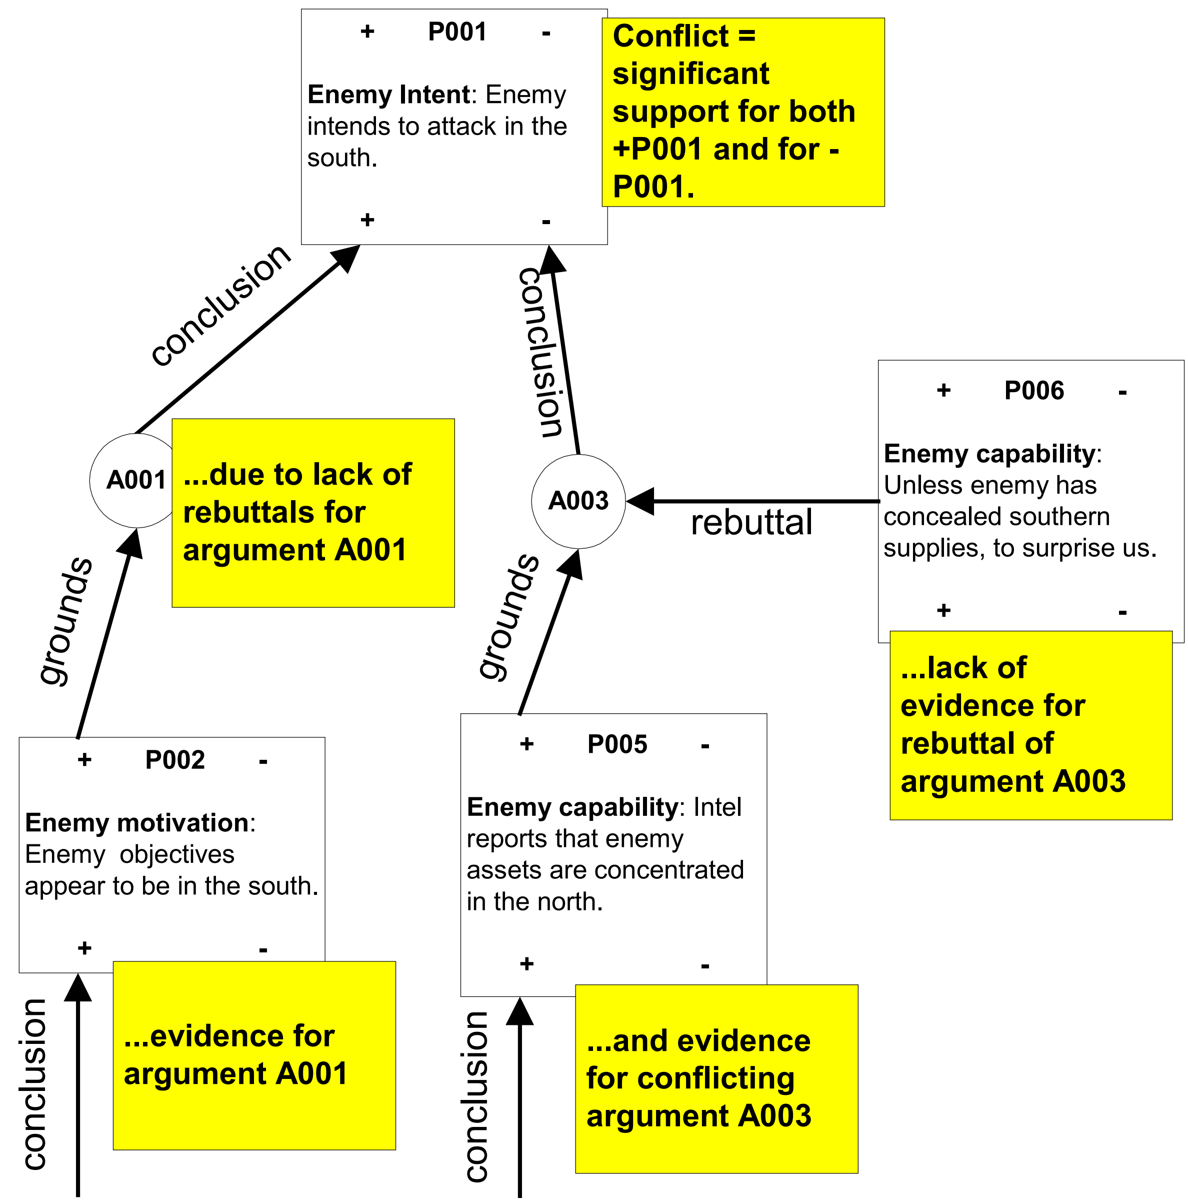

Figure 39. Argument network showing possible causes of conflict at hypothesis P001.

This example illustrates three general strategies for dealing with conflict in a hypothesis of interest:

1. find evidence for a rebuttal to one of the conflicting arguments
2. create new rebuttals
3. find evidence against the grounds for the conflicting arguments

These responses are quite different from those that seemed appropriate in the case of incompleteness.

General reflective skill may both support, and be supported by, the acquisition of more specialized reflective skills in particular domains. On the one hand, decision makers might acquire general skills in part by abstracting or generalizing from specialized experience, although we expect that more explicit training would also be useful. On the other hand, general skills might (i) serve as a substitute for specialized knowledge in unfamiliar domains, and (ii) provide an entering wedge for more efficient learning of domain-specific knowledge in new domains.

Advantages of Discriminating Uncertainties

These concepts –general measures of local uncertainty, domain-specific measures of culpability or general roles in arguments, shifting attention, and clamping activation – provide a starting syntax for metacognitive critiquing and correcting strategies. A particular metacognitive strategy can combine these elements in any number of ways, depending on how it is adapted by processes of reinforcement learning to problem domains or personal experience. A common element of both the domain-specific and general strategies that we have discussed is that they are tailored to different types of uncertainty. Based on interviews with decision makers in several domains, we hypothesize that this "uncertainty-specific" aspect of metacognitive strategies tends to be quite general, is learned in a similar way in a variety of domains, and tends to be more effective than strategies that treat all types of uncertainty alike.

The particular hypothesis that we will pursue in this section is fundamental to the R / M model: *In complex or novel domains, it is often better to address different types of local uncertainty sequentially, rather than trying to resolve all uncertainty at once.*

Comparison of Figure 22 and Figure 23 shows that a not unlikely outcome of successfully reducing incompleteness is an increase in conflict. For example, a decision maker who starts out relatively ignorant on a topic, may well find evidence that appears to support opposing conclusions. Similarly, the price for reducing conflict may in some cases be an increase in incompleteness . For example, a decision maker trying to understand and explain an apparent conflict between two lines of reasoning may eventually break the Gordian knot by rejecting an assumption common to *both* conflicting arguments.[40]

It might be tempting to try to avoid situations in which one appears to jump from one problem to another, by aiming to minimize the sum of all three kinds of local indecisiveness (incompleteness, conflict, and lack of resolution) *at the same time*. A decision maker seeking to simultaneously reduce all three measures would try to move lower in the landscape of Figure 25. Notice that, unlike Figure 22 and Figure 23, the combined measure cannot be minimized unless a conclusion is found that has support equal to 1.0, while its negation has support equal to 0.0. The problem must be solved in one fell swoop. Both empirical observation and theory, however, suggest that such a combined search will often go wrong. Addressing incompleteness and conflict simultaneously may lead either (1) to premature closure or (2) to failure to find a solution at all.

---

[40] Does the barber who shaves all and only those who do not shave themselves, shave himself? Argument (1): If he shaves himself, he does not shave himself. Argument (2): If he does not shave himself, he does shave himself. An assumption common to both arguments: *This barber exists.* Drop this assumption, and the paradox disappears.

(1) Problem solving that seeks to minimize the aggregate measure of indecisiveness will be strongly biased toward exploring hypotheses that promise an *immediate* reduction in all uncertainty. As a result, the aggregate strategy will completely discount nodes whose exploration could decrease incompleteness at the expense of increasing conflict, and nodes whose exploration could decrease conflict at the expense of increasing incompleteness. It will steer away from information sources that might lead to the expression of conflicting views, or for evidence that could disconfirm an already supported conclusion. And it will not even try to look for information that might legitimately rebut both of two conflicting arguments. In some decision making contexts, therefore, this strategy is likely to halt complacently with a wrong answer before the problem has been adequately explored.

(2) On the other hand, in some decision making problems, the aggregate strategy may be so ineffective that it fails to find an answer at all. This will be the case if there is no *single* hypothesis, or information collection option, whose further exploration is likely to lead to the elimination of incompleteness, conflict, and low resolution at the same time, without the benefit of an assumption. In such situations, the only available path to a solution will be unforeseeable in advance and indirect, through a sequence of information collection and analysis way stations that have more limited objectives, e.g., collecting information to reduce incompleteness, detecting conflicting lines of argument in the new information, trying to explain the conflict by rebutting assumptions that underlie the arguments; collecting more information to fill new gaps; etc. The information and understanding that are acquired along the way are used to explain apparent inconsistencies, expose assumptions, and construct a new, more valid set of arguments, leading eventually, if all goes well, to a consistent, well-supported solution. The all-or-nothing aggregate strategy, by contrast, may never even get started.

It is important to realize that jumping from one type of uncertainty to another (e.g., reducing incompleteness but increasing conflict or vice versa) does not represent failure of the problem solving process. On the contrary, such outcomes always mean that new information has been added to the mix (either retrieved or collected), even if a final interpretation of its meaning has not yet been arrived at. The cumulative effect of such moves is better knowledge of the problem space: a better grasp of what kinds of evidence are available in the relevant domain, and what kinds of arguments work and which do not. The aggregate strategy deprives decision makers of this information, and – in complex problems – can succeed only by luck.

REFERENCES

Anderson, J.R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.

Brandimonte, M., Einstein, G.O., & McDaniel, M.A. (1996). *Prospective memory: Theory and applications*. Mahwah, NJ: Erlbaum.

Cohen, M.S. (1986). An expert system framework for non-monotonic reasoning about probabilistic assumptions. In J.F. Lemmer & L.N. Kanal (Eds.), *Uncertainty in artificial intelligence* (pp. 279-293). Amsterdam: North-Holland Publishing Co.

Cohen, M.S. (1993). The naturalistic basis of decision biases. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsambok (Eds.), *Decision making in action: Models and methods* (pp. 51-99). Norwood, NJ: Ablex.

Cohen, M.S., Schum, D.A., Freeling, A.N.S., & Chinnis, J.O., Jr. (1984, August). *On the art and science of hedging a conclusion: Alternative theories of uncertainty in intelligence analysis* (Technical Report 84-6). Falls Church, VA: Decision Science Consortium, Inc.

Cohen, M.S., Freeman, J.T., & Thompson, B.B. (1998). Critical thinking skills in tactical decision making: A model and a training method. In J.A. Cannon-Bowers, & E. Salas (Eds.), *Decision-making under stress: Implications for training & simulation.* Washington, DC: American Psychological Association Publications.

Cohen, M.S., Freeman, J.T., & Wolf, S. (1996). Meta-recognition in time stressed decision making: Recognizing, critiquing, and correcting. *Human Factors, 38*(2), pp. 206-219.

Cohen M.S., Parasuraman, R., Serfaty, D., & Andes, R. (1997). *Trust in decision aids: A model and a training strategy*. Arlington, VA: Cognitive Technologies, Inc.

Cohen, M.S., Adelman, L., Tolcott, M.A., Bresnick, T.A., & Marvin, F. (1993*). A cognitive framework for battlefield commanders' situation assessment*. Arlington, VA: Cognitive Technologies, Inc.

Cohen, M.S., Thompson, B.B., Adelman, L., Bresnick, T.A., Tolcott, M.A., & Freeman, J.T. (1995). *Rapid capturing of battlefield mental models.* Arlington, VA: Cognitive Technologies, Inc.

Ellsberg, D. (1988; original pub. 1961). Risk, ambiguity, and the Savage axioms. In P. Gardenfors & N-E. Sahlin (Eds.), *Decision, probability, and utility*. NY: Cambridge University Press.

Ennis, R.H. (1996). *Critical thinking.* Upper Saddle River, NJ: Prentice Hall.

Howard, R. (1960). *Dynamic programming and Markov processes.* Cambridge: MIT Press.

Just, M.A., & Carpenter, A. (1977). *Cognitive processes in comprehension.* Hillsdale, NJ: Erlbaum.

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

King, P.M., & Kitchener, K.S. (1994). *Developing reflective judgment*. San Francisco: Jossey-Bass.

Kuhn, D., Amsel, E., & O'Loughlin, M. (1988). *The development of scientific thinking skills*. San Diego: Academic Press, Inc.

Lukes, G., Thompson, B.B., & Werbos, P. (1990). Expectation driven learning with an associative memory. *In Proceedings of the International Joint Conference on Neural Networks* (IEEE, January).

Mani, D.R. (1995). The design and implementation of massively parallel knowledge representation and reasoning systems: A connectionist approach, (Ph.D. Dissertation), Department of Computer and Information Science, University of Pennsylvania.

Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning internal representations by error propagation. In D.E. Rumelhart, J.L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition: Volume 1: Foundations* (pp. 318-362). Cambridge, MA: The MIT Press.

Shastri, L. (1992). Neurally motivated constraints on the working memory capacity of a production system for parallel processing. In *Proceedings of the Fourteenth Conference of the Cognitive Science Society*, Bloomington, IN., pp. 159-164.

Shastri, L. (1999a). Advances in Shruti--A neurally motivated model of relational knowledge representation and rapid inference using temporal synchrony. *Applied Intelligence*.

Shastri, L. (1999b). Types and quantifiers in Shruti---A connectionist model of rapid reasoning and relational processing. In S. Wermter & R. Sun (Eds.), *Hybrid neural symbolic integration*, Lecture Notes in Computer Science, Lecture Notes in Artificial Intelligence. Berlin: Springer-Verlag.

Shastri, L., & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Brain and Behavioral Sciences, 16*, 417-494.

Shastri, L., & Grannes, D.J. (1996) A connectionist treatment of negation and inconsistency. In *Proceedings of the Eighteenth Conference of the Cognitive Science Society.* San Diego, CA.

Shastri, L., & Mani, D.R. (1997) Massively parallel knowledge representation and reasoning: Taking a cue from the brain. In J. Geller, H. Kitano, & C. Suttner (Eds.), *Parallel processing for artificial intelligence 3.* Elsevier Science.

Shastri, L. & Wendelken, C. (1998). Soft computing in Shruti---A neurally plausible model of reflexive reasoning and relational information processing. To appear in the *Proceedings of Soft-Computing '99*, Genova, Italy.

Toulmin, S.E. (1958). *The uses of argument.* Cambridge: Cambridge University Press.

Ullman, H., & Wade, J. (1996). *Shock and awe: Achieving rapid dominance*. Washington, DC: National Defense University Press.

Werbos, P. (1977). Advanced forecasting methods for global crisis warning and models of intelligence. *General Systems Yearbook*. (Appendix B).

Werbos, P. (1989). A menu of reinforcement learning over time. In W. T. Miller & R. S. Werbos (Eds.) *Neural networks for robotics and control.* Cambridge: MITR Press.