13th ICCRTS: C2 for Complex Endeavors

"Synthesizing Disparate Experiences in Episodic Planning"

Track 4: Cognitive and Social Issues

Authors

Anthony J. Ford

Dr. James H. Lawton

Point of Contact:

Anthony J. Ford
Air Force Research Laboratory Information Directorate
Information Systems Research (RISB)
525 Brooks Road, Rome, NY 13441-4505
315-330-2610
Anthony.Ford @rl.af.mil

Distribution Statement A: Approved for Public Release

Synthesizing Disparate Experiences in Episodic Planning

Anthony J. Ford, Dr. James H. Lawton

Abstract

Many decisions are actually made by synthesizing previous experience. Often, this involves many different experiences coming together to form a feasible solution. This paper presents a statistical model for predicting the outcome of solutions based on multiple experiences. In edge organizations, such as emergency first responders, it often requires the expertise of more than one person to form an approach to a complex problem. Unfortunately, each planner only has access to his or her own memories. We propose to use an artificial intelligence decision aide to help bridge this gap, by reasoning over distributed collections of previous experiences. The key research questions that we address include: How can an artificial reasoner form a plan based on several disparate experiences from different sources? How can we gauge the potential efficacy of such a plan? How can we trust this plan if a clear line cannot be drawn to one author? We will also discuss such critical issues as analogies in planning with disparate experiences, civil-military planning by analogy, trust, provenance, and organizational issues in planning.

Keywords: Case-based reasoning, analogy, coherence, trust, planning

1 Introduction

In this paper, we will present a planning approach that utilizes experience from several different sources to create a solution to a problem by utilizing an epistemic approach called *robust coherence*. *Robust coherence* takes Thagard and Verbeurgt's (1998) characterization of coherence as a constraint satisfaction problem one step further by applying counterexamples from the evolving world to ascertain the validity of information. Under *robust coherence*, trust in a solution is established by correspondence with both the outside world and the collective understanding of a problem. This approach addresses two important requirements for the successful execution of complex endeavors: diverse expertise and trust.

Complex endeavors, by definition, require the participation, expertise, and diversity of several different thinkers and planners (Alberts and Hayes, 2007). *Robust coherence* allows the diverse experiences of planners to be combined into a collective set of possible actions and goals. This collective understanding is used to select the most coherent set of actions and goals, allowing for a shared understanding of the group's course of action and values. This allows otherwise separate collections of experience to work together in new, previously unexplored ways.

One of the biggest hurdles in the execution of complex endeavors is a lack of the perception of competence and trust between decision makers (Alberts and Hayes, 2007). Since each planner must contribute their unique expertise and experience to solve problems, decision makers must be able to believe and rely upon those contributions. In this sense, the need to *believe* experiences is what drives the need to *trust*. Therefore, the issue lies not simply in the explicit sense of *trust* itself, but in a sense of *truth* in the information presented to a thinker. *Robust*

coherence addresses the issue of trust from this epistemic perspective. The establishment of *robust coherence* itself then performs double-duty. It not only combines the experiences, but also establishes the overall trustworthiness of the collective understanding. Because counterexamples from the world are taken into account, the natural solution of the coherence system yields trustworthy actions and goals from each collection of experience.

Using this sense of coherence, decision makers can formulate solutions to planning problems. The example system presented in this paper uses case-based agency to accomplish this goal. Case-based agents are software agents that act in a dynamic, complex world by utilizing past experiences to generate plans and react to change. They perform case-based reasoning (CBR), which is the process of recalling past plans rather than starting from scratch (Hammond, et al. 1996). We will refer here to case-based agents as *agents* for short.

In complex endeavors, the world is understood to be dynamic and uncertain (Alberts and Hayes, 2007). This contrasts with well-understood domains, such as mathematics, where more information implies better models and more certainty. That level of certainty would be better suited for a *complicated*, rather than *complex* endeavor[1]. The *everyday* world is also understood in case-based agency as dynamic and complex (Hammond, et al. 1996). This means that success in a complex endeavor attains the same standard as in case-based agency: adaptability, agility, flexibility, and decisiveness.

To place the research in this paper in context, we should understand the goals of DEEP, or Distributed Episodic Exploratory Planning, a project underway at the Air Force Research Laboratory Information Directorate. This program's research has been focused on emerging concepts for the future of Command and Control (C2). Our first goal is developing a C2 environment that supports the vision of Network Centric Operations (NCO) (as defined by Alberts, et al. 1999). The tenets of NCO are:

- Information sharing
- Shared situational awareness
- Knowledge of commander's intent

Our second goal is developing a distributed C2 environment that supports Cyber Warfare. A key challenge of cyber C2 is the speed at which electrons move, requiring a C2 system of unprecedented response time, global arena, and human expertise that my not be located in a single command center. We currently assume that the cyber domain requires a faster than real-time (or predictive) C2 capability that is not bounded by traditional thinking (i.e., air and space).

The implication of these goals is an all encompassing, global battlespace that requires expertise that is seldom co-located. This transformation requires a vastly new C2 process that can adapt to

[1] Complicated endeavors involve a large number of interrelated elements that often affect one another. The cause-effect relationships between these elements are well known and understood. This contrasts complex endeavors, which involve elements with meaningfully different perspectives and goals. There is a lack of well-understood relationships, and consequences are often unpredictable (Alberts and Hayes, 2007).

the any level of conflict, provides a full-spectrum joint warfighting capability, and can rapidly handle any level of complexity and uncertainty.

The long-term goal of the DEEP project is to develop a prototype system for distributed, mixed-initiative (human and machine cooperative) planning that improves decision-making by applying analogical reasoning over an experience base. The two key objectives of DEEP are:

- Provide a mixed-initiative planning environment where human expertise is captured and developed, then adapted and provided by a machine to augment human intuition and creativity.
- Support distributed planners in multiple cooperating command centers to conduct distributed and collaborative decision making.

The architecture of DEEP was explicitly designed to support the tenets of NCO in a true distributed manner. Because DEEP is not based on any current C2 system, we are able to explore concepts such as combining planning and execution to support dynamic replanning, machine-mediated self synchronization of distributed planners, and experiment with the impact of trust in an NCO environment.

The research in this paper addresses key concerns that face the DEEP project, such as planning with multiple experience bases, and establishing trust in a truly distributed environment. The example system explored in this paper, DEEP's Ensemble Case Knowledge (DECK), uses *robust coherence* to address these concerns. In our example system, case-based agents are arranged into a competent group for a given problem, exchange experiences, establish coherence, and form a solution.

The following sections of the paper describe each step in detail. Section 2 describes the core capability of establishing coherence. Section 3 describes the selection of competent agents for perform this task. Section 4 describes the example system's approaches and algorithms for using coherence for planning and prediction. Section 5 discusses the implementation plan for the example system. Section 6 discusses related works, and Section 7 concludes and highlights future work.

2 Establishing Truth

In case-based agency, agents exchange and utilize experiences to solve problems. This is also true for complex endeavors, where self-synchronizing decision makers use shared awareness to operate in an environment of competence, trust, and interdependence (Alberts and Hayes, 2007). The need to exchange these experiences drives the need to trust. For this reason, trust is intimately connected to truth. In other words, if a group of agents must solve problems by exchanging their experiences, then the mechanism that establishes their mutual trust is the same mechanism that allows them to perform this exchange. This exchange of experiences is what manifests and drives trust. In this paper, we will address the issue of trust as an epistemic issue, concerned with the establishment of truth.

How can an agent establish the truth of various experiences that did not come from its case base? We will be approaching this question by forming a basis of truth based on coherence, then applying critical rationalism to that basis to avoid circular reasoning and group-think.

In practical terms, an agent will present a problem, which will be addressed by different agents by allowing each agent to suggest experiences for adaptation. That set of experiences will be challenged by critical rationalism. This same basic process is repeated to predict the outcome of the resulting plan. We will begin with an explanation of *coherence* and conclude this section with a discussion of *critical rationalism* and how to apply it.

2.1 Coherence

Thagard and Verbeurgt (1998) discuss coherence in philosophical, psychological, and computer science terms. Philosophically, coherent knowledge is knowledge that is mutually supportive in an overall context of justification. In simple terms, establishing coherence involves establishing constraints, and satisfying those constraints by sorting elements into either accepted or rejected sets. Solving this constraint satisfaction problem allows a reasoner to determine what information is mutually supportive in an overall system of beliefs. This becomes more complex in the establishment of these constraints, and the many algorithms that are available to solve the resulting constraint satisfaction problem.

This approach requires the establishment of positive constraints and negative constraints. Certain relations can be characterized as coherent (such as explains, associates, or facilitates), while others denote incoherence (such as incompatible, contradictory, or inconsistent). Elements that are related by a coherent relation are positively constrained, while elements related by an incoherent relation are negatively constrained.

Fulfilling the constraints established by these relations consists of sorting elements into the appropriate set, accepted or rejected, based on their constraints. Positively constrained elements are either both accepted or both rejected. The two elements that are positively constrained must be sorted into the same set. Negative constraints are satisfied by accepting one or the other element involved. The two elements that are negatively constrained cannot be sorted into the same set. Logically, this is equivalent to treating positive constraints as an AND operation, and negative constraints as an XOR operation. However, this system attempts to maximize constraint satisfaction, which means that not all constraints need to be met, just as many as possible.

Each constraint has a strength value. That strength denotes how valuable that constraint is when satisfied. Coherence is attained when the set of elements are sorted into accepted and rejected sets in such a way that maximizes the sum of the strengths of the constraints that are satisfied. This means that some constraints might be violated, because the strength of that constraint is not sufficient to choose that constraint over other, stronger ones.

Relation	Constraint	Explanation	Operation
Coherent	Positive	These elements support one another, either both are accepted or both are rejected.	AND
Incoherent	Negative	These elements cannot coexist, if one is accepted then the other is rejected	XOR

Table 1: Types of Coherence Constraints

We can use this notion of coherence to choose a set of experiences that will be acceptable to use in planning or prediction, because they are coherent with the world view established by the group of cooperating agents. In this way, suggested experiences exist in an overall context established by the whole group of agents. Expertise is exchanged collectively, leading to shared understanding of a problem.

While this form of coherence may appear to be useful to establish truth, there is the danger of forming a coherent truth that is circular in nature. That is, forming a set of experiences that mutually support one another as coherent, but do not actually correspond to the outside world and how it really works. Moreover, there could be multiple coherent systems of experiences which lead to alternate world views. Some or none of these equally plausible interpretations could correspond with reality. For this reason, we must be able to establish coherent collections of experience which also correspond with the current reality of a situation.

2.2 Critical Rationalism

"It is an old maxim of mine that when you have excluded the impossible, whatever remains, however improbable, must be the truth" (Sir Arthur Conan Doyle, 1892)

To address these problems, we turn to another form of reasoning to help inform coherence as an epistemology: a more deductive view of truth which seeks to refute theories based on inconsistency with evidence. Reid and Griffin (2003) discuss this method of reasoning as *critical rationalism*.

Under critical rationalism, theories are postulated and stood to the test of falsification. In other words, theories are considered which are potentially falsifiable, and then compared to a set of observations. If the theory holds up to this set, compared to competing theories, it is considered the least untrue (rather than most true). The measure of this aspect of *truth* is known as *verisimilitude* (Reid and Griffin, 2003). Logically, critical rationalism is based on deduction, rather than induction. This means that verisimilitude measures the degree to which a theory is able to stand up to criticism based on what it deduces should be true.

Reid puts this understanding of knowledge forth as a way to allow military planners to communicate evidence. Under this approach, the utility of information is based on its value in refuting theories, and the importance of doing so. Because an exhaustive search of the complicated world for a complete set of counterexamples is impossible, theories are ranked

based on their level of repudiation, not on their level of truth. Communication establishes a reasonable level of verisimilitude for theories, understanding that a complete measure is unattainable.

We can use verisimilitude to better inform a coherent set of experiences by attempting to locate information that refutes some of the aspects of the experiences. By doing this, we establish the degree of false-ness in those experiences for facing a current problem, and avoid the pitfall of blindly applying experience. We will refer to this approach as *robust coherence*.

DECK uses coherence informed by critical rationality to create set of coherent, robust experiences which address a specific problem. The counterexamples for experiences allow us to examine how those experiences' utility is inhibited by facets of the ever-changing world. In this way, critical rationalism can also indicate critical conditions in the world, allowing for the discovery of new goals. In the following section, we will examine the mechanisms DECK employs to accomplish a system of robust coherence. Specifically, we will examine how DECK forms ensembles of agents, establishes a coherent set of experiences to solve a problem, and combines those experiences to form a single decision.

3 Forming an Ensemble

3.1 Selecting Agents

The first step to addressing problems using multiple case bases is choosing which case bases to consider using. There are several perspectives on how this can be done, and how that ensemble should communicate to solve problems. Here we will examine ways of forming and managing an ensemble of case-based reasoners.

Leake and Wilson (1999) discuss one of experience-based reasoning's main shortcomings: the fact that experience might be wrong. He discusses two forms of consistency that could be measured in a case-based reasoning system: *problem-solution regularity* and *problem-distribution regularity*. Problem-solution regularity expresses the regularity with which a case-based reasoner will retrieve solutions from cases that are acceptable for solving current problems. Problem-distribution describes the reliability with which you can expect problems you face now are anything like problems faced in the past.

Problem-solution regularity is specifically defined by attempting to analyze two different functions to see if one approximates another. *Pdist* (problem distance) defines the distance between a problem faced now and a problem from a past case. In other words, it characterizes the behavior of a case-based agent's similarity metric for case recall. *Rdist* (real distance) defines the distance between solutions an agent finds acceptable for its problem and solutions that can be found in a past case. In order for full problem-solution regularity to exist, the cases that *Pdist* chooses must be the same cases chosen by *Rdist*. In other words, the cases the similarity metric chooses are the same cases that would be chosen if you had looked only of the best solution in a case base. Leake and Wilson define this not as a momentary measure, but rather as a way to analyze the trends of a system's retrieval efficacy. This way you can analyze the regularity over time. The most difficult part of utilizing this method is determining *Rdist*, which reflects the goals and expectations of a CBR agent's case retrieval.

Problem-distribution regularity describes the expectation that problems an agent faces over a span of time are addressed by similar cases within the evolving case base. As with the previous perspective on regularity, this requires an analysis over time. As a set of problems are addressed by the case-based agent, the distance from each problem to its closest case is compared to a set threshold. If the agent continues to locate cases within this threshold, then it maintains problem-distribution regularity. The distance is again defined by *Pdist*, which reflects a CBR retrieval algorithm. As before, this measure reflects a trend of performance for a case base over time.

This is a useful perspective to have because a user's needs may change over time more quickly than a given case base can encompass the proper experiences to react. Knowing this can allow an agent to attempt to locate a new source of cases (working in conjunction with problem-solution regularity analysis) or modify the existing source of cases (prompting case base maintenance).

These measures of regularity can be used to form an ensemble of case-based agents to work together on a problem. In other words, case bases can be chosen to take part in planning based on the desired levels of regularity they have based on the problem and the expectations of the best solution.

Having a high problem-distribution regularity indicates that this case base faces similar problems to the one faced by the originating agent who faces a present problem. This means that the case base can be viewed as a *specialist* in this problem domain.

In contrast, having high problem-solution regularity indicates that this case base produces acceptable solutions, regardless of which kinds of problems they are for. In this way, the case base can be seen as an effective *generalist*, in that it usually produces acceptable solutions, but for a possibly wide range of domains.

Depending on the current problem and the discretion of the originating agent which faces that problem, an ensemble with the right mix of specialists and generalists is useful in addressing a problem from many different perspectives.

DECK uses these measures of regularity to determine the right mix of agents to collaborate with given what CBR task it is presently performing. If the task at hand is to form a plan, then the originating agent's *Rdist* will reflect what a good plan looks like. If the task at hand is to perform a prediction, then the originating agent's *Rdist* will reflect what is expected from a well-founded prediction. These two tasks are explained in further detail in Section 4.

3.2 Agent Communications and Collaboration

Once a group of agents is established as the group that will be solving a task, how these agents interact is an important issue. This is especially relevant in complex endeavors. Communication, intent, and understanding are important to agility in C2 and execution. Here we will discuss ways for agents to communicate and coordinate their efforts for a task.

Multiple agents can use models of competence to make a decision about which agent to collaborate with for assistance in a task (Martin, et al. 1999). An agent can be described as

competent when it is able to make these decisions, informed by its own competence model and models of other agents' competence.

Describing the exact competence model depends on the task. It can be closely related to the problem-solution regularity, problem-distribution regularity, or other forms of regularity for a case base, as described before. However, because the actual values of functions like Rdist from the originating agent are not known by other agents, they would use local or collective approximations of those functions to determine the value of models that require them.

We can use this environment to describe different cooperation modes. These modes indicate how two agents cooperate with each other to accomplish CBR tasks (Martin et al., 1999). Between two agents, cooperation modes indicate the following:

- Which agent originates the task?
- Which agent describes the problem solving method?
- Which agent applies their computational resources?
- Which agent contributes the experience base?

An agent can choose to retain or relinquish different levels of authority in the performance of a given CBR task. For example, if an agent knows exactly which features of a given case are important, it can prescribe the similarity metric to another agent to perform a case recall task. Conversely, if an agent is unable to realize how best to perform a case recall task, it can leave it up to a helping agent to use its local similarity metric. This decision could be applied to many tasks that require a domain-specific strategy to perform. Additionally, we can determine additional cooperation modes that take advantage of the delegation of the four aforementioned attributes across two or possibly more agents.

Agents in DECK make decisions about how they interact based on their competence models. If the agent that originally posted the problem does not have the competence necessary to accomplish a task, it will allot maximal authority to competent helper agents in the ensemble. If the agent is more competent than other agents, it can assist the helpers, taking advantage of their unique experiences and/or resources. These decisions are made on the fly to allow for emergent leadership based on each agent's merit in different tasks and domains.

In practical terms, this establishes how DECK collects the initial set of experiences from each agent in an ensemble. If the agent that originally posted the problem is hard-pressed to solve the problem on its own, DECK allots more authority to assisting agents in the ensemble. For different steps of the CBR cycle this means that different methods of recalling experiences from a case base are communicated. The most competent agents for a given problem help determine the initial set of experiences that need to be made coherent. Once that initial set of experiences is identified, the agents need to work together to refine that set into a coherent collection. This coherent collection denotes the collective view of a problem, and will be the set of experiences put to the scrutiny of emerging counterexamples.

4 Planning and Prediction in DECK

This section describes two tasks that DECK performs in the context of DEEP, planning and prediction. Planning is accomplished by the recall of previous experiences from several sources. These previous experiences are placed into a system of coherence, which is solved to determine the best set of actions and goals for a problem. Prediction is accomplished by the recall of outcomes based on actions and goals. These possible outcomes are also placed into a system of coherence and solved for the most likely outcome to occur based on a collective understanding. We will first discuss the planning task, and then proceed to the prediction task using the implementation of DECK as a running example.

4.1 Deliberative Coherence

Thagard and Millgram (1995) describe an application of the theory of coherence (discussed in Section 2.1) to the world of deliberative planning. This approach, called *deliberative coherence*, refines a set of actions and goals by selecting actions and goals which are most coherent. It seeks to establish positive and negative constraints based on *facilitate* and *incompatible* relationships between factors (actions and goals). These relationships define positive and negative constraints which are resolved to select both actions and goals to undertake as a plan. DECK utilizes deliberative coherence to form new plans based on previous experiences recalled by an ensemble of competent agents. Rather than encoding actions and goals from scratch, DECK populates the system of coherence with actions and goals from previous experiences, and then determines the most coherent set of actions and goals.

Thagard outlines a set of principles that describe the interaction of different factors in a system of coherence. This denotes how actions and goals interoperate, and will provide the starting point to apply deliberative coherence to planning using case-based reasoning. The six principles are outlined directly below (Thagard, 1995):

P1. *Symmetry*. Coherence and incoherence are symmetrical relations: If a factor (action or goal) F1 coheres with a factor F2, then F2 coheres with F1.

P2. *Facilitation*. Consider actions A_1 ... A_n that together facilitate the accomplishment of goal G. Then

(a) each A_i coheres with G,
(b) each A_i coheres with each other A_j, and
(c) the greater the number of actions required, the less the coherence among actions and goals.

P3. *Incompatibility*.

(a) If two factors cannot both be performed or achieved, then they are strongly incoherent.
(b) If two factors are difficult to perform or achieve together, then they are weakly incoherent.

P4. *Goal priority*. Some goals are desirable for intrinsic or other non-coherence reasons.

P5. *Judgment*. Facilitation and competition relations can depend on coherence with judgments about the acceptability of factual beliefs.

P6. *Decision*. Decisions are made on the basis of an assessment of the overall coherence of a set of actions and goals.

4.2 Planning in DECK

New plans can be formed from previous experience by using actions and goals from recalled experiences as the initial factors in a system of deliberative coherence. This system is populated with facilitation and incompatibility relations that allow the system of coherence to be solved. This means that deliberative coherence does not have to be utilized from scratch, but rather can be accomplished by an ensemble of competent case-based agents.

This is accomplished by establishing relationships between factors inside and outside of each experience. As different agents suggest experiences from their own case bases, DECK interprets the actions and goals from these experiences in a system of deliberative coherence. This reasoning establishes the positive and negative constraints between these portions of experience, allowing a collective set of actions and goals emerge as coherent. This collective set can then be adapted and de-conflicted as a cohesive plan.

In order to accomplish this, DECK makes certain assumptions about the properties of cases and how they can be interpreted as a system of deliberative coherence. These assumptions are the guiding principles that allow DECK to establish ways to interpret suggested cases as member of a system of deliberative coherence. These assumptions are stated below:

A1. *Completeness*. The factors (actions and goals) stored in a case and the meaningful relationships between them are sufficient to describe the experience the case represents.

A2. *Closure*. If a factor is stored in a case, then it is related to at least one other factor in that case in a meaningful way.

A3. *Inference*. Meaningful relationships can be inferred between factors by using other information within a case, or knowledge of a case's structure.

A4. *Simplicity*. Inferences which discover relationships between factors will be applied prudently to avoid undue complexity and contradiction with explicit information.

A5. *Proximity*. Inferences involving "similar" factors will only be applied for factors which are similar to a degree, *s*, or higher.

Given these assumptions, DECK utilizes three key mechanisms to infer relationships between factors stored in a case. These mechanisms allow DECK to accept cases that do not explicitly store facilitation and incompatibility information. The mechanisms are implemented by the DECK prototype, and allow DECK to place cases at different levels of detail into a system of deliberative coherence. These mechanisms are:

C1. *Structure*. Actions stored in a case facilitate the met goals of that case, and are incompatible with the failed goals of that case.

(a) The degree of strength for these relationships, *d*, is dependent upon the number of actions in the case, *a*, such that: $d=1/a$.

C2. *Effect Transitivity*. Information about effects and how actions achieve or avoid them can indicate indirect relationships to other factors.

(a) If an action achieves an effect, and a factor requires that effect, then the action facilitates that factor.
(b) If an action avoids an effect, and a factor requires that effect, then the action is incompatible with that factor.
(c) The degree of strength for these relationships, *d*, is dependent upon the degree to which the effect is achieved or avoided, *e*, and the number of factors that require that effect, *f*, such that: $d=e/f$.

C3. *Competition*. Actions from different cases that facilitate a similar met goal are incompatible with each other. They do not need to be performed together.

(a) The degree of strength for these relationships, *d*, is dependent upon the similarity of the met goals, *s*, and the number of actions that facilitate those met goals, *a*, such that: $d=s/a$.

Using the principles of deliberative coherence, the assumptions made about case-based reasoning, and the mechanisms described to populate coherence systems with experience, we can now describe an algorithm to allow DECK to plan using deliberative coherence. This algorithm takes as input a set of cases which have been recalled from different competent agents. The algorithm applies the various principles, assumptions, and mechanisms at its disposal to return a coherent set of actions and goals which describe a plan. The algorithm is outlined below:

1. For each suggested case:
    a. Add all factors (actions and goals) to the system of coherence.
    b. Add all facilitation and incompatibility relationships that are stored in the case to the system of coherence.
    c. If effect and requirement information exists, apply C2 (*Effect Transitivity*) to the factors in the system of coherence.
    d. If A2 (*Closure*) is not satisfied:
        i. Apply C1 (*Structure*) to infer additional relationships.

2. Determine similar goals by applying A5 (*Proximity*), and apply C3 (*Competition*) to the actions which facilitate them.

3. Remove redundant and/or contradictory relationships, as required by A4 (*Simplicity*):
    a. Apply P4 (*Goal Priority*) and P5 (*Judgment*) where possible.
    b. If two factors are related in *opposite* ways simultaneously, the stronger relationship remains with its strength reduced by the strength of the weaker relationship.

c. If two factors are related in the *same* way simultaneously, only one description of that relationship is necessary, with a combined strength of the similar relationships.

4. Convert facilitation (coherent) and incompatibility (incoherent) relations to positive and negative constraints as per P1 (*Symmetry*), P2 (*Facilitation*), and P3 (*Incompatibility*).

5. Weight constraints based on emerging counterexample information. Constraints with more counterexamples are less likely to be satisfied.

6. Solve the constraint satisfaction problem, and form a plan as per P6 (*Decision*).

In DECK, agents share recalled experiences to address a set of new goals. These experiences may be relevant to the set of goals from each agent's individual perspective, but in reality may not be fully relevant from a collective perspective. Using the above algorithm, DECK establishes the interrelationships between these experiences. Using this populated system of deliberative coherence, deliberative coherence is used to refine a set of actions and goals that represent a collectively established plan.

Throughout this entire process, these experiences are treated like hypothesis by a system of critical rationalism. That is to say, as counterexamples that contradict elements of the experiences enter the system, they weaken the strength of the relations posited by the system of deliberative coherence. Because the strength of the relationships indicate the strength of the constraints in the constraint satisfaction problem, relations with more counterexamples are less likely to be upheld in the constraint satisfaction problem. In this way the solution to the constraint satisfaction problem achieves a system of robust coherence that follows the view of truth established by collective experience and also adheres to the evolving reality of the situation.

4.3 Explanatory Coherence

In order to predict the outcome of a proposed set of actions, DECK uses case-based reasoning supplemented by another form of coherence called explanatory coherence (Thagard, 2005). Rather than simply querying case bases to find similar actions from the past and directly applying the solution, DECK recalls experiences from the past and uses them as possible hypothesis about how that outcome came about. This utilizes past experiences in their explanatory capacity. This is to say that a past case denotes an instance of past events. Those events may be remembered without any explanation as to why they occurred. However, this does not mean that given a large number of experiences from different sources, we cannot begin to deduce how the world around us works. In this paper, we do not take the approach of induction to determine these explanations, but rather apply coherence in a similar way that was just described in use with planning. We will begin by outlining the established principles of explanatory coherence, and supplement those principles to ones that apply directly to case-based reasoning. Below are the seven principles of explanatory coherence (Thagard, 2005):

E1. *Symmetry*. Explanatory coherence is a symmetrical relation. If proposition P coheres with Q, then Q coheres with P.

E2. *Explanation*.

(a) A hypothesis coheres with what it explains, which can be evidence or another hypothesis.
(b) Hypotheses that together explain some other proposition cohere with each other.
(c) The more hypotheses it takes to explain something, the lower the degree of coherence.

E3. *Analogy*. Similar hypothesis that explain similar pieces of evidence cohere.

E4. *Data Priority*. Propositions that describe the results of observation have a degree of acceptability on their own.

E5. *Contradiction*. Contradictory propositions are incoherent with each other.

E6. *Competition*. If P and Q both explain a proposition, and if P and Q are not explanatorily connected, the P and Q are incoherent with each other. P and Q are explanatorily connected if one explains the other or if together they explain something.

E7. *Acceptance*. The acceptability of a proposition in a system of propositions depends on its coherence with them.

4.4 Prediction in DECK

In this system, propositions are related to each other by either *explain* or *contradict* relations. Explanation is a coherent relation, and contradiction is an incoherent relation. The purpose of explanatory coherence, as stated before, is to determine the most coherent set of propositions to accept, based on the maximal solution to a constraint satisfaction problem. This system in DECK will be constantly updated with information from the world that denotes counterexamples to some of the relations. Because the principles for explanatory coherence include *data priority*, we can exercise this principle to insure that our adherence to critical rationalism is maintained in our establishment of explanations to use for prediction.

Once the explanations are established, they will be applied to deduce the most coherent outcome that can be expected. These explanations will be established in experiences that do not store in a way analogous to the method described for planning. Rather than establishing *facilitation* and *incompatibility* relations, the principles are adjusted to establish *explanation* and *contradiction* relations.

4.5 DECK Summary

In order for CBR to be relevant in the world of complex endeavors, a CBR planning system needs to address a fundamental concern: sometimes we do not know how to tell people what we really want. That is, just because an explicit goal is met, it does not follow that an approach is necessarily acceptable. In traditional case-based reasoning, a previous experience is recalled and applied, but what if this experience contains implicit goals, constraints, assumptions, or even deception that cannot be challenged or utilized by just recall alone? More importantly, how can another agent be assured that the recalled experience is even relevant?

Here we have presented DECK, which addresses these concerns by forming ensembles of case-based agents which establish a set of coherent experiences to solve a problem. These coherent experiences are subject to counterexamples from the environment, which further define what experiences are truly applicable to the situation. These experiences are combined to perform planning and prediction tasks by applying more specific principles of coherence.

5 Implementation Plan

Implementation of the DECK prototype is underway, and is scheduled to be completed by June 2008. Initial efforts will center on using experience to populate a system of deliberative coherence. This will be followed by work on agent selection, competence modeling, and the use of counterexamples in coherence. After that, effort will be placed behind the case-based prediction capabilities of DECK and DECK's integration into the overall DEEP architecture. The three stage approach allows for potentially eye-opening ancillary research, side projects, and other opportunities.

The initial developmental objective will be to form a system of deliberative coherence using actions and goals from previous experience. Results from this initial objective will be presented at the symposium.

6 Related Works

There are several papers in the case-based reasoning literature which are related to the endeavor DECK is undertaking. Utilizing multiple case-based reasoners is a powerful approach to solving problems, and there are a myriad of different ways decisions can be made within the context of a group. Here we will discuss different approaches to forming ensembles, combining cases, and performing planning and prediction using case-based reasoning.

6.1 Forming Ensembles

Leake and Soojiamurthi (2001) highlight the need for multiple, mutually supportive case bases. They describe two basic mechanisms for utilizing multiple case bases: case dispatching and cross-case-base adaptation. Case dispatching requests another, more capable case base handle a CBR task. Cross-case-base adaptation uses weights to establish a bias through which adaptation can occur across case bases.

While this approach is interesting and would give DECK an otherwise absent adaptation techniques, DECK instead uses a more flexible communication paradigm. However, it does show that it is possible to take a flexible approach in the adaptation task; delegating responsibility to agents that are the most capable for a certain facet of adaptation.

Plaza and Ontanon (2001) present an alternative approach, using multi-agent collaboration policies for CBR. Within a distributed CBR environment, characterized by each agent using their own similarity metrics for case retrieval, they discussed different modes of collaboration.

Each collaboration mode involves asking other agents for a Solution Endorsement Record and taking a poll for the solution class that receives the most votes. A Solution Endorsement Record pairs solutions suggested by an agent with the cases from which each solution was derived. The voting mechanism works by allowing each agent one vote. That vote can be distributed fractionally across several solutions based on the number of supporting cases. This way, agents with larger case bases do not have an advantage. Each agent balances its recommendation over the different solutions it derives before reporting that distribution back as its vote.

While these polices are designed to maintain the autonomy of each agent, it is easy to imagine several points of departure and innovation from the base polices discussed. The biggest trade-spaces involved are the amount of agent communication, the number of distinct solutions, and how well each agent can gauge its own competence for a task. DECK could build on its initial recall task by using different collaboration modes to retrieve a more coherent set of cases.

6.2 Case Combination

An interesting approach to case retrieval in a distributed CBR system utilized multi-agent negotiation to decide upon a composite case to reuse (Prasad, et al. 1996). Under this approach, agents would operate asynchronously by performing a variety of negotiation tasks. The interesting thing about this approach is how the retrieval process is framed. Rather than a search operation simply split up for many agents to accomplish, retrieval becomes a constraint satisfaction problem.

Each agent seeks local matching cases based on the feature set and constraint set that they know about. When attempting to combine these local solutions during a merge or extend operation, the constraints of other agents might be violated. Agents then negotiate by offering their constraints as feedback and incorporating those constraints as the acceptance of this feedback.

This dialogue becomes even more interesting as conditions lead to the relaxing of certain constraints (hard or soft), the adjustment of adaptation or search methods, or the refinement of parameters that describe a problem. This interplay affects the reliability, uncertainty, quality, and cost of a solution. This leads to unique and interesting composite cases, which no individual agent could have created on their own (with a limited perspective on the problem).

DECK could leverage part of this approach in the negotiation of constraints, and how they could become hard and soft. In an implicit sense, this already occurs because the system of coherence is attempting to maximize the total weight of constraints satisfied, which means some are "softened" de facto. However, allowing agents to negotiate constraints during the initial case retrieval might minimize the number of constraints that are overlooked by coherence, making the satisfaction problem easier. Also, the intelligent combination of partial cases may give DECK an additional option in case adaptation.

Purvis and Pu (1995) present a technique to case combination which represents cases as constraint satisfaction problems (CSP). Under this approach, the retrieval phase of CBR is accomplished via the structure mapping approach, recalling cases from the past that addressed similar constraints. These recalled cases not only contribute the structure of the CSP itself, but also some initial solution values. The adaptation phase, then, is really a constraint satisfaction

problem. Purvis and Pu utilize the *minimum conflicts heuristic* (Minton et al., 1992) to address the consistency of one solution, as well as the merging of sub-problems. Because the constraint satisfaction problem heuristic employed requires no specific domain knowledge, this approach's adaptation step is free from that knowledge requirement, making it capable of addressing problems from any (possibly multiple) domains.

DECK can learn important lessons from these approaches, but it is important to note that because DECK utilizes a case representation that is not explicitly based on constraints, we cannot apply these approaches directly. The greatest area for improvement from these approaches lies in refining the initial case recall performed to populate the system of coherence.

6.3 Planning and Prediction

There are interesting ways we can further supplement DECK's ability to perform planning and prediction tasks for DEEP. Here we will examine pragmatic centrality and abductive reasoning to define important elements in an experience to use in planning and prediction.

An important aspect of approximating *Rdist*, the distance of a suggested solution from an agent's ideal solution (Section 3.1), is considering what information in an experience is pragmatically distinguishable from the rest of the experience. Here, we can use an algorithm to establish the *pragmatic centrality* of an experience (Thagard, 1990).

Under this approach, we can determine if certain elements are important based on what relationships or constraints those elements exist within. Thagard discusses three such structures: problems, explanations, and arguments.

When encountering a problem, we can consider the goals of that problem to be important. In terms of pragmatic centrality, we can highlight as important elements that form a chain from the starting conditions of the problem to its goal state. This chain is followed backwards until a terminal *root* item is found. This item is marked as important. Pragmatic centrality could be a great asset to DECK in how it can inform the weight of certain constraints in the system of coherence, and how it could improve the initial case recall task.

When attempting to explain phenomena to form a prediction, we can highlight important elements of explanations by following similar chains as we did for problems. That is, we follow the chain of reasoning back from the goal of the explanations through the known relationships to the starting state. This highlights the important facts and elements in explanations. A similar approach can be also used for arguments, highlighting as important elements related by "if" from the conclusion of the argument to the premise of that argument.

Abductive reasoning can be used to establish the explanation of phenomena, but can be explosively complicated in terms of the chaining of explanations. Case-based reasoning can be used to apply abduction by recalling past explanations rather than postulating explanations from scratch (Leake, 1993). Leake highlights some important issues in applying a case base of explanations. Most importantly, that everyday problems deal with domains with weak theories where incomplete explanations are acceptable and sometimes necessary. This follows closely with our use of CBR in the *everyday world* for complex endeavors.

Case-based recall of explanations can take these issues into account. In this environment, we can focus the recall and adaptation of explanations based on the anomalies which prompted explanation, and the overall goals under which the explanation was required.

By indexing explanations based on anomalies, we can recall abductive hypotheses that better suit the nuances of the situation, rather than any and all explanations that might be plausible. The knowledge of anomalies - what is surprising about phenomena - indirectly reflects the immediate goals in forming an explanation. The information that surprises an agent about a situation indicates, intuitively, what actually needs to be explained. If some facet of an event does not appear anomalous, then explanation is not necessary for that facet.

Moreover, by using the overarching goals of an agent trying to explain phenomena, we can even better focus the explanation to the simplest theory that not only explains the phenomena but also does so in a way that suits the agent's goals. Unless explanations are focused on an agent's goals, then all explanations would be equally plausible, and equally useful to that agent.

DECK can use abduction to provide a more advanced way to recall explanations for outcome prediction. Rather than attempting to derive the explanation relationships from the structure of the case, we can use anomalies (and pragmatics) to define the case retention process. This is a possible direction not explored directly in this paper.

7 Conclusion and Future Works

In this paper we introduce DECK, an approach to case-based agency that utilizes the establishment of truth as an approach to combining disparate experiences. DECK uses measures of regularity to form an ensemble of agents. These agents use models of competence to define the parameters for the initial case recall task. Using these experiences, DECK establishes a system of coherence based on the task at hand. That coherence is subject to counterexamples from the world which define the robustness and relevance of the experiences. Once the system of coherence is solved as a constraint satisfaction problem, a collection of experience emerges which is based on a collective sense of truth that has stood the test of the evolving world.

Future directions for this research include extending the system of coherence to understand a variety of relationships. We could also focus on establishing coherence when relationships could be either coherent or incoherent based on the context. For agent ensembles, one future area is to use dynamic, pickup ensembles of reasoners based on unfolding events. In order to perform this kind of reasoning in other areas of DEEP, we could research applying explanatory coherence and critical rationalism to other areas of the DEEP architecture. As far as the reasoning itself, we could try to use derivational replay of coherence systems to see if relationships maintain their level of coherence and incoherence over time. Another important area to explore is challenging some of the domain and agents assumptions with more sophisticated supplemental reasoning. All of these areas focus on making DECK more versatile, robust, and flexible to an evolving world and evolving needs in planning and prediction. DECK represents a powerful opportunity to give decision makers access to other thinkers' memories, a unique advantage in establishing trust and maintaining an edge in strategy and tactics in complex endeavors.

Citations

Alberts, David S, John J. Garstka, & Frederick P. Stein (1999), Network-Centric Warfare: Developing and Leveraging Information Superiority, *DoD Command and Control Research Program Publication Series*. Washington, DC.

Alberts, Davis S, & Richard E Hayes (2007). Planning: Complex Endeavors. *DoD Command and Control Research Program Publication Series*. Washington, DC.

Doyle, Arthur Conan (1892). The Adventure of the Beryl Coronet . *The Strand Magazine : an illustrated monthly, Vol. 3 (Jan.-June 1892)*, 511-525.

Hammond, Kristian J, Mark J Fasciano, Daniel D Fu, & Timothy Converse (1996). Actualized Intelligence: Case-based Agency in Practice. *Applied Cognitive Psychology*, 10. S73-S83.

Leake, David B (1993). Focusing Construction and Selection of Abductive Hypotheses. *Proceedings of IJCAI '93*, 24-29.

Leake, David B and David Wilson (1999). When Experience is Wrong: Examining CBR for Changing Tasks and Environments. *Proceedings of ICCBR-99*.

Leake, David B and Raja Sooriamurthi (2001). When Two Case Bases are Better than One: Exploiting Multiple Case Bases. *Proceedings of ICCBR 2001*.

Martin, Fransisco J, Enric Plaza, & Josep Lluis Arcos (1998). Knowledge and Experience Reuse through Communication among Competent (Peer) Agents. *International Journal of Software Engineering and Knowledge Engineering.*

Minton, Steven, Andy Philips, Mark Johnston, & Philip Laird (1992). Minimizing Conflicts: a Heuristic Repair Method for Constraint Satisfaction and Scheduling Problems, *Artificial Intelligence, 58*, 161-205.

Plaza, Enric and Santiago Ontanon (2001) Ensemble Case-Based Reasoning: Collaboration Policies for Multiagent Cooperative CBR. *Proceedings of 4th International Conference on Case-Based Reasoning*.

Prasad, MV Nagendra, Victor R Lesser, & Susan E Lander, (1996). Retrieval and Reasoning in Distributed Case Bases. *Journal of Visual Communication and Image Representation, Special Issue on Digital Libraries, Vol 7, No. 1*. 74 - 87

Purvis, Lisa and Pearl Pu (1995) Adaptation Using Constraint Satisfaction Techniques, In Veloso, M. & Aamodt, A. (eds.) *Topics in Case Based Reasoning, Proceedings of the First International Conference on Case Based Reasoning, LNAI Series, Springer Verlag*.

Reid, Darryn J and Ralph E Giffin (2003). A Woven Web of Guesses, Canto Three: Network Centric Warfare and the Virtuous Revolution. *Proceedings of 8th ICCRTS*.

Thagard, P, K J Holyoak, G Nelson, & D Gochfeld (1990). Analog Retrieval by Constraint Satisfaction. *Artificial Intelligence, 46*, 259-310.

Thagard, P. and Millgram, E. (1995) Inference to the best plan: A coherence theory of decision. In A. Ram & D. B. Leake (Eds.), *Goal-driven learning*. 439-454.

Thagard, P and K Verbeurgt (1998). Coherence as Constraint Satisfaction. *Cognitive Science, 22(1)*. 1-24.

Thagard, P (2005). Testimony, credibility, and explanatory coherence. *Erkenntnis*, 63, 295-316.