



AFRL-RI-RS-TR-2010-102

DISTRIBUTED EPISODIC AND ANALOGICAL REASONING (DEAR)

RAYTHEON BBN TECHNOLOGIES CORPORATION

April 2010

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE**

NOTICE AND SIGNATURE PAGE

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

This report was cleared for public release by the 88th ABW, Wright-Patterson AFB Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2010-102 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/
GENNADY R. STASKEVICH
Work Unit Manager

/s/
JULIE BRICHACEK , Chief
Information Systems Division
Information Directorate

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

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 data sources, gathering and maintaining the data needed, and completing and 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 Service, Directorate for Information Operations and Reports, 1215 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.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) APRIL 2010		2. REPORT TYPE Final		3. DATES COVERED (From - To) March 2008 – March 2010	
4. TITLE AND SUBTITLE DISTRIBUTED EPISODIC AND ANALOGICAL REASONING (DEAR)				5a. CONTRACT NUMBER FA8750-08-C-0099	
				5b. GRANT NUMBER N/A	
				5c. PROGRAM ELEMENT NUMBER 62702F	
6. AUTHOR(S) Alice Mulvehill				5d. PROJECT NUMBER 558S	
				5e. TASK NUMBER 8C	
				5f. WORK UNIT NUMBER P1	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Raytheon BBN Technologies Corporation 10 Moulton Street Cambridge, MA 02138-1119				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL/RISC 525 Brooks Road Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S) N/A	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-RI-RS-TR-2010-102	
12. DISTRIBUTION AVAILABILITY STATEMENT Distribution Approved for Public Release; Distribution Unlimited. PA# 88ABW-2010-2145 Date Cleared: 20-April-2010					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT An AFRL program, Distributed Episodic Exploratory Planning (DEEP), is developing a mixed-initiative decision-support environment where commanders can readily access and leverage historical data from distributed sources for use in decision making. Episodic reasoning paradigms and specific Case Based Reasoning (CBR) technology are being considered as methods to facilitate the use of analogical reasoning and past experience [Ford & Lawton, 2008]. The research conducted in the DEAR (Distributed Episodic Analogical Reasoning) project builds on previous research conducted by Raytheon BBN Technologies Corp. (BBN) for the DEEP project [Mulvehill, Deutsch, & Rager, 2007] and is intended to further influence the design, development and implementation of CBR and analogical reasoning in the AFRL DEEP program.					
15. SUBJECT TERMS Analogy, Episodic, Experience, Mixed-Initiative, Case Based Reasoning, Agent, DEEP, CBR, DEAR					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 35	19a. NAME OF RESPONSIBLE PERSON Gennady R. Staskevich
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) N/A

Table of Contents

1. Introduction.....	1
2. Background	3
2.1 Episodic Reasoning	3
3. Technical Approach.....	5
4. Technical Discussion.....	7
4.1 Multiple Distributed Actors and Multiple Cases	8
4.1.1 Coherence	12
4.1.2 Collaboration and Negotiation.....	14
4.2 Analogical reasoning	16
4.3 Distributed CBR Management.....	20
5. Recommendations	21
6. Conclusion.....	23
7. References	26
8. Acronyms	26
Appendix A. Additional Reviewed Papers	30

List of Figures

Figure 1. Example of a three level model of social contexts.	9
Figure 2. Distributed Problem Solving Use Case	11
Figure 3. DEEP Design and Development Approach.....	24

List of Tables

Table 1. Summary of Gasser Research on Collaboration and Coherence	16
Table 2. Summary of Analogical Reasoning Research.....	18
Table 3. Analogical Reasoning Issues and Recommendations	21
Table 4. Distributed CBR Management Issues	21
Table 5. Multiple Actors and Case Base Issues.....	22

1. Introduction

An AFRL program, Distributed Episodic Exploratory Planning (DEEP), is developing a mixed-initiative decision-support environment where commanders can readily access and leverage historical data from distributed sources for use in decision making. Episodic reasoning paradigms and specific Case Based Reasoning (CBR) technology are being considered as methods to facilitate the use of analogical reasoning and past experience [Ford & Lawton, 2008]. The research conducted in the DEAR (Distributed Episodic Analogical Reasoning) project builds on previous research conducted by Raytheon BBN Technologies Corp. (BBN) for the DEEP project [Mulvehill, Deutsch, & Rager, 2007] and is intended to further influence the design, development and implementation of CBR and analogical reasoning in the AFRL DEEP program. The three primary research objectives of the DEAR project are:

1. Identify, develop and recommend technology necessary to leverage case-based reasoning (CBR) technology and analogical reasoning in a distributed environment.
2. Identify and evaluate issues constraining the employment of CBR technology in a distributed environment.
3. Provide design and architectural recommendations, with an emphasis on how different subjective perspectives can influence the annotation, representation and usage of one or more episodes as they are evolved by multiple participants within a net-centric planning and execution environment.

Commanders in military Command and Control (C2) environments generally operate within a group, each with a prescribed role and protocols for how to interact; therefore, the research in DEAR has been focused on issues associated with group decision-making and the use of historical data by the group both as a whole and as individuals within that group. In addition to investigating how historical data is used, our research has examined how roles and the individual perspectives of the actors¹ involved in the group problem-solving effort can influence what each individual contributes and what is ultimately used to generate a plan or Course of Action (COA).

For our research, we made several assumptions including: any actor can be distributed in space, and potentially time; each actor will leverage their own experience to solve the problem; and each actor's experience may or may not be stored and represented electronically in an episodic memory or case base. We also acknowledged the importance of software support for C2 decision-making. Software support includes any tool that supports the human problem solver, including intelligent software agents that are capable of reasoning and adapting to a problem space working either autonomously or using mixed-initiative interactions with a human.

¹ Note: the term actor and agent are used interchangeably in this document to refer to either a human or software agent that is participating in a group decision making effort.

The product generated in DEEP will be a military COA. That COA should be “coherent”. By coherent, we mean that the product must conform to some legitimate model of a product that is of relevance to the problem being solved, e.g., a COA or a plan such as a military logistics deployment plan. During the course of our research, we investigated several existing approaches for coherence and also designed and partially developed an approach for use in DEEP that applies a constraint based coherence strategy based on the work of Thagard [2002]. Coherence is described in more detail in Section 4.1.1 of this report.

This final project report provides findings and recommendations that are a result of the selective research that was conducted during this project. While most of the reviewed research papers are referenced in this document, Appendix A provides a categorized list of some additional papers that were reviewed. Many of these papers address potential future research for the DEEP program. While recommendations for the DEEP program are provided throughout the report in the form of *questions to answer* and *things to do*, Section 5 presents recommendations that require answers, while Section 6 concludes with the things to do.

2. Background

When software systems are designed and developed with a CBR approach, the intent is that the system will help a user to solve a current problem by providing the user with access to related historical experience. Because the past and the present are not exactly equivalent, the previous experience/case must be *revised* to fit the needs of the current problem. Each time a problem is solved, that experience is added to the case base, thus allowing the user to learn over time. For a more complete description of CBR technology, see the report “Case Based Reasoning for DEEP: Observations and Recommendations” by Mulvehill, Deutsch, & Rager [2007].

For the DEEP program, there is a requirement that the architecture must support the access and usage of historical information:

“The Distributed Episodic Exploratory Planning (DEEP) project is a mixed-initiative decision support system that utilizes plan experiences, encoded into a case base reasoning (CBR) system, to suggest courses of action for new situations. It is being implemented as a distributed multi-agent system, using agents to maintain and exploit the experiences of individual commanders as well as to transform suggested past plans into potential solutions for new problems. The agents interact through a common knowledge repository. The primary challenge of the DEEP project is translating the experiences collected from good (or potentially bad) command decisions into a form that is understandable by a computer and amenable for use in mixed-initiative planning. The key is to represent knowledge in a form that facilitates inferencing (i.e., drawing conclusions from knowledge).” [Ford and Carozzoni, 2007]

2.1 Episodic Reasoning

While the term *case based reasoning* is generally used to describe a specific type of reasoning in automated systems, the term *episodic reasoning* is used to describe a similar activity in humans. Both case based and episodic reasoning rely on the retrieval and reuse of past experiences (i.e., case bases or episodes).

Episodic reasoning in humans is complicated by a variety of factors, including: the sensitivity, accuracy and effectiveness of the sensors used to acquire data; the cognitive ability of the individual to use the sensed data; the role of the human in a particular problem solving context; the goals of the human in the decision making context; the influence of multiple, active or non-active, and possibly competing goals; and the influence of previous experience via training or hands on experience². While human memory has been an active area of research for decades, in DEAR, we were particularly influenced by the research of Schacter [2001] who’s work focuses on how decision making problems can be influenced by human memory and perception; and the research of Conway [2005] whose work emphasizes the interconnectedness of self and memory and a conceptual framework called the Self-Memory System.

² In an automated CBR system, these factors are also present but can often be controlled through design and architecture decisions.

In automated systems, episodic memory can be represented as a case base [Shank, 1982; Weber, 1996]. Here an episode is represented as a case with indices that describe the salient features of that experience. CBR storage and retrieval techniques are often used to create decision support tools that enable users to take advantage of previous recorded experience. In DEEP, CBR is being investigated to create a decision support system that would support the development of quick and efficient responses to a crisis situation. *Crisis action planning* is a knowledge-intensive activity that is complicated because of the speed required, the uncertainty and ambiguity associated with the problem context, and the experience level of either the human decision maker and/or the automated system that is responding to the situation. Additionally, there is often little time to plan from scratch so early crisis response planning tends to heavily leverage past experience. From a theoretical point of view, given a human expert with an extensive episodic memory, problem solving can be interpreted as a form of pattern matching over events in time. Klein [1998] has termed this process *recognition primed decision-making*. From a more practical point of view, responding to a crisis situation involves both recognition and action. The patterns in a person's episodic memory that enable recognition-primed decision-making are representative of methods employed in CBR approaches to index, store and retrieve cases in a case base.

Most CBR systems incorporate modules to support the process that has been defined as the *CBR Cycle* [Aamodt & Plaza, 1994]. This cycle is comprised of the following functions: *retrieve*, *reuse*, *revise*, and *retain*. The CBR cycle can be used by a single problem solver who leverages his/her own experience, or the CBR cycle can be implemented in a distributed environment in a variety of ways to include:

- One case base and multiple distributed users,
- Multiple case bases and a single user,
- Multiple case bases and multiple users.

Currently DEEP has been developed to support multiple case bases and multiple users. In the rest of this paper we will review some issues associated with the current DEEP approach that DEEP researchers will need to evaluate as they continue to evolve DEEP to more thoroughly leverage CBR and analogical reasoning technology.

3. Technical Approach

The research conducted over the course this project was based on collaborative work and frequent interactions between BBN and DEEP research staff. A selective review of the literature was conducted in the areas of CBR and episodic reasoning, with a particular focus on how these technologies can be applied in a distributed environment, handle and leverage analogical reasoning, and be managed over time.

Our research has been incremental. As our research progressed, we began to identify certain research issues that we considered to be of particular interest to DEEP. For example, to provide DEEP with functionality to support multiple actors with multiple cases within a distributed environment, our research indicates that DEEP researchers will need to evaluate (1) how the perspective and self reference of the actors affects their case base development, (2) the use of ontologies for supporting interoperability across non-homogeneous actors, and (3) issues associated with case selection and merging from multiple case bases.

Distributed Environment Interactions: We began our research by considering the requirements for a distributed agent environment. In a distributed environment, there are generally many participants. Each participant will likely have a slightly different perception of a given problem solving context and each participant will also have unique historical experiences. For example, an operations officer and an intelligence officer will each take away a different view of a military engagement because of task responsibility and their own experiences. These differences in perspective will influence how the episode is stored and annotated³. The work of Mantovani [1996] presents a model that describes how the role, multiple goals and tools available for decision making can influence how decisions are made. Our research to date indicates that a variant of the Mantovani model could be useful in DEEP for determining how a commander should evaluate contributions from actors in a distributed problem solving context. The research of Mantovani is described in more detail in Section 4.1 of this paper.

Some of the perceptions of individual actors in the group may be based on an analogical evaluation and interpretation of the problem episode by the actor. Like CBR, analogical reasoning is a method that can be used by actors to learn and to solve problems. Like CBR, analogical reasoning leverages historical experiences (episodes) with pattern matching on aspects of the current problem with previous learned problems. The work of Dehghani et al [2009] presents a good introduction to the use of analogies in problem solving. But analogical reasoning can be more complicated than basic CBR and recording how an analogy is formed and how it contributes to the interpretation of a problem can be difficult. A paper by French [2002] provides a good overview of the benefits and technical challenges associated with analogical problem solving. Analogical reasoning research is described in more detail in Section 4.2 of this paper.

³ For more information on the creation of case features see Leake and Wilson [1999]

In many problem solving environments, like the COA development environment that could be supported by DEEP for crisis action planning, human problem solvers will likely interact with intelligent software or software agents to solve their problems. The work of Gasser [1991] on multi-agent systems (MAS) describes some of the issues associated with distributed problem solving in human/human or human/software groups. See Section 4.1.2 of this report for more detail. Because both human and non-human problem solvers often interact in a social context, the social aspects of collaborative computing also need to be considered. The work of Prietula and Carley [1999] and Anderson [2009] provide many insights into these issues.

Net-centricity: The DEEP program also has an objective to support the net-centric paradigm. This network-centric objective implies that distributed users must be able to access, use and share information in novel distributed ways, e.g., shared situation awareness, knowledge of the commander's intent. Research into multi-agent systems (MAS) by Gasser [1991] provides insights and requirements for how human and software agents can contribute and share data. From the perspective of distributed CBR, Plaza and McGinty [2005] suggest two dimensions along which distributed approaches to CBR systems can be defined: how knowledge is *processed* in the system and how knowledge is *organized* in the system. The processing may be accomplished by a single or multiple humans/agents and the knowledge may be organized within a single case base or in multiple case bases. For DEEP, there are multiple agents and the knowledge is organized into multiple case bases. This approach introduces a number of problems that need to be better understood in order to guarantee proper behavior and useful performance (see Section 4.3 for more details). These problems include: the capture of case solution rationale; usage of rationale in a distributed environment; merging and adaptation of cases from disparate case bases; and issues associated with the creation, updating, and maintenance of cases and case features in a distributed environment. When analogical reasoning is involved, explanations of how an analogy was used may facilitate better understanding and acceptance of the analogy by the group. In addition, explanation by any one member may enable other participants to better grasp and utilize the benefits of a *good* analogy. (Note, we have not considered the negative consequences of using a bad analogy, however in Section 4.1.1 of this report, we discuss how coherence mechanisms can be used to constrain the use of episodic and analogical data from multiple agents).

4. Technical Discussion

In the very early stages of our research, we decided to conduct a small experiment to test some of our assumptions about how past experience, goals, and context can influence the development of perspective and the associated decision making. In the experiment, a set of images about birthday party items (e.g., cake, balloons) were presented to *individual* subjects who were asked to describe what they saw. We conjectured that since most, if not all of the subjects, had probably experienced birthday events with cakes and balloons, that the presentation of these selected pictures would invoke their memories/episodes and hence descriptions of past birthday events, which would then bias the interpretation of subsequent images.

After the subjects viewed the birthday images and gave their descriptions, they were then presented with a picture of a woman moving cakes off shelves (or putting them on shelves). The subjects were asked to describe what the woman was doing. Most of the subjects said that the woman was possibly working in a bakery, perhaps the origin of the birthday cake. It was hard for anyone to be sure if she was putting cakes on the shelf or removing them from the shelf since the image was basically a single clip from some on-going context. The results indicate that the pictures of the birthday events established a “context” by which subsequent images (the lady with the cakes) were interpreted.

The results from the experiment seem to have confirmed our assumptions. To check our results and to elaborate on the influence of context, we conducted another experiment. Here a *group* of people were presented with the birthday images. Like the individual participants, the group was asked to describe what they saw. As members of the group described what they saw, the general agreement was that the scenes were about a birthday party. Next, the picture of the woman moving the cakes was presented, and the group was asked what the woman was doing. Most of the people in the group said that the woman was in a bakery, perhaps selling the birthday cake.

Next, the group was told that the woman in the picture was the wife of a suspected drug dealer and that the bakery was suspected of being a drug transfer site. Very quickly, members of the group began to offer ideas about how the birthday party and other images could be related to drug trafficking. The birthday images appeared to take on new meaning as a function of a new perspective. A larger schema seemed to form among the group about drug trafficking. Members of the group began to offer revised explanations about the birthday scenes and how they could be related to the new context. So, it appeared that the context did influence how the data was being interpreted. But it also appeared that the group was trying to create a coherent story about the birthday party and the cake with respect to the new drug trafficking context. Afterwards, while evaluating the experiment results, we realized that the images presented in the experiment were single scenes with gaps in time. We proposed that if some of the gaps were filled with clips of scenes either before or after a particular scene, that there would be more certainty about what was happening in subsequent scenes. Research in episodic and analogical reasoning suggests that historical data is often used to fill in the gaps [Kokinov, 2001; Padovotz et al, 2006]. Could it be that members of the group were collaboratively trying to use their own experiences to form a coherent story? Researchers like Wang and Gasser [2002] suggest that teams often do form shared perceptions as a result of collaboration within a Multi Agent System. Research on collective cognition [Paul, 2001; Rettberg, 2005; and Aleman-Meza et al, 2005] also supports this finding.

4.1 Multiple Distributed Actors and Multiple Cases

When multiple agents (human or software) interact to solve a problem, does the role that an agent takes in a problem solving context constrain the problem solving actions? In the cake experiment, nobody had a particular role. But in the DEEP COA environment, each agent will typically have a role, e.g., logistics officer, intelligence officer, etc. Will the role of the agent influence how aspects of the problem are perceived, interpreted and acted upon? Our intuition leads us to believe that the answer to both of these questions is “yes”. To explore this more, we hypothesize that the perspective of a given agent is a function of policies and data access associated with a current role, plus historical experiences obtained from participation in one or more previous roles.

We also advocate that cultural models are associated with a role. In our work, we define a cultural model as a knowledge base that describes the language, the tasks and the policies and/or rules that are required for interaction with others from the position of a particular role. For example, an intelligence officer will have certain tasks to perform (as defined by some procedural document). During collaboration, members of the group will expect and/or depend on certain products (possibly in a specified format) from the intelligence officer. The research of Giuseppe Mantovani [1996] provides a model that may be of value to DEEP researchers for better understanding and evaluating the importance of *role*. Figure 1 describes Mantovani’s three tiered model. In the model, each participant’s view of a problem can change as a function of his/her participation (construction of the problem, interpretation of the situation, and local interaction with the environment). Mantovani emphasizes the importance of artifacts (tools as defined in Level 3) that are available for interaction by an agent with the environment.

According to Mantovani, these artifacts and/or tools will also influence and/or constrain how a problem is viewed and subsequently how it is solved.

In the Mantovani model goals and interest are introduced in the 2nd level. It is at this level that subconscious goals or long range unspecified goals and objectives can influence both problem perception and how tasks are selected and performed. Current work by Prietula and Carley [1999] describes how personalities and motivation, e.g., benevolent vs. exploitative can influence both the contribution by an actor and how that actor participates within a group. For a CBR system, this type of information can be represented, to some degree, with user profile information that could be captured as indices on a case. The goals of the actor, as well as the active goal within the problem solving group could be represented with an ontology of explicit/implicit or specified/implied goals and actions.




Figure 1. Example of a three level model of social contexts.

One of the most critical aspects of episodic reasoning is the notion that an episode is associated with an actor, from a “self” perspective [Struck and Ganger, 2003]. In a distributed environment, some mechanism needs to be developed to store the collection of related episodes, each annotated with the perspective of the actor who participated in that episode. Classic CBR systems can implement self perspective through indices with weights. The problems arise when another actor with a different self perspective wants to utilize another agents’ case base. Level 1 is where Mantovani defines interaction with others within a social context. He defines cultural models to be at this level and in our interpretation of his work the cultural models are socially approved actions within a specified problem solving context that describe and delimit actions for some role within a given problem solving context.

If the Mantovani model were to be implemented in DEEP, we recommend that the implementation allow for elements of all three levels to be active in parallel. For example, DEEP representation should recognize that a tool may be used in a particular situation because it is the required tool (by policy), or because the actor perceives an opportunity (associated with a higher-level goal) for using a different tool to achieve the goal. The choice of the tool could also be influenced by the actor's historical experience, e.g., a more experienced actor may prefer one particular tool instead of a tool that a novice actor would choose to use. Actions and associated tools for a given role may also be defined by the cultural model that is associated with the role within the problem solving context.

For DEEP, the Mantovani model could also be used to support the representation of each of the actors involved in a problem solving environment. This would allow the actors to collaborate and provide solutions from their perspective and historical experiences. However, it would not guarantee a useful COA unless some coordinating agent was available to manage contributions. Although the group may, over time, form a coherent product (as was observed in the cake experiment), another approach for guaranteeing the coherence of a product such as a COA within some specified amount of time is to introduce a coordinating agent. For this function, we propose a “watcher” agent that observes the actors in the group and uses its own cultural model, along with some coherence maintenance methods, to combine their inputs into a *coherent* product.

Figure 2 presents a use case that leverages the multiple perspective models of Mantovani. In this use case, multiple actors, each in a particular role with goals, cultural models, and past experience are engaged in creating and selecting a COA in response to a natural disaster, e.g., earthquake. For the sake of discussion, we imagine that the individuals in the problem are located in distributed settings but can easily exchange information through one or more communication tools. We assume that an ontology or other terminological/language translation tools are available to constrain and translate communication within the group.

Course Of Action Generation




Figure 2. Distributed Problem Solving Use Case

In the use case, each participant has a role associated with COA development and specific tools that they can use to support problem solving. For example, the Logistics Officer will have a set of tools that support the generation of logistics support for the disaster. The Logistics Officer will also be able make decisions about the resources that he/she has available. Each participant has his/her own unique history (episodic memory) that can be referenced and each has a cultural model that describes how that participant should act (according to existing policies, procedures and/or rules) in the role he/she has. Here the Logistics Officer may have policies that stipulate that search and rescue teams should be provided early in the disaster and that transport resources should be provided to evacuate US citizens from the affected area. In order to evaluate experience over time and to introduce concepts of cultural models, we include two agents, Agent 2 and Agent 3. These agents are set up to have previous experience in a given role but from a different problem solving scenario. As a result of their past experience, our approach assumes that their cultural models will be similar. We believe that the common experience and background between Agent 2 and Agent 3 will also result in better communication and agreement between these two agents, and perhaps an increase in the overall effectiveness of their contribution to the overall COA.⁴

While we assume that agents can communicate with each other, in our use case, we are mostly concerned with the agents providing input to the “watcher” agent, who is then responsible for coordinating the input. In this use case, the main role of the watcher agent is to maintain coherence; and the main goal is to collect COA input from each agent and select what COA should be executed. In our example, the “watcher” agent can use constraint satisfaction techniques and a coherence mechanism to construct and select a COA for simulation and/or execution. Over time, the “watcher” agent will also develop a history of its experience in the task and with the members of the group. This experience will likely bias the watcher’s expectation of agent behavior and contributions in subsequent interactions.

4.1.1 Coherence

The theory of coherence is useful for integrating different pieces together to make a whole that obeys policies, rules and/or general constraints, e.g., temporal dependencies. In our use case, this could involve the integration of sub plans as they are created by each agent to address different aspects of the disaster relief problem. This could also involve the selection of alternative plans, each created from a different perspective and historical experience. The theory of coherence is best explained by the work of Thagard [2002].

⁴ Note: How long ago the role was experienced by each agent may be of relevance. If it was more than some amount of time, we assume that some of the current tools and practices may be novel to the person who participated in the role in the past. For example, technological tools change, so the current agent might choose to accept information from blog data (a media potentially not available to the less current agent). This may require the agent who uses the blog data to explain the usage of blog data and provide rationale on why it is relevant to the other actors who have not had this experience. Note: certain “personality” factors may also influence how successful the agent is at introducing and proving the usage of new technology. Some of these characteristics are being explored in more detail in the ELICIT experiment [Anderson, et. al, 2009].

There are several ways to compute coherence. One method is to utilize coherence graphs as described by Joseph et. al. [2008]. A coherence graph can be used to describe the relations between nodes of a graph. In our example, a node could be the COA input and/or plan contribution from a single planning agent. Some nodes will have more weight in terms of believability. For example, the COA input from the USAID planning agent could have the most weight in a disaster relief COA in the initial stages of the relief effort, especially if that agent has current experience with the country in need.

Our approach to coherence maintenance is focused on the issue of how to reconcile the different personal and cultural perspectives of multiple agents in order to make a single group decision or product such as a plan of action. In order to evaluate the usefulness of a coherence approach, we designed a coherence mechanism based on the *constraint* approach of Holyoak and Thagard [1989]. We also experimented with the usage of coherence graphs to support the development of a coherent COA. Several researchers [Pasquier et. al. 2003] have attempted to unify the theory of coherence with the Belief, Desire and Intention (BDI) construct, especially as an architecture for multi-agent systems; in DEAR, BDI constructs were reviewed but not used.⁵

Initially, we reasoned that a coherence mechanism would only be needed for the “watcher” agent to support the integration of input from multiple actors. As our research progressed, we decided that each actor should have local coherence mechanisms to guarantee that their contribution to the overall problem is rational and useful. In Ford and Mulvehill [2009] we describe how individual agent history and culture can influence experience-based reasoning. In that paper, we compare our approaches to traditional experience-based reasoning techniques (e.g., CBR, analogical reasoning and episodic reasoning), and highlight important issues and future directions that we believe are necessary to better understand social decision making. Results from an experiment are also presented in the paper that describe how shared problem solving presents issues of consensus and buy-in for various stakeholders and how a constraint based coherence mechanism can be used to ensure coherence in joint action among a group of agents with improved performance. More recent work is being done to quantify the personalities of the agents to determine how those factors influence social interaction and COA development. Additionally, our approach advocates that integration coherence (watcher) agent can develop its own case base of coherence experiences so that, over time, this experience will influence how it selects and merges contributions from individual actors.

⁵ For a good review of BDI and coherence theory see Joseph and Prakken [2009]. In their paper, they describe their research in which two coherence driven agents aim to reach agreement. The authors describe how individual goals and social goals affect the interaction between the agents. Their work provides a method for representing some of the hierarchical model for a given agent as expressed by Mantovani. In their paper, the authors also describe how protocol or rules can govern agent interaction. They introduce the notion of a moderator – a construct analogous to our “watcher” agent.

4.1.2 Collaboration and Negotiation

In the use case presented in this paper, each agent provides input to the watcher agent who can use negotiation and coherence maintenance techniques to resolve problems from divergent input. Alternatively, each planning agent could collaborate and negotiate with each other in order to generate their input to the watcher agent (like the behavior in cake-drug trafficking experiment). This could possibly result in a more coherent product and might require less work by the watcher agent. Research described in a paper by Shum [1997] describes some of the aspects of group problem solving. Shum specifies that “memories are not simply retrieved according to some database model - they are *reconstructed* in the context of who is asking and for what purpose. Knowledge is constructed to serve particular needs at a particular time.” Shum’s research offers an approach for capturing the memory that is created by the group. Shum’s research coined the term *organizational memory* to represent group memories and Shum conjectures that knowledge management approaches may be useful for storing and leveraging this type of memory. The work of Shum also emphasizes the importance of negotiation and argumentation by the members of the group. Shum identifies certain factors, similar to the ideas of Mantovani about how goals, assumptions and the agendas of each of the team members can influence the group decision making process. In their paper they also introduce a representation language that was developed by Lee [1990] for representing an argument network. For DEEP researchers, the work of Shum and the representation language of Lee could be useful for defining mechanisms for collaboration and for specifying a memory for the watcher agent.

The research of Giampapa and Sycara [2006] also offers an approach for managing collaboration among a group of agents. Their work focuses on the negotiation aspects of collaboration. In their paper, they define a negotiation strategy as a mapping from input information about the environment to a sequence of decisions or strategy that specifies the action at each step as it is conditional on the negotiation history. For DEEP, this would imply that the *watcher* agent maintains a negotiation history and/or that each of the planning agents maintain negotiation histories. Giampapa and Sycara [2006] propose a negotiation model that could be implemented in DEEP composed of three modules: “*single-threaded negotiations; synchronized multi-threaded negotiations; and dynamic multi-threaded negotiations. The single-threaded negotiation model provides negotiation strategies without specifically considering outside options. The model of synchronized multi-threaded negotiations builds on the single-threaded negotiation model and considers the presence of concurrently existing outside options. In a synchronized multi-threaded negotiation process a negotiator participates in multiple bilateral negotiation threads with different, simultaneous negotiation opponents. The negotiator can reach an agreement in at most one of these threads and is aware of all the threads at the beginning of the process. Dynamic multi-threaded negotiations build on the synchronized multi-threaded model but introduce uncertainty to the threads.*”

Other work by Shintani, Ito and Sycara [2000] describes how negotiation can be implemented in a multi agent system. In their approach, the social aspects of negotiation and the tradeoff between reaching consensus and maximizing one's own expected payoff are considered. In the work by Huang and Sycara [2002], methods for representing the subjective belief and personalities among the agents in group decision making are provided.

The importance of collaboration has also been studied in the CBR community. Research in collaborative CBR (C-CBR) [McGinty and Smyth 2001] provides some solutions for how a collection of homogeneous CBR agents, each having the same CBR capabilities but differing in their problem solving experience, can interact to solve problems, e.g., generate a route. The paper by Plaza and McGinty [2005] provides an overview of this work.

Many of the issues about collaboration and negotiation that have been presented in this section are reinforced by the work of Gasser [1991]. While some of the research we have cited focuses on either software agents or human agents, the work of Gasser [1991] focuses specifically on issues of collaboration between both human and artificial agents within a distributed multi-agent framework. Because this perspective resonates with the requirements of DEEP, Table 1 highlights the issues that we believe are of most relevance for DEEP. Note: the page numbers reference pages in the Gasser document.

Table 1. Summary of Gasser Research on Collaboration and Coherence

Guidance and/or Open Issue	Description	Page
Coherence among collaborating agents	"How does one ensure that agents act coherently in making decisions or taking actions, accommodating the non-local effects of local decisions and avoiding harmful interactions?"	p. 110
Coordination	"...to coordinate their actions, intelligent agents need to represent and reason about the knowledge, actions, and plans of other agents."	p. 109
Ignoring or removing an agent's influence	"Assuming that there was a need, how would one go about removing the influence of a particular agent in a situation, e.g., disabling it or discrediting its knowledge)"	p. 115
Predictability	"...in an asynchronous and open distributed system, no message can be guaranteed to lead to the same set of behaviors twice"	p. 116
Episodic history	"Action is a particular commitment to doing things in a particular way, conditional upon the actor's particular knowledge of the situation."	p. 116
Multiple perspectives	"Multiple perspectives are a fundamental feature of any multi-agent system, simply by virtue of differing commitment histories and local circumstances....If multiple perspectives are basic, disparities in perspectives are an issue...moreover, multiplicity of perspectives raises the issue of the impossibility of global conceptions."	p. 117
Compromise	"In some applications, e.g., nuclear seismic analysis, it is essential that each node avoid compromising its local set of beliefs and assumptions by integrating faulty or malicious messages from other sensing nodes. Each node must maintain local autonomy and arms-length relationships while incorporating useful information generated by others."	p. 125
Global coherence	"Global coherence should be conceptualized as the situated outcome of a negotiation – as long as agents collectively reach agreement, and agree that they have, their actions are coherent."	p. 126
Information Transfer across Contexts	"Transporting representations raises problems of completeness and of interpretation in a new context (e.g., transporting a concept may strip it of its context and render it un-combinable)."	p. 127
Resource balancing	"Include a collection of checks and balances (plurality) in a DAI system, so that different participants have control over different resources in critical interactions and no participant can be ignored..."	p. 133

4.2 Analogical reasoning

Analogical reasoning is a very powerful approach for solving problems and for learning. An overview of analogical reasoning is provided in Mulvehill, Deutsch and Rager [2007]. In that paper, we conclude that analogies are often used in educational settings to help teach new concepts [Gentner & Stevens, 1983; Goswami, 2001]. For DEEP, the goal of analogy is also educational. Each of the planning agents can use analogies of their own or published historical military scenarios to support the analysis of the current scenario. If the analogy is appropriate, it may help to educate or prompt the planner about the nuances of the current problem and to predict possible behaviors.

Analogies are typically conceived of as involving a mapping between two domains called the source (base) and the target. Hall [1989] lists four abstract processes that are widely considered to be necessary for analogical reasoning:

1. recognition of a source, given a target description;
2. elaboration and evaluation of the mapping between the two;
3. transfer of information from the source to the target;
4. consolidation (i.e., learning) of the outcome.

In order to perform the mapping, some type of classification about the domain must be available. For example, in order to determine if a *hurricane* is similar to an *earthquake* one can consider these as subclasses of a more abstract class like *environmental catastrophe*. Analogical reasoning researchers like Gentner and Stevens (1983) caution that “the inheritance of characteristics in analogies is only partial.” For planners in the use case presented in this paper, the overlapping similarities between a hurricane and an earthquake can be useful for quickly determining how to respond to an environmental catastrophe in general but will need to be refined to support the specific type of disaster.

Analogies can also help provide missing information or fill gaps in a developing plan. However, recent research by Kokinov [2009] cautions that analogies can be false and perhaps unreasonable. Imagine a false analogy and a strong personality in a collaborative session. How would the false path of reasoning be discovered, and could the coherence mechanisms of a watcher agent be used to correct the errors? Analogies can also lead to distractions. Therefore, any design of a watcher agent for DEEP will need to have mechanisms to help the agent maintain focus on the problem solving context.

Regardless of the potential pitfalls of analogical reasoning, analogies are powerful and are commonly used by humans to support problem solving. Analogical reasoning paradigms have been utilized by many researchers in the development of systems and have often been combined with CBR and/or episodic reasoning techniques. For a review of research in analogical reasoning see Keane [1988] and French [2002]. Table 2 presents a sampling of research in analogical reasoning. The first column indicates the time frame that the research was conducted, the 2nd column describes what the research is, the 3rd column is the name of a program or main theme of the research and the last column presents the name of a researcher commonly associated with the work. The table is not meant to be all inclusive. Many of the entries in the table are described in more detail in French [2002].

In Table 2, four analogical reasoning approaches are marked with a * to indicate that they satisfy some of the primary objectives of this project (see section 1.0) and are candidates for further investigation by DEEP researchers. These include the work of Holyoak and Thagard [1989] on ACME; the work of Forbus, Gentner and Law [1995] on similarity based retrieval (MAC/FAC); the work of Hummel and Holyoak [2005] on the LISA system; and the work of Kokinov [2001] on the AMBR system. A brief summary of each of these systems is provided here with a pointer to a reference for additional information.

Table 2. Summary of Analogical Reasoning Research

Year	Research Area	Program Name	Prime Researcher
1965	Analogy Making	Argus	Reitman
1968	Geometric figure analogies	analogy	Evans
1969	Cognitive and computational analogy with Long Term Memory (LTM)	JCM	Becker
1978	Mapping between source and target	Transfer Frames	Winston
1979	LTM Knowledge stored as rules	ANA	McDermott
1983	Learning by analogy using means ends analysis		Carbonell
1983	Modeling of Analogy Making	Structure Mapping Theory (SMT)	Gentner
1984	Agent based approach to analogy making	Copycat	Hofstadter
1985	CBR and Analogy Making	Mediator	Simpson
Year	Research Area	Program Name	Prime Researcher
1986	Concept Formation	CARL	Burstein
1989	Analogy making with parallel activation of a neural network	ACME*	Holyoak and Thagard
1989	Computational implementation of SMT	Structure Mapping Engine (SME)	Forbus
1990	Learning by analogy with larger domains	Prodigy/Analogy	Veloso and Carbonell
1991	Analogical Retrieval Engine	MAC/FAC*	Gentner and Forbus
1993	Using Analogies for metaphor understanding	ACT-R	Anderson
1997	Connectionist Model of analogy making	LISA*	Holyoak
2000	Micro Agent Architecture for episodic re-collection	AMBR*	Kokinov

- **ACME:** Provides an architecture for analogy making that is implemented in a neural network like structure. With ACME, constraints are set up between the source and the target and excitatory and inhibitory links are used as weights to implement constraints. The goal is to maximize the coherence of the constraints to produce the most coherent analogies. For more information see: Holyoak and Thagard [1989].

- **MAC/FAC:** Provides a similarity based method that is based on psychological models providing reminders about structural similarities. MAC stands for *Many Are Called* and FAC stands for *Few Are Chosen*. In this system, an analogy is a match based on a common system of relationships. MAC/FAC employs a two stage process where a set of likely candidates are initially chosen, then a more computationally expensive method is used to find the more likely candidate. The reminders are generally based on surface similarity rather than on structural similarities between the possible solutions. MAC/FAC has been used as a front end to an implementation of Gentner's structure mapping theory called the structure mapping engine (SME). For more information see: French [2002] and Forbus, Gentner, and Law [1995].
- **LISA:** A connectionist model with primitives that are connected in a neural network. Weights in the neural networks are used to facilitate learning and inference. LISA uses a guided pattern recognition algorithm to support analogical retrieval, mapping and inference. According to the researchers "a fundamental aspect of human intelligence is the ability to acquire and manipulate concepts defined by systematic *relationships* among multiple objects....relational thinking involves the ability to see analogies between superficially disparate situations and to form more general schemas or relationally defined concepts." For more information see: Hummel and Holyoak, K. J. [2005].
- **AMBR:** Provides an architecture for facilitating collective memory. It uses an approach to support the blending of episodes in support of episodic reasoning. AMBR also has mechanisms that address partial memory and memory distortions. "The mapping process in AMBR does not start after the old episode is retrieved but runs in parallel to it. This makes it possible for the already established partial mapping to guide the episode construction in such a way that the old episode is reconstructed in directions which allow better alignment between the base and the target." [Kokinov, 2001] For more information see: Kokinov and Zareva-Toncheva [2001].

Lessons learned or approaches utilized in these systems could benefit DEEP researchers. For example, problems with merging contributions from multiple case bases will be an issue in DEEP, and the approach used in AMBR could be applied to support this process. Approaches for supporting learning and feature mapping can be derived from the work of LISA and MAC/FAC. Issues in developing an architecture that supports analogical reasoning can be derived from the work done on ACME.

One of the key problems in analogical reasoning within a distributed system is the issue of balancing individual and group optimality [Gasser, 1991]. In DEEP, we anticipate that groups of agents will be able to use analogies derived from the past experiences of one or more agents to solve problems together. The research of Gasser indicates that effective shared analogical problem solving will increase as stronger relationships among the agents evolve. Also, according to Carley [1991] "as agents interact, they acquire information that changes the way they perceive the world, their actions, attitudes, and beliefs. Two agents are more likely to interact if they both believe that they are more similar to each other than they are to others in the group".

4.3 Distributed CBR Management

A review of the literature on distributed CBR is provided in Mulvehill, Deutsch and Rager [2007]. Because DEEP allows for multiple case bases to be utilized for constructing a COA, our research in this project has been focused on case base management and/or maintenance, and case base selection (from local versus remote case bases). Some insights into case base management, especially maintenance issues are provided in Leake and Wilson [1999]. For issues on problems associated with case selection across a distributed set of case bases see Leake & Sooriamurthi [2001, 2002].

Case selection becomes more difficult when analogies are used because the mapping between the current problem context and an analogy is highly dependent on the determination of features in both the source and the target domain descriptions. A historical description of how this problem has been approached by many researchers is described in French [2002].

Trusting the accuracy of any given case base is also an important issue to consider in case base maintenance. The research of Keil [2008] offers some insights into misperceptions that actors often have on the accuracy of their own problem solving accuracy and into how likely one actor is to use the historical experience of another actor. According to Kiel *“knowing when we need to defer to another’s expertise is intimately related to how complex we think the phenomenon being explained is. If the phenomenon seems trivial and relatively self-evident to any reasonable observant and thoughtful person, then it may be inappropriate to bring in an expert who might only muddy the waters rather than shed insight”*. Leake and Wilson [1999] also address issues with incorrect data in case bases.

5. Recommendations

The research conducted during this project was driven by the following three objectives:

1. Identify, develop and recommend technology necessary to leverage case-based reasoning (CBR) technology and analogical reasoning in a distributed environment.
2. Identify and evaluate issues constraining the employment of CBR technology in a distributed environment.
3. Provide design and architectural recommendations, with an emphasis on how different subjective perspectives can influence the annotation, representation and usage of one or more episodes as they are evolved by multiple participants within a net-centric planning and execution environment

Throughout this report recommendations related to each of these objectives have been provided. The recommendations identified both *questions to answer* and *things to do*.

In this section, several tables are provided to highlight research issues and to present certain questions that must be answered in for order the DEEP architecture and framework to better support CBR, episodic reasoning and analogical reasoning in a distributed setting and to leverage the contributions of agents with differing perspectives. The next section concludes with recommendations on things to do, including a list of two research topics that are potential candidates for continued research.

Table 3 presents the issues that need to be addressed in order to leverage analogical reasoning.

Table 3. Analogical Reasoning Issues and Recommendations

Analogy formation	In order to exploit the use of analogical reasoning in DEEP, issues associated with analogy formation and usage need to be explored. The research of Lee et. al. [2009], Dehghani et. al. [2009] and Kokinov et. al. [2009] provides solutions and issues to be considered.
Similarity Functions	Conceptual tool sets and similarity functions are needed to support analogy formation and, in turn, prediction generation within a simulator. The research of Forbus et. al. [1995] offer recommendations for how to compute similarity.
Features and Ontologies	A strategy for examining how features specified by the operator may be tied to ontologies to support analogical reasoning is important for comparing analogies.

Table 4 presents issues that need to be addressed in order to use and maintain historical data embodied as cases or episodes in a distributed environment.

Table 4. Distributed CBR Management Issues

Shared Episodes	Methods and/or tools that track “shared” episodes with the perspectives from each actor are required to support the integration of data from multiple distributed case repositories. Maintaining reference to “self” is important.
Multiple Case Base Usage	Reasons to prefer multiple case bases over a single monolithic case base are documented in Leake & Sooriamurthi [2002]; and d’Aquin, Lieber, & Napoli [2005].
Case Base Archiving	Some form of archiving or compression of episodes is required in an operational system.
Selective Forgetting	Some form of selective forgetting (or case deletion) will be required, especially if the data source is found to be untrustworthy or corrupt.

Table 5 lists issues that need to be examined closely in order for DEEP to benefit from multiple perspectives.

Table 5. Multiple Actors and Case Base Issues

Perspective	The notion of perspective may become extremely critical for the DEEP program, especially in the predictive area as different perspectives by actors will offer different futures of a situation. How the multiple perspectives are blended is a topic of research and some of the work by Kokinov [2001] may offer insights. Methods to prevent one perspective from dominating a situation are also of importance. Research in social groups (Prietula and Carley [1999]) may offer insights.
Self	One of the most critical aspects of episodic reasoning is the notion that an episode is associated with an actor, from a "self" perspective. In a distributed environment, where multiple episodes are used to create a single product, some mechanism needs to be developed to keep track of how episodes from the multiple sources are related, each annotated from the perspective of the actor.
Ontologies	The use of global versus multiple ontologies to support the use of multiple case bases is important for supporting integration of cases from different actors. Relevant research includes: Bouquet, Giunchiglia, van Harmelen, Serafini, & Stuckenschmidt [2004] and d'Aquin, Lieber, & Napoli [2005].
Case Selection and Merging	In a distributed case base environment, case base selection and merging (from local versus remote case bases) becomes an important issue. Issues associated with this are documented in Leake & Sooriamurthi, [2001].

6. Conclusion

The content for a typical episodic reasoning system is the recorded experiences of a particular operator, e.g., personal memory or mental model. Episodic reasoning systems can also store the *collected* experiences from a set of operators who are performing a given operator task, e.g., a lessons learned document or some protocol that has been developed as a result of multiple experiences by one or multiple actors. In the DEEP project there is a requirement to leverage recorded experiences that are created by a set of distributed users from a variety of work positions, with varied goals and perspectives across time (historical, current, and projected).

In this report, two methods have been proposed to support the construction of a COA with the use of multiple case bases. One method involves a group of problem solvers who interact with each other to *form* the COA. As an example, consider how Wikipedia pages are constructed. The other method is more formal and involves an *integration agent* that leverages coherence maintenance capabilities to merge input from each of the participants into a ‘coherent’ COA.

With either approach, each actor is responsible for episodic retrieval and construction of a useful solution to the current problem. Each actor should also be able to apply coherence maintenance routines locally in order to guarantee that a contribution is contextually relevant and complete. When the integration (watcher) agent is responsible for the construction of the aggregated solution, coherence mechanisms should be used to explicitly support the construction of the output. The coherence mechanism must provide a model, protocol, rules, constraints or some problem sensitive workflow to guide and constrain case contribution merging or blending such that the result is “coherent”, relevant, and in a useful form (representation) for simulation and/or execution.

This process will be complicated if the representations used by the participating actors are not similar. Ontological mapping tools may be used to resolve some inconsistencies and to handle mis-matches across case features and/or plan schema fragments from contributing actors. Additionally, the COA (product) of the coherence agent should adhere to some formal representation – optionally a representation that can be used in the simulation and/or execution environment.

Since the entire system is comprised of a group of agents (both software and human) that have roles and responsibilities in a problem solving context, a model, such as the one offered by Mantovani [1996] could be used to represent the motives of each actor. Since each actor has social characteristics that influence actor to actor interactions, certain social factors need to be evaluated to determine how they influence the overall process. The work of Prietula and Carley [1999] may provide a good starting position for continued research on these issues.

The research reviewed in this effort indicates that as a group continues to interact together their ability to “form” outputs and/or “construct” collective memories will increase. This tendency will likely affect the behavior of each contributing agent and of the integration (watcher) agent. Research by Gasser [1991] on distributed artificial intelligence (DAI) may serve as guidelines for DEEP implementation.

New computing methods need to be specified and defined in order to resolve some of the problems introduced by collective memories and by the requirement to support reasoning by agents with multiple perspectives across past, present and future. Our research indicates that in a distributed environment, each observer's perception/analogy pair should be a function of the role and goals of that observer. Figure 3 presents a suggested design and development approach for DEEP developers to consider as they evolve the DEEP technology. This approach highlights four required capabilities: case base usage and maintenance; publish and subscribe behavior; perspective integration; and learning.

1. Case Base Usage and Maintenance Each actor creates, evolves and maintains one or more case bases. - An actor may partition their cumulative experiences into multiple case bases, each associated with a particular subject, e.g., eating at restaurants, job assignment role, etc. - The case base(s) can be used by the creating actor to solve similar problems, or can be used by other actors. - Other actors may use the case base of a particular actor but they cannot write back to the source case base, instead they write into their own case base(s) or into some common "collective" case base.	2. Publish and Subscribe In group problem solving, each actor offers (publishes) case base data that is relevant to the current problem. - The actor may allow other actors to access or subscribe to its case base or the actor may create a case base for others to share (similar to a database view). - Integration of data from multiple case bases is the responsibility of the retrieving actor.
3. Managing Perspectives Multiple actors can present solutions to a current problem from their perspective (and associated bias). The actors may also interact with each other, so social issues and social predispositions become relevant. - Multiple perspectives can be integrated by a coherence agent (watcher) to form the product. - The coherence agent uses a domain problem model and its own episodic memory to interpret and/or evaluate input. - The result is a COA that will likely need to be refined (details added) and evaluated (via some simulation method and/or plan critic tool) before it is considered "coherent".	4. Using Feedback to Learn Feedback is used to update each actor's case base. - The COA can be evaluated through simulation or real time execution. - Execution feedback should be provided to all of the participating actors and the coherence agent for use in evaluating usefulness/effectiveness and influencing subsequent retrievals, e.g., do this again; don't do this, do this in some particular order.

Figure 3. DEEP Design and Development Approach

During the course of this project two research issues (learning and trust) were identified that are considered candidates for future research. Listed here are some questions for each area that are of relevance for the DEEP research and development:

Learning

- How would a coherence agent learn?
- Are there different types of coherence agents, e.g., integrator, checker, critic?
- What should be stored in the episodic memory of a coherence agent?
- Should the coherence agent of a specific type have its own goal hierarchy as a way to maintain coherence and check constraints?

Trust

- How should trust be represented and/or computed in DEEP?
- Should each actor be profiled?
- Can trust change over time and per problem, e.g., some actors are more trustworthy in a particular problem and role? This could be related to personality types such as those described by Prietula and Carley [1999].

Several papers are recommended in Appendix A that may help in answering some of these questions.

Additionally, a paper by Ford and Mulvehill [2009] was written during the course of this research project that integrates a lot of what was learned during this project and is considered a candidate for continued research. In particular, the paper describes a coherence framework that could be implemented in DEEP. The framework allows for the sharing of experiences and analogies by multiple actors. Each agent is responsible for individually maintaining coherence with its own experiences and the problem at hand. Socially, the agents attempt to maintain coherence with each other within a particular problem solving context. Our research assumes that the group can establish norms and expectations, and deliberate about when those norms should change. As the group more closely coheres and relationships strengthen, our research to date indicates that increased trust may lead to the presentation and possible use of more abstract analogies and more liberal changes in norms. In subsequent research, we have begun to model the personalities of the individual actors to determine how those personalities influence the construction of the COA and how they should be evaluated by the watcher (integration) agent during the construction and revision of a COA.

7. References

- Aamodt, A. & E. Plaza, E. (1994). Case-based reasoning: Foundational issues, methodological variations, and system approaches. *Artificial Intelligence Communications*, 7, 39-59.
- Aleman-Meza, B., Halaschek, C., Arpinar, I. B., (2005). Collective Knowledge Composition in a Peer-to-Peer Network, In: Encyclopedia of Database Technologies and Applications, Laura C. Rivero, Jorge H. Doorn and Viviana E. Ferragline (Eds), Idea-Group Inc., April 2005.
- Anderson, C., Wynn, D., McEver, J., and Ruddy, M. (2009). Exploring the Effects of Individual Characteristics on Organizational Performance using the ELICIT Experimentation Platform". ICCRTS-09.
- Bouquet, P., Giunchiglia, F., van Harmelen, F., Serafini, L., & Stuckenschmidt, H. (2004). Contextualizing Ontologies. *Journal of Web Semantics* 1, 1–19.
- Conway, Martin A. (2005). Memory and the Self. *Journal of Memory and Language*, Vol. 53, pp. 594 – 628.
- d'Aquin, M., Lieber, J., and Napoli, A. (2005). Decentralized case-based reasoning for the semantic web. In Y. Gil, E. Motta, V. R. Benjamins, & M. A. Musen (Eds.), *The semantic web – ISWC 2005* (pp. 142-155). Berlin: Springer-Verlag.
- Dehghani, M., Gentner, D., Forbus, K. Ekhtiari, H., and Sachdeva, S. (2009). "Analogy and Moral Decision Making". *Analogy* 09, July 24 – 27, 2009, Sofia, Bulgaria, <http://nbu.bg/cogs/analogy09/proceedings/>
- Forbus, K. D., Gentner, D. and Law, K. (1995). MAC/FAC: a model of similarity-based retrieval. *Cognitive Science*, Vol. 19, pp. 141-205.
- Ford, A. J., and Carozzoni, J. (2007). "Creating and Capturing Expertise in Mixed-initiative Planning. ICCRTS-07,
- Ford, A. J., and Lawton, J. H. (2008). Synthesizing Disparate Experiences in Episodic Planning. *Proceedings of the 13th ICCRTS: C2 for Complex Endeavors*.
- Ford, Anthony J. and Mulvehill, Alice M. (2009). Collaborating with Multiple Distributed Perspectives and Memories. *Proceedings of the Second International Workshop on Social Computing, Behavioral Modeling and Prediction*, Phoenix, Arizona.
- French, Robert M., (2002). "The computational Modeling of Analogy-Making", *Trends in Cognitive Science*, Vol. 6, No. 5.
- Gasser, Les, (1991). Social Conceptions of Knowledge and Action: DAI Foundations and Open Systems Semantics, *Artificial Intelligence*, Vol. 47, pp. 107 – 138, Elsevier.
- Gentner, D. & Stevens, A. L. (1983). Introduction. In D. Gentner & A. L. Stevens (Eds.), *Mental models*, (pp. 1-6). Hillsdale, NJ: Lawrence Erlbaum.
- Giampapa, Li. J., and Sycara, K. (2006). Bilateral Negotiation Decisions with Uncertain Dynamic Outside Options, *IEEE Transactions on Systems, Man, and Cybernetics*.
- Goswami, U. (2001). Analogical reasoning in children. In D. Gentner, K. J. Holyoak, & B. N. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science*, (pp. 437-470). Cambridge, MA: The MIT Press.
- Hall, R. (1989). Computational approaches to analogical reasoning: a comparative analysis. *Artificial Intelligence*, Vol. 39, pp. 39 – 120.
- Holyoak, K. J., and Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, vol 13, pp. 295-355.

- Huang, P. and Sycara, K. (2002). A Computational Model for Online Agent Negotiation". *Proceedings of the 35th Hawaii International Conference on System Sciences*.
- Hummel, J. E., and Holyoak, K. J. (2005). Relational Reasoning in a Neurally Plausible Cognitive Architecture, An Overview of the LISA project. American Psychological Society, Vol. 14, No. 3.
- Joseph, S., and Prakken, H. (2009). "Coherence-Driven Argumentation to Norm Consensus". *ICAIL-09*, Barcelona, Spain.
- Joseph, S., Sierra, C. and Schorlemmer, M. (2008). Coherence Based Framework for Institutional Agents, by Coordination, Organizations, Institutions, and Norms in Agent Systems III, COIN 2007 International Workshops. Revised Selected Papers, Springer, Volume 4870, p.287-300.
- Keane, M., (1988). Constraints on Analogical Mapping: A Comparison of Three Models <http://www.scss.tcd.ie/publications/tech-reports/reports.93/TCD-CS-93-24.pdf>.
- Keil, Frank C. (2008). Getting to the Truth, Grounding Incomplete Knowledge". *Brooklyn Law Review*, Vol 73:3.
- Klein, G. (1998). Sources of power: How people make decisions. Cambridge, MA: The MIT Press.
- Kokinov, B., Feldman, V., and Petkov, G. (2009). "Analogy-Making Automatically Produces False Memories For Both Mapped Situations". *Analogy 09*, July 24 – 27, 2009, Sofia, Bulgaria, <http://nbu.bg/cogs/analogy09/proceedings/>
- Kokinov, B., Zareva-Toncheva, N., (2001). "Episodic Blending as Result of Analogical Problem Solving". *Proceedings of the 23rd Annual conference of the Cognitive Science Society*, London: Erlbaum, pp. 510-515.
- Leake, D. & Wilson, D. (1999). "When Experience is Wrong: Examining CBR for Changing Tasks and Environments". *Proceedings of ICCBR-99*.
- Leake, D. B. & Sooriamurthi, R. (2001). When two case bases are better than one: Exploiting multiple case bases. In *Case-Based Reasoning Research and Development: Proceedings of the Fourth International Conference on Case-Based Reasoning, ICCBR-01*. Berlin: Springer-Verlag.
- Leake, D. B. & Sooriamurthi, R. (2002). Managing multiple case bases: Dimensions and issues. *Proceedings of the Fifteenth International Florida Artificial Intelligence Research Society (FLAIRS) Conference*, pp. 106-110. Menlo Park, CA: AAAI Press.
- Lee J. (1990). "SIBYL: A Tool for Managing Group Design Rationale". *Computer Supported Cooperative Work*. Los Angeles, CA, ACM Press, New York, pp. 79-92.
- Lee, H. S., Holyoak, K. J., and Lu, H. (2009). Integrating Analogical Inference with Bayesian Causal Models". *Analogy 09*, July 24 – 27, 2009, Sofia, Bulgaria, <http://nbu.bg/cogs/analogy09/proceedings/>
- Mantovani, Giuseppe, (1996). Social Context in HCI: A New Framework for Mental Models, Cooperation, and Communication, *Cognitive Science*, Vo. 20, No. 2, pp. 237 – 270, Ablex Publishing Corporation, NJ.
- McGinty, L., and Smyth, B. (2002). "Shared Experiences in Personalized Routh Planning", in *FLAIRS-02 Proceedings*.
- Mulvehill, A., Deutsch, S., and Rager, D. (2007). "Case Based Reasoning for DEEP: Observations and Recommendations, BBN Report 8474, Cambridge MA.

- Padovitz, A., Zaslavsky, A, Loke, S. (2006). "Merging Context Perspectives: An Approach to Adaptive Agent Reasoning in Pervasive Computing Systems", Proceedings of the Fourth Annual IEEE International Conference on Pervasive Computing and Communications Workshops.
- Paul, S. (2001). Collective Memory Support in Negotiation: A theoretical Framework, Proceedings of the 34th Hawaii International Conference on System Sciences.
- Pasquier, Phillipe, Chaibe, draa-Brahim. (2003). "The Cognitive Coherence Approach for Agent Communication Pragmatics" in AAMAS'03.
- Plaza, E. & McGinty, L. (2005). Distributed case-based reasoning. *The Knowledge Engineering Review*, 20, 261-265.
- Prietula, M., and Carley, K. (1999). Exploring the Effects of Agent Trust and Benevolence in a Simulated Organizational Task". *Applied AI*, 13, 321-338.
- Rettberg, Scott. All Together Now: Collective Knowledge, Collective Narratives, and Architectures of Participation". (2005). retts.net/documents/cnarrativeDAC.pdf
- Schacter, D. (2001). *The Seven Sins of Memory*, Houghton Mifflin.
- Shank, R. (1982). *Dynamic Memory: A theory of Reminding and Learning in computers and People*. Cambridge University Press.
- Shintani, T., Ito, T., and Sycara, K. (2000). Multiple Negotiations among Agents for a Distributed Meeting Scheduler". *IEEE*.
- Shum, Simon Buckingham. (1997). Negotiating the Construction and Reconstruction of Organizational Memories". *Journal of Universal Computer Science*, Vol 3, No. 8.
- Struck, J. and Ganger, G. (2003). Human Organization Analogy for Self-* Systems, CMU, First Workshop on Algorithms and Architectures for Self-Managing Systems. In conjunction with Federated Computing Research Conference (FCRC). San Diego, CA. June 11, 2003. Supersedes Carnegie Mellon University SCS Technical Report CMU-CS-03-129, <http://www.pdl.cmu.edu/PDL-FTP/SelfStar/CMU-CS-03-129.pdf>
- Thagard, Paul. (2002). *Coherence in Thought and Action*, MIT Press.
- Weber, G. (1996). Episodic Learner Modeling. Vol. 20, No. 2, pp. 195-237, *Cognitive Science*.

8. Acronyms

ACME	Analogical Constraint Mapping Engine
AMBR	Associative Memory-Based Reasoning
BDI	Belief Desire Intention
C2	Command and Control
CBR	Case-base Reasoning
C-CBR	Collaborative CBR
COA	Course of Action
DAI	Distributed Artificial Intelligence
DEAR	Distributed Episodic and Analogical Reasoning
DEEP	Distributed Episodic Exploratory Planning
LISA	Learning and Inference with Schemas and Analogies
MAC/FAC	Many Are Called and Few Are Chosen
MAS	Multi Agent System
USAID	United States Agency for International Development

Appendix A. Additional Reviewed Papers

Argumentation

- Bex, F. J., Prakken, H., and Verheij, B. (2007). “Formalising Augumentative Story-based Analysis of Evidence”. *ICAAIL '07*, Palo Alto, CA, June 4 – 8, 2007.
- deMaat, E. and Winkels, R. (2009). A Next Step towards Automated Modeling of Sources of Law”. *ICAAIL 2009*, Barcelona, Spain.
- Li, C., Giampapa, J., Sycara, K. Bilateral Negotiation Decisions with Uncertain Dynamic Outside Options. Proceedings of the First IEEE International Workshop on Electronic Contracting (WEC'04), p.54-61, July 06-06, 2004 [doi>10.1109/WEC.2004.5]
- Shintani, T., Ito, T., and Sycara, K. (2000). Multiple Negotiations among Agents for a Distributed Meeting Scheduler”. *IEEE*.

Learning

- Bozsahin, H. C., and Findler, N. (1992). Memory-Based Hypothesis Formation: Heuristic Learning of Commonsense Causal Relations. Vol. 16, No. 4, pp. 431-454, *Cognitive Science*.
- Craw, S., Wiratunga, N., and Rowe, Ray, C. (2006). Learning Adaptation Knowledge to Improve Case Based Reasoning. *Artificial Intelligence*. Vol. 170, pp 1175 – 1192, Elsevier.
- Fu, Wai-Tat and Anderson, J. “Solving the Credit Assignment Problem: The Interaction of Explicit and Implicit learning with Internal and External State Information”. Proceedings of the 28th Annual Conference on the Cognitive Science Society, Hillsdale, NJ.
- Garalevicius, Saulius, (2005). Using Memory – Prediction Framework for Invariant Pattern Recognition”. Final Report, Temple University, December 2005.
- Garalevicius, Saulius. (2007) Memory – Prediction Framework for Pattern Recognition: Performance and Suitability of the Bayesian Model of Visual Cortex”. Temple University, AAAI-07.
- Gasser, L. and Wang, J., (2002). Mutual Online Concept Learning for Multiple Agents. *AAMAS02*, July 15-19, Bologna Italy.
- Ke, W., Mostafa, J., and Fu, Y. (2007). “Collaborative Classifier Agents: Studying the Impact of Learning in Distributed Document Classification”. *JDCL-07*, Vancouver, British Columbia, Canada.
- Landy, D. H., Taylor, E. G., (2009). Variation Among Individuals in How Structure Affects Similarity”. *Analogy 09*, July 24 – 27, 2009, Sofia, Bulgaria, <http://nbu.bg/cogs/analogy09/proceedings/>
- Macedo, L. and Cardoso, A. (2004). Exploration of Unknown Environments with Motivational Agents. *AAMAS'04*, July 19-23, 2004.
- Ortanon, S., and Plaza, E. (2007). “Learning and Joint Deliberation through Argumentation in Multi-Agent Systems”. *IFAAMAS-07*.

Memory

- Anand, P. and Sternthal, B. (1991). Perceptual Fluency and Affect Without Recognition”, by Memory and Cognition, Vol. 19-3, pp 293 – 300.
- Merikle, P. and Reingold, E. Measuring Unconscious Perceptual Processes”, in R. F. Bornstein and T. S. Pittman (Eds.), *Perception without Awareness*, pp. 55 – 80, New York: Guilford Press.
- Stanfill, C., and Waltz, D. (1986). Toward Memory-Based Reasoning”. *ACM*.

Xia, Q. and Rao, M. (1999). Dynamic case-based reasoning for process operation support systems. *Engineering Applications of Artificial Intelligence*, Vol. 12, pp 343 – 361.

Social and/or Cultural Issues

Bernard, M., Backus, G., Glickman, M., Gieseler, C., and Waymire, R. (2009). Modeling Populations of Interest in Order to Simulate Cultural Response to Influence Activities”. *Social Computing and Behavioral Modeling*, by H. Liu et al, (eds), pp. 42 – 49, Springer Science.

Latour, B., and Woolgar, S. (1979). *Laboratory Life: The Social Construction of Scientific Facts*, Sage, Beverly Hills, CA.

Zackland, Manual. (2003). *Communities of Action: A Cognitive and Social Approach to the Design of CSCW Systems*. GROUP’03, November 9-12, 2003, Sanibel Island, Florida, USA.

Trust

Chan, K., Ivanic, N., Rivera, B., and Bowman, E. (2009). “A General Framework of Human Trust in Networks”. ICCRTS-09.

Cho, J., and Swami, A. (2009). Towards Trust-based Cognitive Networks: A Survey of Trust Management for Mobile Ad Hoc Networks”. ICCRTS-09.

Other

Plan Critics

Argall, B., Browning, B., and Veloso, M. (2007). Learning by Demonstration with Critique from a Human Teacher”. HRI’07, Arlington VA.

Ontologies

Colomb, R., and Dampney, C. N. G., (2005). An approach to ontology for institutional facts in the semantic web. *Information and Software Technology*, Vol. 47, pp. 775-783.