12TH ICCRTS

“Adapting C2 to the 21st Century”

Title:

Creating and Capturing Expertise in Mixed-Initiative Planning

Track 4: Cognitive and Social Issues
Track 1: C2 Concepts, Theory, and Policy
Track 8: C2 Technologies and Systems

Authors
Anthony J. Ford
Joseph A. Carozzoni

Point of Contact:
Anthony J. Ford
Air Force Research Laboratory Information Directorate
Systems & Information Interoperability Branch (IFSE)
525 Brooks Road, Rome, NY 13441-4505
Anthony.Ford @rl.af.mil
315-330-2610

Abstract

Compared with prior engagements, commanders today are exposed to a battlespace that is more dynamic and less predictable. With increasing frequency, commanders are confronted with an array of problems whose solution requires knowledge beyond their military training. In these novel situations, decision makers often rely on their past experiences incorporating a process best described in research as analogy-based reasoning and/or recognition primed decision making. While the relevancy of the experience is based on the individual, a key goal would be to capture and exchange relevant experiences between individual decision makers. The shocking events of 11 September 2001 may have been less shocking to anyone with experience serving in the Pacific theater of operations toward the end of World War II and experienced Kamikaze warfare. This paper describes work in progress at the USAF Research Laboratory Information Directorate to capture, develop, and provide an experience-based reasoning system to commanders during mixed-initiative planning. The objective of this work is to provide a rich database of experiences for the commander to compare to the current situation. The research described in this paper is aimed at developing a computational representation for episodic models, and reasoning on those models for retrieval and experience extraction.

Introduction

By using the experience of the past, we can seek insight on the present, and shape the future. Experience-based reasoning draws upon the past in order to adapt past successes to current problems. Applying this sort of reasoning to war planning is an exciting opportunity to tap into a powerful human process and augment that process with the power of technology. By using a computer to assist in the storage, retrieval, and application of experience, we can enhance decision making with another take on an already occurring human process. The challenge becomes knowing how to store experiences in a computationally reasonable way, and how to match those experiences to new situations as they arise. We will explore and expand upon different technologies and ideas to engage this task. First, we will explore how to store an experience itself. Then, we will seek to adapt a plan representation to an episode representation to store these cases. After that, we will look at different ways to pull experiences back out of the case base to use again. Finally, we will look ahead to future possibilities for this research.

Case Based Reasoning

The primary means of experience-based reasoning we will be exploring is case-based reasoning. In case-based reasoning (CBR), experiences are split up into discrete segments of time known as cases. These cases contain the context of a problem, the solution employed, and the results of that solution. Given a set of cases, known as a case-base, an AI can apply the four R's of case-based reasoning. When a new problem arises, the reasoner can *retrieve* the most relevant cases. Then it can *reuse* those cases to come up with a proposed solution. Then it can *revise* that solution to better apply to the current situation. Then, after a solution is tried, it can *retain* the solution and results back in the case-base. By using these four simple steps, a case-based reasoner learns over time by experience, rather than formal rules. The asset of this form of reasoning is that even if the solution does not succeed, that experience can be useful in the future as an example of what *not* to do. Additionally, storing the context in which the solution was attempted can provide further insight later, rather than trying to strip those subtleties away from a solution as one would do while generalizing in a more rule-based form of AI.

Modern Planning

In a fast paced conflict like the GWOT, finding solutions to novel problems occurs every day. Even though it's hard to imagine looking into the past to find solutions for a conflict that appears so new, true insights can be gained by having a deep knowledge of the past. Additionally, as new solutions are found to these novel problems, using case-based reasoning would store those experiences for future use. In other words, if you can best a tactic once, and you can always remember what you had to do, best practices can arise that come from the concrete ground of experience. Rather than waiting months or years for a formalized counter-tactic to be devised, using case-based reasoning (especially if it can be done in groups) can employ the school of hard knocks to get inside of an adversary's decision loop. In other words, using case-based reasoning as a conflict continues allows decision makers to gain faster insight into potential approaches to a problem, giving them the ability to make progressively faster and more effective decisions.

Capturing an Experience

Sometimes life's greatest lessons are learned through hard experience. Perhaps one of the greatest abilities of the human mind is its sensitivity to the world around it in respect to time. As we move forward, we gain new experiences which forge our beliefs, skills, and characters. This is especially true a high-stress environment like battle. How do we translate the experiences gathered from good (and bad) command decisions into a form that is understandable by a computer to use in mixed initiative planning? For this, we have to turn to the field of knowledge representation, with an eye for this unique domain.

To represent a commander's decision, there are a few things that need to surround it in order to provide context. First and foremost, the actual plan itself must be stored. However, in a vacuum it can provide little insight. We must also know *why* that plan was chosen. First, we must then know what the commander wanted to accomplish; what objectives the commander hoped to achieve. Second, we must know what constraints were placed on the

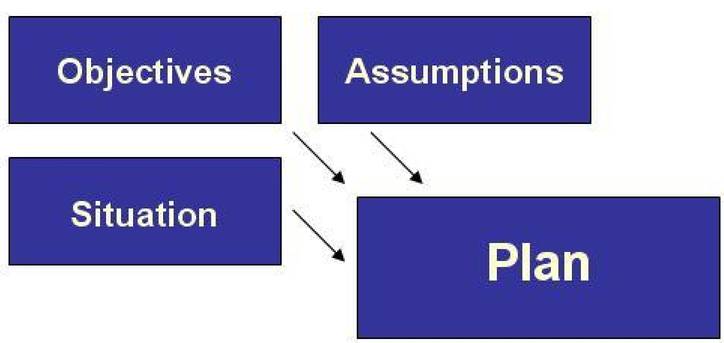

commander based on the world (s)he lives in. The resources you have available, the geospatial (and non-geospatial) layout of the operations environment, and other dimensions which define your overall situation have a huge impact on what specific decisions can be made. Finally, it's also important to denote what was assumed by the commander at the time of planning. Assumptions drive decision making to a huge degree, since assumptions denote the commander's stance towards the adversary's intent, the adversary's possible actions, available intelligence information, and the proposed efficacy of the actions the commander undertakes. Figure 1 shows the overall areas of information that need to be captured to represent a plan and its context.

Once the plan is formed, the actual execution denotes how the experience actually occurs. To represent this, we need to capture a few different kinds of information. First, the actual actions the commander's force undertook in accordance with their plan. Then, we need to know what events took place in reaction to the blue force's actions. In other words, although there was some assumption of actions other factions would take, we need to know what actions they actually took when the rubber hit the road.

Now that we've captured context, plan, and actual events of the experience, we need a way for the case to be characterized in terms of its overall results. We need to be able to characterize the outcome of the experience. Several metrics could be used to explain the overall outcome of a case. These include the objectives that were met and failed, the assumptions that were true and false, the costs incurred by all factions involved, and any plans and actions that cannot be undertaken due to the selection of this plan (opportunity cost). Figure 2 shows these additional elements in an experience.

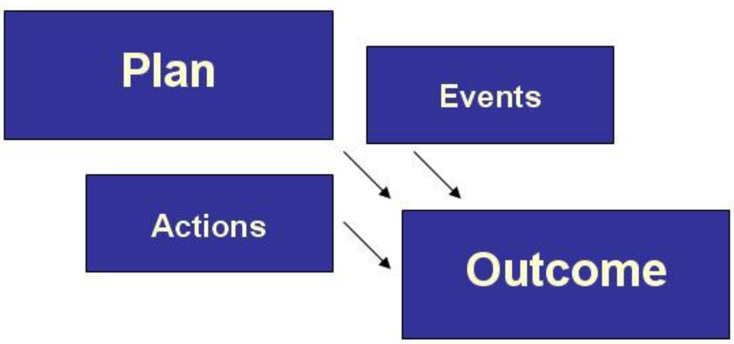

With all of these elements, we now have a reasonable flow of an experience from beginning to end. Temporally, this approach to capturing an experience is agnostic; an experience from this perspective could be five minutes or two years. This presents not only the challenge of intelligent experience capturing, but also the flexibility to formulate the exact case building methodology based on the specific use in mind.

Objectives

Understanding a decision depends heavily on what the decision was trying to accomplish. Knowing why someone does something can give you understanding of the action in relation to intention. Otherwise, we would all appear to be acting for no reason whatsoever. In order to avoid this, we need to be able to represent intent for the sake of reasoning. How do we even begin to break down the mighty 'why' behind human actions? There are a couple of different approaches we could take.

First, we could allow the commander to articulate the ideal world state that they would prefer to see. In other words, we could store the world state that has been adjusted to show what the mission was supposed to accomplish. This would require that the commander knows exactly how they would like the world to look, with a great level of detail and specificity (or rather, with the same level of detail as the world state model). This approach could be thought of as supplying the *vision*, or overall goal end state.

A second approach would be to allow the commander to enunciate objectives in a more direct, almost rule-based way. Statements such as ‘establish air superiority’ denote a characterization of the end state. In other words, there is not just one specific instance of the world where air superiority is established. There are many ways to approach that state and many methods by which to achieve it. Any end state, regardless of the details, that satisfies the objective is acceptable. This approach leaves more flexibility for the decision maker, since several various world states can satisfy the objectives. With this flexibility, the commander can make decisions that fit into their own situation.

The computational challenge for this approach is determining how to characterize the overall state in such a way that we can determine if an objective has been attained. In other words, how do we know that ‘establish air superiority’ has been successful? How do we understand air superiority in terms that can be measured and understood by a computer? In order to answer this question, we need to be able to take a natural-language concept like ‘establish air superiority’ and break it down into a concise meaning.

Representing Information

Thus, the challenge becomes: what information in a military planning situation can be captured and stored as a case? To examine this question, it will take a perspective on what a military plan consists of, and what sort of information it depends upon. As we have examined, there are several broad areas which make up an episode (Objectives, Assumptions, Situation, Plan, Actions, Events, Outcome). What sort of data can we collect in each of these areas, and how can that data be stored in a meaningful way that allows us to draw analogies for episodic reasoning?

Regardless of how we decide to represent Objectives, the information required will eventually boil down to storing a way to know when the Objective has been met. In other words, while the exact information stores depends on the approach to representing intent, what the information should tell you (based on the nature of Objectives) is when a goal is accomplished.

In order to properly codify our assumptions, three basic things must be recorded. First, what the assumption is about. Next, what the assumption is. Finally, the level of confidence in the assumption. The first piece of information is fairly straightforward. The second piece of information affords us some flexibility. We could store information in an attribute-value tuple. We could store information as plain text, relying on natural language processing to interpret meaning. We could even store the entire entity itself; as to indicate that its very existence and everything about it is an assumption. All of these approaches afford us the ability to draw analogies.

The final item in codifying assumptions is possibly more challenging. Laying down a number to explain your level of confidence is a difficult feat without resorting to a totally subjective (if not trivial) assessment. In other words, it can be extremely difficult for people to interpret their 'gut feeling' of trust and turn that into a number that can be used in computational reasoning. There are several approaches that can be used to help ground a confidence score in reality. We could use the reliability of the means by which the evidence for the assumption was drawn. For example, if the assumption relies upon sensor information from a sensor with 98% reliability, then that 0.98 can contribute to the confidence in that data on a reasonable way. For evidence gathered in a non-mechanical way, we could abstract this point of view into an overall framework for trust measurement. We can examine the history of the source as an accurate one, assigning a number to the percentage of correct information items. We can use an evidence combining theory (such as Dempster-Schaefer) to coalesce the overall reputation of the source, by using the trust of other individuals in this source. There are a variety of methods by which we could assign trust to a source.

Several things within a situation, or world state, can be easily quantified. What is present and in what quantity, the temperature outside, the altitude of terrain, etc. Several challenges befall us, however, when we start seeking to model less easily visible and concrete things. Social dynamics, public opinion, cultural factors, etc. can present considerable challenges. Several opportunities exist for analogies, especially since simple similarity metrics can work well for the easily quantifiable aspects of a situation. A more thorough structural analysis could be performed for other types of information, such as social networks.

In an abstract sense, we can imagine a plan to be a collection of objectives and the actions that will be undertaken to satisfy those objectives. However, other information can be stored to give a plan more meat on its bones. Cost, for example, is an important factor when choosing one plan over another. The effects of the plan, beyond those that satisfy the objectives, are also important to track.

The things that we can model about an action include who is performing the action, what resources they are using, where they are located, when the action is carried out, and what effect the action has. Basically, while codifying an action, you can ask all the basic questions: who, what, when, where, why, and how.

In many ways, Events are a lot like Actions. However, the key difference between them is that the events we observe are performed by other people with their own, unknown agendas. Essentially, we can ask all of the same questions we can with an action, except for 'why'. Because we cannot know the exact motives of another person, the reason for their action is always an assumption.

The outcome of the plan can be a fairly straightforward thing to capture, in the short term. The objectives that were satisfied, the resources expended, and the assumptions that panned out: these are all fairly straightforward things. However, some aspects of an outcome are not so easy to capture. For example, what plans or actions cannot be taken due to the result of this plan (opportunity cost)? Beyond that, what are the second and tertiary effects of this plan? How long should I pay attention to possible effects from my plan? These are all deep concerns when trying to determine the efficacy of a plan. What seems like a good outcome now, might actually pan out into a disaster down the road.

Now, let's examine the object oriented structure that we can seek to place this information into to create an episode representation. We examined specifically the Core Plan Representation, developed under DARPA and the ARPI initiative.

Core Plan Representation

In order to facilitate sharing information between different planning systems, an effort was undertaken to develop a common plan representation (Pease 1996). Through several iterations, the effort spawned several innovative ideas in knowledge representation for war planning, as well as various iterations for an object oriented plan representation. The fourth version of the Core Plan Representation (CPR) is now documented on TekKnowlege's website, and includes feedback from comments made throughout the effort's history (Pease 1998).

The CPR work has given us a huge amount of foundational knowledge to work from in developing an episodic representation. Most interestingly, since the effort includes a large amount of documentation, we can open up the design process and pick apart differences between each version. In doing this, we can adapt CPR to our specific needs.

Figures from the 1996 release and the 1998 release are available in Appendix A. These class diagrams will be examined and adapted to provide an episodic representation, using the best contributions from each version.

Adapting CPR

Although the Core Plan Representation was evolved over the years of its study, we found it necessary to understand the overall flow of the representation, rather than just the finished product. This required that we examine several iterations of the representation, ranging in chronology from 1996 to 1998. From the disparate but highly related versions, we were able to pick apart various ideas in knowledge representation and customize how we wanted our episodic representation to operate. Because of the nature of the ARPI work, several elements of the representation were open to further growth and change, giving CPR the flexibility we required to apply it to the task of representing cases in mixed-initiative planning. Remember that the point of this representation is to allow a computer to understand the episode for the sake of reasoning, as well as capturing the experience of a human planner.

Moving forward chronologically, there were several intriguing elements from the initial 1996 report on CPR that are highly interesting and potentially very useful. Namely, the inclusion of Facts and Assumptions, which are subclasses of Annotation, help the planner and computer understand the subtle nature of the source of information. Rather than simply tacking on information to an encapsulated element in planning, it can be understood more concretely that information is either derived from evidence with certainty (a Fact) or something the human planner believes is true (an Assumption). This allows the computer-based reasoning to understand how confident it can be in handling the information.

In the same token, the inclusion of Imprecision and Uncertainty in the 1996 version is very important to further articulating the confidence in information. Although the ARPI group admitted that the exact nature of modeling Imprecision and Uncertainty were wider spanning research topics than they were willing to tackle, the inclusion of the classes themselves leaves the opportunity to expand upon the idea. By including a concrete representation of the doubt of the source of information (Uncertainty) and the doubt in the exactness of the information itself (Imprecision), we can carry these ideas forward by utilizing Fuzzy Logic. In other words, even though it is difficult to capture Imprecision and Uncertainty, there are paths by which a computer can understand them mathematically in terms of its internal reasoning using Fuzzy Logic.

From the 1998 version, we can see the addition of a super-class for all represented plan elements called PlanObject. This is a very useful addition, since the Object Oriented nature of the model lends itself well to these types of generalizations. In other words, now that we have a common parent class to all planning elements, we can understand the structure as a more flexible one. PlanObjects now exist in relationships, specialized for their specific function, but the fields and methods they have in common can now be understood from the perspective of one class. This also aids in implementation, since writing code once for all represented PlanObjects when they all have something in common is simply more efficient.

Another welcome addition from the 1998 version is SpatialSpec. It defines a location where an action takes place. However, the current structure allots only one SpatialSpec per Action. In reality however, thinking from a non-geospatial point of view, every Action can take place in a huge variety of places. Not only can an action take place on the ground, but also in cyberspace, in the realm of a political structure, on the lines of a power grid, or any other huge variety of ‘places’ that are defined by the context of the action. Because of this, SpatialSpec will need to be dealt with specifically in terms of adapting the concept to fit this thinking.

An especially useful addition from the 1998 version of CPR is Role. Using Role, we can understand the exact part that is played by the action. Using this information, we can easily draw analogies for that Action by looking for Actions with similar Roles. In other words, by understanding what function an action played in the overall scheme of an operation (such as ‘Fire Support’ or ‘Transport’), we can make a quick and efficient search in the case-base to find similar actions based on what role they accomplished.

Using this mentality, deeper pruning of those results are possible, because the 'first cut' of the Actions are already known to perform the role that is needed. This will become clearer in terms of the MAC/FAC algorithm, and the establishment of partial orders to use similarity on qualitative values, both explained later.

A final, especially useful contribution from 1998 is the Entity class. This class allows Actors, Resources, and Roles to interrelate to each other. In other words, we can view one item from a variety of perspectives by allowing Entity to 'bridge the gap' between its various facets. While one real-world item may be considered an Actor and Resource, for example, we can model it as an Entity, and treat Actor and Resource as 'perspectives' on that single item. This is an exciting opportunity to model highly complex and subtle items without the limitations of only one of the three perspectives. This also allows us to create more perspectives with the ability to interrelate them in the overall CPR scheme.

Now that we've identified specific items within the Core Plan Representation that are particularly interesting, we can work to expand the representation to encompass an entire episode of an experience, rather than just the plan itself. As discussed before, we need to adapt CPR to represent the Situation, Objectives, Assumptions, Plan, Actions, Events, and Outcome of an experience. Figure N shows the level to which we need to change CPR to allow it to represent whole episodes. Some features need to be adjusted, meaning that they already exist and simply need to be customized. Other features need to be articulated, which means that they exist, but not to a sufficient degree of fidelity or development, thusly requiring further research and refinement to use. Other features still need to be created, meaning that they go outside the scope of the original CPR context, thusly requiring that they be made from scratch.

New classes in CPR

Outcome

To represent more than just the plan involved in a military operation, CPR needed to be adapted to show the whole story: from the planning, to the action, to the results. This would allow us to model a whole experience; one that can be stored as an episode for the purposes of Episodic Reasoning. In order to capture results, the Outcome class was created.

Keeping track of the success or failure of an experience is a tricky thing. There are many important things to consider. Not only are the goals the endeavor was undertaken with important, but also the costs of following that particular plan. It's not only a matter of success, but also of pragmatic impetus. In terms of comparing one plan to another, shear accomplishment of goals leaves no basis to discern. In order to track the many and varied possible costs of a plan, a second new class (contained within Outcome) can be created called simply *Cost*.

However, tracking the objectives achieved and cost incurred is not the entire outcome either. Forcing the adversary to spend their resources is also an important factor in the success of a military operation. This is especially so in terms of Fourth Generation Warfare, where affecting the will of the adversary can be strongly influenced by the effect of their own costs. Knowing the red line past which an enemy will not spend any more to oppose you is vital to developing plans and strategies over time. For these reasons, not only the commanders' own costs, but also the cost of their adversaries should be tracked.

In the same grain of responding to important measures of outcome in military planning, another key area of this type of planning is the effective use of good intelligence to produce valid assumptions about the situation. For example, if a patrol tells you that they spotted ten enemy troops at the top of the next hill, your planning should reflect this observation. In fact, this observation could be an indicator of something else not yet seen. For example, seeing ten troops might indicate that there is an enemy base nearby, or that there are one hundred more troops waiting to march once they know the way is clear, or that there is a convoy out of gas ahead with ten hitchhikers looking for a gas station, etc. One piece of information could spawn thousands of assumptions. It is up to a good commander to make the best assumptions possible given the situation. At the very least, in this example, there are ten enemy troops on or near the next hill. This piece of information already can influence planning, without a terribly wide leap of faith. Either present more force to the hill to prepare for an enemy force, or avoid that hill all together. At least you know that the hill is not empty, if nothing else. Keeping track of a commander's assumptions and the degree to which they were actually true is a good way to capture one of the essential skills in good planning. Since CPR already includes the *Assumptions* class, simply maintaining a collection of them within the *Outcome* class, along with corresponding truth values, is adequate to take the truthfulness of assumptions into account.

Another important piece of the planning puzzle is to know what you *could have* done if you had not undertaken this particular plan. Some opportunities are particular, and don't come knocking again. Time constraints, resources limitations, etc. leave only a certain window of opportunity for many plans. This opportunity cost should be taken into account when recording an experience. Squandering useful opportunities can have a huge impact on the success of an overall campaign. To keep track of these costs, we can use the existing CPR objects of *Plan* and *Action* to keep track of full plans or single actions that cannot be undertaken because of the events unfolding in this experience. Of course, many of these cannot simply be known at the outset, so these should only be stored if they were already prepared plans that are now impossible, or obviously actions that were noted to be operationally advantageous that are now impossible.

With these four main areas (Objectives, Costs, Assumptions, and Opportunity Cost) in play, we can now begin to develop the *Outcome* class, which will extend *PlanObject*. This is shown below in Figure 3.

Outcome
ObjectivesMet : Objective
ObjectivesFailed : Objective
FriendlyCost : Cost
AdversaryCost : Cost
AssumptionsTrue : Assumption
AssumptionsFalse : Assumption
OpportunityCostP : Plan
OpportunityCostA : Action

Using this class as a single point of collection for various types of information about the success or failure of a military plan allows objective scoring mechanisms to examine the case from their own perspective. For example, if a member of the Red Cross was looking at a plan, they may gauge the success of the experience as a function of casualties and medical supplies used. From a different perspective, an intelligence collector may gauge the success of a plan based on what information sources have become unavailable because of the plan compared to how many new information sources were created/revealed. The success of a plan is all a matter of perspective; however each perspective requires similar kinds of information with which to establish the context.

Going beyond simply storing information, we can apply some forms of objective scoring. If someone were looking at this plan not from a specific point of view, but rather from a broad and general sense, they may want to see some overall measures of this experience's success. To accomplish this, we can look at how important each element was to the commander's vision of the plan. For example, if certain objectives were more important that others, they will simply have a higher weight. If certain assumptions are more critical and certain than others, they could have a higher weight. Then, when determining the overall, general success of the experience, we can simply use those weights to determine the ratio between what elements were successful and not successful. In other words, we can take a simple ratio to see which of the most important objectives were met. We can apply the same sort of reasoning to costs. Applying a weight to resources can allow us to judge by simple ratio the comparative cost between commander and adversary. Constructing ratios based on importance gives a way to examine the success of the plan in terms of what the commander doing the planning deemed important. These ratios can provide a quantitative way to broadly look at the success of a plan. This can be used to scan through an experience base quickly to 'weed out' certain plans, and focus attention on others based on success in one or more areas. You can look at highly effective plans, in terms of accomplishing goals. Also you could seek out plans with good use of judgment and intelligence information with a high rate of true assumptions.

You can look for plans with a cost ratio that is very high; plans that are exceptionally costly to the adversary. Or you could look for plans that leave your options the most open for future planning. Feasibly, using these weighted scores you could aggregate them to find the best overall plan (by the numbers).

Cost

Being able to tell how much a plan costs is an important factor in plan selection. However, the current version of Core Plan Representation has no way to take this into account. We must first understand what cost is, in this context. Cost is a resource that was expended to enable your own actions. Cost differs from simple negative effects in that it implies expenditure in something that you already own and control (in some way). In this context, costs can be expected to occur, or could arise unexpectedly, but all pertain to resources that are under your command. In a sense, cost is an element of control, because the resources expended to enable your actions can be perceived as a management of risk. Being able to perform a cost-benefit analysis gives a potential plan further context. Representing those costs in our plan representation further enables reasoning upon plans and cases. However, this challenge is two pronged. First, we must be able to represent costs themselves. Next, we must be able to differentiate between expected costs and real costs.

In order to represent costs, we must first examine the variety of things that could be expended in the course of a plan. Then, we must have the ability to understand the level of fidelity with which we can model those various factors in regards to the level of fidelity of the overall plan. For example, although employing a tank in battle requires fuel, ammunition, and other supplies, those materiel items were all at some point purchased. From that abstract perspective, you can express that cost in simple dollars and cents. However, from a more concrete perspective, if the tank needs to be deployed rapidly, and purchasing that fuel would take more time that you have, then you need to rely on the materiel *on hand* to accomplish a task. From that angle, it’s not the dollars you spent that counts, but rather the value of the actual materiel in terms of its limited supply. The logistical ability expended to attain supplies in that case is worth more than the actual cost of the materiel itself in a vacuum. Depending upon the situation, there is a difference between the simple dollar cost and the *value* of the materiel used in a plan.

Extending the scope of this kind of thinking, simple dollars and cents aren’t even enough to explain the variety of costs that can be incurred during a plan. Casualties, for example, are one of the hardest costs in warfare. Since every human life is important, how can we simply attach a dollar value to it? Another cost that could be incurred could be expressed in terms of political capitol. Sometimes when you have to call in a favor, you might find your relationship with that person strained. You obviously can’t put a dollar amount on someone’s trust, respect, or friendship. Obviously, we must examine methods of qualifying and quantifying cost that go beyond simple monetary exchanges.

This is a good opportunity to use the PMESII framework to help conceptualize the different kind of costs that can be incurred by a plan.

Dimension	Types of Costs
Political	Primacy of your control over the situation
Military	Casualties, Logistics stock expended, Equipment lost, etc.
Economic	Money expended
Social	Relationships strained (with others, public, leadership, etc.)
Information	Intel revealed to adversary
Infrastructure	Transportation means inhibited

From a political point of view, the costs you could incur from a plan would have to do with your control over the situation. Do you have the support of your superiors? Are coalition partners cooperating? Do local civil authorities allow you to operate? Are they cooperative? Does your own force still adhere to your command? These concerns are primarily concerned with relationships, but of an official rather than personal nature. The parties concerned could still trust and respect you, but their political obligations drive their behavior otherwise.

Costs from the military perspective are fairly simple and straightforward. Casualties, materiel consumed, and the like are easily quantifiable and have been used as metric for cost throughout military history.

The same goes for costs from the economic perspective. Tracking the money spent on an operation is a straightforward process.

Costs from the social perspective can create a considerable challenge in terms of measurability. Sometimes it is hard to gauge if public support is on your side, or if morale is at a certain level, or if the contacts in your rolodex are really behind you. Capturing and quantifying that kind of information can be difficult, if not entirely subjective. However, using technologies like Social Network Analysis, this seemingly impossible task can be accomplished.

Measuring cost from the information perspective can be both challenging, but also fairly straightforward at the same time. It’s straightforward because you generally know what sort of information you had to make known in some way in the course of carrying out your plan. It can be challenging because you have no idea what other people perceive of that information. In other words, you may know full well you had to use an otherwise unknown capability in order to surprise an opponent, but you don’t know if the opponent realized that’s what was going on. On the other side, you may believe you are operating quite unpredictably, but a clever opponent has determined something about your patterns and doctrine through several observations. It’s hard to tell what exactly another person knows. In the face of this uncertainty, perhaps the best way to track information cost is to understand what information you know that you are making known, assuming the adversary has perceived it.

Costs from an infrastructure perspective are also fairly straightforward to measure, much like military and economic costs. All that is required is the overall capabilities in infrastructure that existed before, and how the plan has inhibited that in any way.

Event

One of the major drawbacks of the Core Plan Representation was its lack of support for representing red force actions. In other words, although it had all sort of things in mind for planning from your own perspective, it had no support for tracking what you believed an adversary might do, or what they actually did. To alleviate this, we took an approach that used similar classes for modeling assumptions about the adversary, and actual adversary actions. In this section, we will discuss modeling actual adversary actions using adaptations to Core Plan Representation.

Under this approach, a new class very similar to Plan was formulated, called Event. Event is used to hold information about what actions some other group undertook, without direct knowledge of their intentions or world state. In this sense, it is similar to Plan, but trimmed down to accommodate the lack of knowledge of the observed group's subjective universe.

Location

In order to model concepts of location that are not strictly geo-spatial, we had to explore what sort of construct could be used to represent SpatialSpec in CPR. As you remember, each Action can contain only one SpatialSpec. Therefore, we are left with two options. First, we could store more data inside of SpatialSpec as to leave a place for the many and varied aspects of location that we could conceive. Second, we could alleviate the restriction in the Action class by allowing it to store any number of instances of information to denote locations. Because it is not known how many different aspects of a location might be modeled a priori, we decided that each class should store only one type. This makes each instance highly cohesive, since it only stores data on one aspect of location. Therefore, we lifted the restriction on the Action class, allowing it to store any n SpatialSpec. Also, because of the changed nature of the SpatialSpec, to allow several different aspects of locations to be modeled, we also renamed the class Location.

Now, let's examine some ways we can facilitate the reasoning upon these experiences using meaning and analogies to better facilitate the CBR process.

Meaningful Information

Determining what information a computer can understand is one of the biggest challenges facing knowledge representation. Given the vast pattern matching and semantic abilities of the human brain, it can be difficult to prune out what exactly can be understood by an entity without any ability to intuit context, subtlety, or inferred meaning. However, given the field of military planning, we have at least a stepping stone into forming information that can be reasoned upon by a computer. The challenge is: how do we ourselves understand the context of military planning in such a way that can be operationalized by both a computer and a human in mixed-initiative planning? That is, how can we make military planning information understandable to a computer, without condensing it to the point of losing its human meaning?

One of the biggest challenges in the traditional computing world is qualitative information. Computers are great with arithmetic, and humans enjoy the concreteness afforded by 'pinning down' something numerically, but how to we reconcile that to a world that refuses to be pinned down? Determining qualitative information, that is taking an abstract concept and acknowledging a discernable concrete quality or trait, is a powerful human ability. Even so, not every situation affords the opportunity even to surmise qualitative inferences. Every day humans are faced with situations where they can't just pin things down in the chaos, so how can we expect computers to be any better?

One way we can begin to understand qualitative information is in relative terms. In other words, whether we say 'big' or 'small' we are talking about two values in terms of size. These two terms are relative; we would not know what things were big unless there were things that were small to compare them to. This brings up the crux of the approach, that the two values are *comparable*. Given that they are, we can set up an ordering scheme (either strict or partial ordering depending upon the situation) that a computer could understand.

Another way to allow a computer to understand qualitative information is taking the simplest approach possible. That is allowing each and every piece of information to be tagged with qualities. These qualities would be simple strings that could be compared directly, or arranged in a taxonomy that would allow them to be compared in terms of pre-determined meaning. In other words, we can create a dictionary of qualities, and arrange them in some reasonable tree (or graph) that allows them to carry meaning in terms to their relationships with each other. This, of course, would require tremendous space and effort to create. However, there is already ongoing ontological research that covers this arena of knowledge representation extensively (e.g. Cyc).

Analogies

An important area of research in artificial intelligence is the field of understanding and forming analogies. In a sense, our power to create relationships between symbols, experiences, and situations is at the root of our ability to apply that experience to our everyday life. We would not know what experiences were valid to draw upon if we did not know how to relate the past to the present. In the same token, we would not know how to adapt those situations to the present context if it were not for that same power. Thus, if we are to truly accomplish mixed-initiative planning, we need to understand how to enable a computer to draw analogies.

There are several approaches we can take. If we think for a moment, we can easily conceive of why two items, say apples and oranges are actually similar. There are several ways to approach that similarity, most notable three.

First, apples and oranges share some similar physical characteristics. They are both sweet, both about the same size, and both have an outer skin. We can examine those characteristics through measurement and observation. In this same sense, if we are trying to draw an analogy in war planning, we can say that a WWII-era Sherman tank and a modern M1A Abrahams tank are similar. Both have guns, both have treads, and both have crews of more than one person. They share these characteristics, which we can determine through an examination of the two items. This of course takes not only time, but also the ability to measure those characteristics. All of this is under the assumption that the two sets of characteristics can even be compared.

A second approach to determining if apples and oranges are similar is examining them not by looking at their attributes, but how they are organized and classified in various taxonomies. In this example, apples and oranges can both be classified as 'fruits', so in this sense they are similar. They can also both be classified as member of the plant kingdom, although they are from different (phylum?) genuses. In this sense, the two are similar, but not as similar as they could be. From this perspective, knowing the taxonomical structure of how we classify plants and animals gives us the ability to granulate the level to which two things are similar. In this context, cats and dogs would be more similar than cats and bamboo plants. Because cats and dogs share a common classification that is at a deeper level of the taxonomy, we can consider them more similar. The asset of this approach is that we don't need to be able to measure of compare the attributes of the items to be matched, as long as they are already classified in some taxonomical structure. This is also the drawback of this approach; items must already be classified, or else you must find a way to classify them yourself, which would require an examination of its attributes (which returns you to the first approach). A military planning example of this would be determining that our Sherman and Abrahams tanks are again similar, because they are both classified as Main Battle Tanks. In this sense, the Sherman and the Abrahams are more similar to each other than the Abrahams and the F-16 are, because the F-16 is a fighter jet and not a Main Battle Tank.

However, following these two approaches blindly may not be enough to adequately make analogies. A fighter jet and a main battle tank actually do share some things in common, even if they are classified differently. Both are able to lay destruction upon targets. Both are capable of being a show of force. Both are capable of protecting soldiers (and airmen) from harm. Even though both items may go about these capabilities differently, they can in a sense be interchanged for each other in the proper situation. Examining not just the physical characteristics of an item is not enough for drawing analogies. Especially in military planning, where using available resources to utilize a capability is supremely important. It is not as important that you have two similar vehicles, what is important is that the mission is accomplished effectively.

Pulling away from an equipment-centric point of view, even two different approaches can be analogous to each other if they accomplish the same effect. For example, a show of force flight, a pamphlet drop, and a bombing run might actually all accomplish the same goal. Knowing this allows the commander to plan in such a way that optimizes costs, benefits, and risk to make more effective decisions.

We can, then, examine analogies not simply as the similarity between the objects involved, but in terms of capabilities and effects. From this perspective, analogies are drawn based on what you want to accomplish, not just on the attributes that you can examine. Going deeper, we can think of those capabilities and effects as relationships between the actor and the affected. If we know the relationship we desire, and the target object, we know what sorts of capabilities are required to satisfy that relationship. From this, we can draw analogies from our past experiences to see what exact items have satisfied this relationship in the past.

Structural Mapping

Important work in this style of analogy building has been ongoing. One interesting algorithm is the Structural Mapping Algorithm, invented by Diedre Gentner and Ken Forbus (Forbus et al. 1989). This algorithm relies on forming an analogy not just on the objects involved, but rather on their relationship to each other. In other words, SME utilizes the structure of a representation, rather than simply the attributes of the constituent parts.

More specifically, SME seeks to make analogies that adhere to the *systematicity principle*, which stipulates that matches are preferred when they involve higher-order relations rather than several lower-order facts. In other words, SME matches based not only on attribute matches (like similarity), but more importantly on true analogy, where the matching elements fit into a system.

Many are Called, Few are Chosen

An interesting approach to drawing analogies in an efficient and interesting way is the Many are Called but Few are Chosen (MAC/FAC) approach designed by Gentner and Forbus from Northwestern University (Gentner & Forbus 1994). It combines a speedy similarity based approach with a deeper structural matching approach to optimize analogy drawing. This approach is interesting in two ways. First, it can make episodic reasoning more efficient in terms of matching and retrieval. Second, the approach itself can help a commander understand his or her priorities in planning. Let us first briefly describe the approach itself, and then this important second benefit. After, we will explore various ways the approach could be adapted in novel ways, prompting future research.

In the MAC stage, each case is assigned a vector of numbers. This vector contains the number of occurrences of a given description of *functors* (predicates, functions, connectives) contained within the case. In other words, if a description contained *[DESTROYS, BIG]*, then each case would be assigned a vector that simply counted the number of occurrences of the function *DESTROYS* and the predicate *BIG* in the case. This is done regardless of the structure of the case; it is simply a count. Then, the MAC chooser takes a look at the vectors compared with the description and selects the best match and everything within 10% of that. These cases are passed onto the FAC stage to examine their structure.

In the FAC stage, the pruned case-base is subjected to structural matching with the current situation, as with the traditional Structural Matching approach, using similarity based matching. The combination of the two approaches allows the otherwise thorough and possibly costly Structural Matching algorithm to operate in a concise and focused way; eliminating wasteful work. However, the MAC/FAC approach operates on the assumption that a 'first-pass' with a desired description will eliminate cases with little value to the case to be matched. Bear in mind that if this assumption does not hold, then there is a possibility to prematurely eliminating useful results. However, Gentner and Forbus' experiments with the algorithm are suggestive that it is well-founded.

Summary

Case-based reasoning is a method by which we can use the past to determine possible solutions to a problem. However, the foundation of a good case-based reasoner is a solid case-base. Formalizing how an experience is captured, stored, and reasoned upon is the brick and mortar upon which a potentially powerful decision aide could be built. By exploring and adapting existing plan representation work, we can formalize a case structure for military planning. Then, by employing analogy-based matching algorithms, we can begin to understand how to make insightful connections between the past and the present, in order to shape the future.

Future Work

Further research into this topic can branch into formalizing the measurement of trust and confidence in information for making assumptions, determining the right level of fidelity in information capturing to ensure good case matches, determining the best way to capture a world state for provide context, methods by which to capture cases from written (or lived) historic war stories, and a variety of other areas.

References

Forbus, Ken, Dedre Gentner, and Brian Falkenhainer. "The Structure-Mapping Engine: Algorithm and Examples." Artificial Intelligence 41(1989): 1-63.

Forbus, K., Gentner, D. & Law, K. 1994b MAC/FAC: A model of similarity-based retrieval. *Cognitive Science* **19**(2), 141-205.

Pease, Adam. "Core Plan Representation - Request for Comment Version 4." Core Plan Representation. 06 NOV 1998. Teknowledge. 5 Feb 2007 <http://projects.teknowledge.com/CPR2/Reports/CPR-RFC4/>.

Pease, Adam, and Capt Todd M Carrico. Object Model Working Group Core Plan Representation. Dayton OH: United States Air Force Armstrong Laboratory, 1996.

De Mantaras, Ramon Lopez et al.. "Retrieval, reuse, revision, and retention in case-based reasoning." The Knowledge Engineering Review 00(2005): 1-32.

Appendix A: CPR Version Charts

1996:

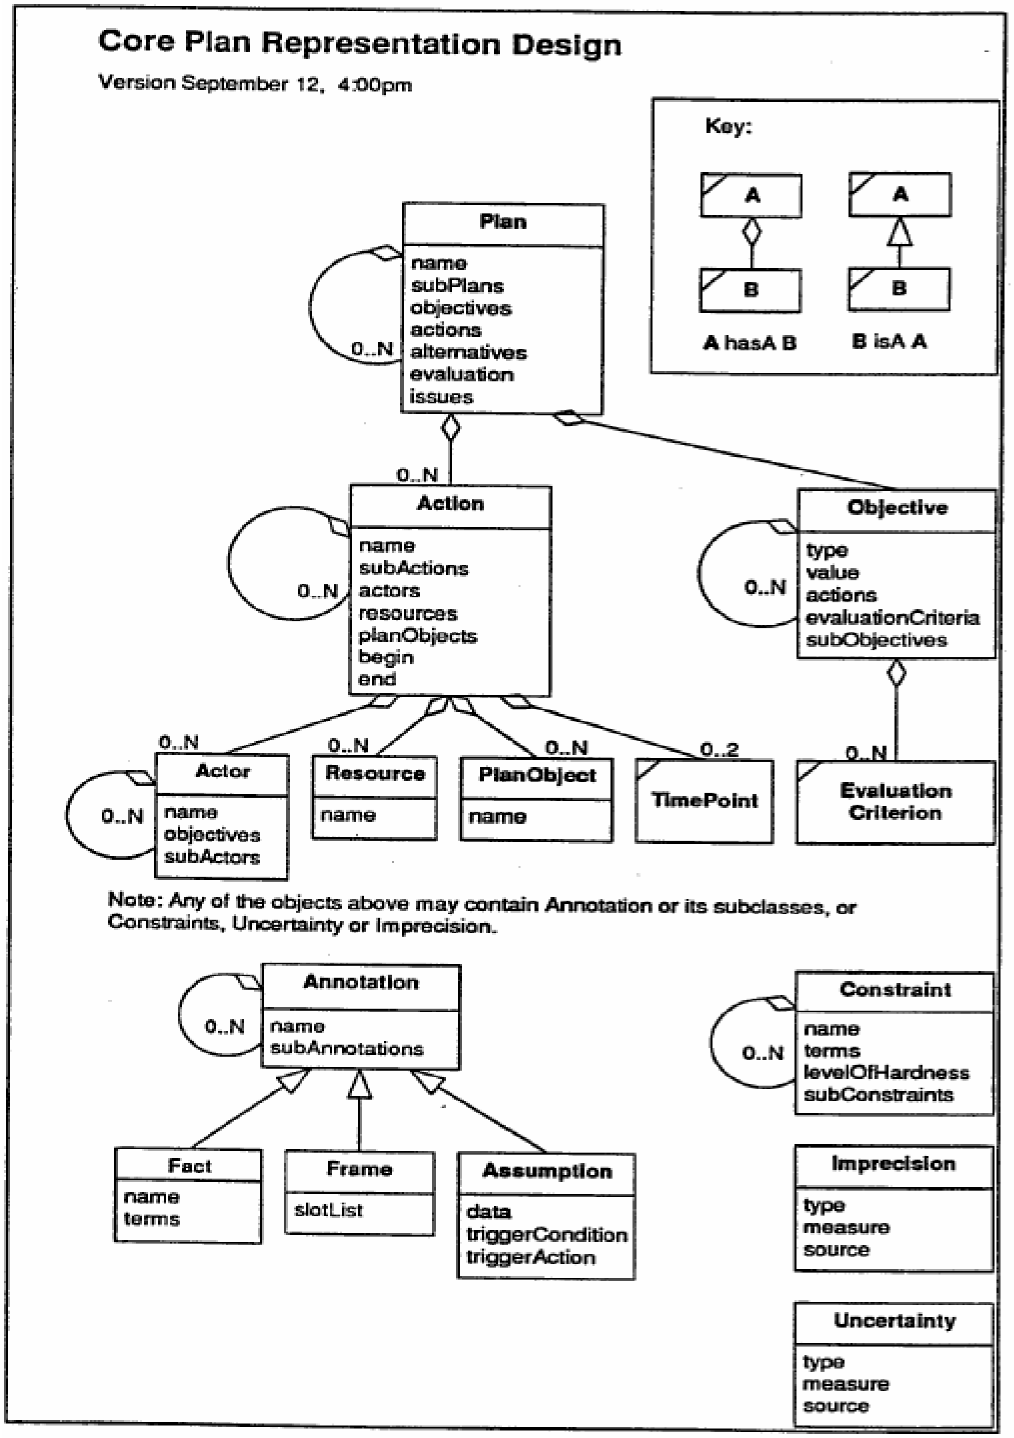

1998:

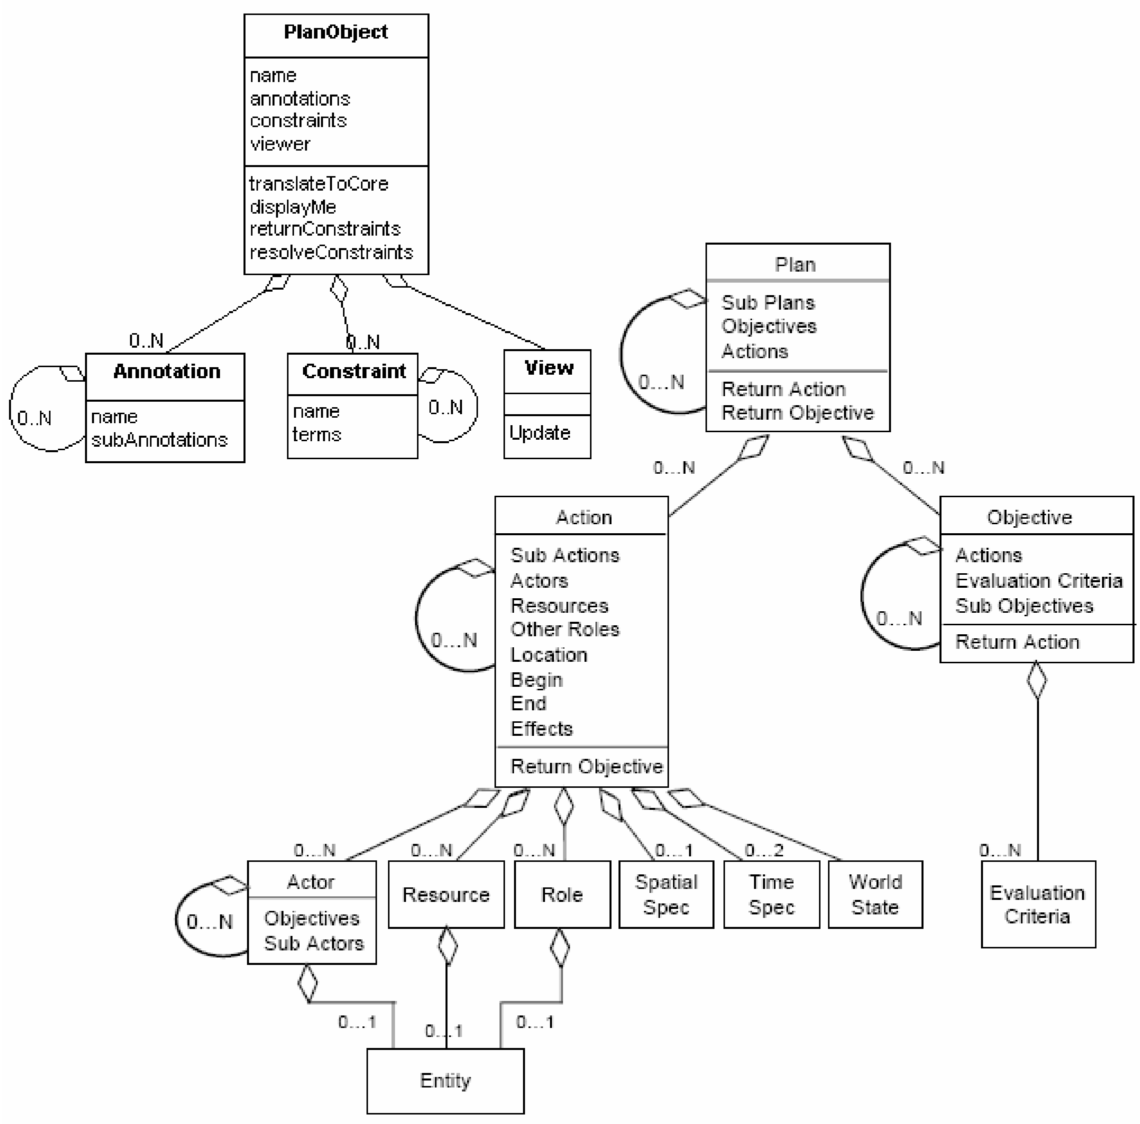

PlanObject
name
annotations
constraints
viewer
translateToCore
displayMe
returnConstraints
resolveConstraints
0..N
0..N
Annotation
name
subAnnotations
0..N
Constraint
name
terms
0..N
View
Update
Plan
Sub Plans
Objectives
Actions
Return Action
Return Objective
0…N
0…N
0…N
Action
Sub Actions
Actors
Resources
Other Roles
Location
Begin
End
Effects
Return Objective
0…N
Objective
Actions
Evaluation Criteria
Sub Objectives
Return Action
0…N
0…N
0…N
0…N
0…1
0…2
0…N
Actor
Objectives
Sub Actors
0…N
Resource
Role
Spatial Spec
Time Spec
World State
Evaluation Criteria
0…1
0…1
0…1
Entity