Using Guidelines to Constrain Interactive Case-Based HTN Planning*

Héctor Muñoz-Avila[†‡], Daniel C. McFarlane[‡], David W. Aha[‡], Len Breslow[‡], James A. Ballas[‡], & Dana S. Nau[†]

† Department of Computer Science
University of Maryland
College Park, MD 20742-3255
lastname@cs.umd.edu
(301) 405-2684 | FAX: 405-6707
‡ Navy Center for Applied Research in AI
Naval Research Laboratory
Washington, DC 20375
lastname@aic.nrl.navy.mil
(202) 404-4940 | FAX: 767-3172

Abstract. This paper describes *HICAP*, a general-purpose, interactive case-based plan authoring architecture that can be applied to decision support tasks to yield a hierarchical course of action. It integrates a hierarchical task editor with a conversational case-based planner. HICAP maintains both a task hierarchy representing guidelines that constrain the final plan and the hierarchical social organization responsible for these tasks. It also supports bookkeeping, which is crucial for real-world large-scale planning tasks. By selecting tasks corresponding to the hierarchy's leaf nodes, users can activate the conversational case-based planner to interactively refine guideline tasks into a concrete plan. Thus, HICAP can be used to generate context sensitive plans and should be useful for assisting with planning complex tasks such as noncombatant evacuation operations. We describe an experiment with a highly detailed military simulator to investigate this claim. The results show that plans generated by HICAP were superior to those generated by alternative approaches.

1 Introduction

Planning a course of action is difficult, especially for large hierarchical organizations (e.g., the U.S. Navy) that assign tasks to elements (i.e., groups or individuals) and constrain plans with guidelines (e.g., doctrine). In this context, a concrete plan must adhere to guidelines but should also exploit organizational knowledge where appropriate (e.g., standard procedures for solving tasks, previous experiences when reacting to unanticipated situations). Case-based reasoning (CBR) can be used to capture and share this knowledge.

* Presented at the 1999 International Conference on Case-Based Reasoning

In large planning environments, automatic plan generation is neither feasible nor desirable because users must observe and control plan generation. We argue that, rather than relying on an automatic plan generator, users prefer and can greatly benefit from the assistance of an intelligent plan formulation tool with the following characteristics:

- *Guidelines-driven*: Uses guidelines to constrain plan generation.
- *Interactive*: Allows users to edit any detail of the plan.
- *Provide Case Access*: Indexes plan segments from previous problem-solving experiences, and retrieves them for users if warranted by the current planning scenario.
- *Perform Bookkeeping*: Maintains information on the status of and relations between task responsibilities and individuals in the organizational hierarchy.

This paper describes HICAP, a general-purpose plan formulation tool that we designed to embody these characteristics.[1] HICAP (Hierarchical Interactive Case-Based Architecture for Planning) integrates a task decomposition editor, HTE (Hierarchical Task Editor) (Muñoz-Avila et al., 1998), with a conversational case-based planner, NaCoDAE/HTN. The former allows users to edit and select guidelines for refinement, while the latter allows users to interactively refine plans encoded as hierarchical task networks (HTNs) (Erol et al., 1994). Refinements use knowledge of previous operations, represented as cases, to augment or replace standard procedures.

The following sections describe the application task, HICAP's knowledge representation, its architecture, its empirical evaluation, and a discussion of related work.

2 Planning Noncombatant Evacuation Operations

Noncombatant evacuation operations (NEOs) are conducted to assist the U.S.A. Department of State (DoS) with evacuating noncombatants, nonessential military personnel, selected host-nation citizens, and third country nationals whose lives are in danger from locations in a host foreign nation to an appropriate safe haven. They usually involve the swift insertion of a force, temporary occupation of an objective (e.g., an embassy), and a planned withdrawal after mission completion. NEOs are often planned and executed by a Joint Task Force (JTF), a hierarchical multi-service military organization, and conducted under an American Ambassador's authority. Force sizes can range into the hundreds and involve all branches of the armed services, while the evacuees can number into the thousands. More than ten NEOs were conducted within the past decade. Publications describe NEO doctrine (DoD, 1994), case studies (Siegel, 1991; 1995), and more general analyses (e.g., Lambert, 1992).[2]

[1] Implemented in Java 2, the HICAP applet can be run from www.aic.nrl.navy.mil/hicap. HICAP was introduced in (Muñoz-Avila et al., 1999), which did not include the evaluation described here.

[2] See www.aic.nrl.navy.mil/∼aha/neos for more information on NEOs.

The decision making process for a NEO is conducted at three increasingly-specific levels: strategic, operational and tactical. The strategic level involves global and political considerations such as whether to perform the NEO. The operational level involves considerations such as determining the size and composition of its execution force. The tactical level is the concrete level, which assigns specific resources to specific tasks.

JTF commanders plan NEOs in the context of doctrine (DoD, 1994), which defines *general* guidelines (e.g., chain of command, task agenda) for designing strategic and operational plans; tactical considerations are only partly addressed. Doctrine is abstract; it cannot account for the detailed characteristics of specific NEOs. Thus, JTF commanders must always adapt doctrine to a NEO's specific needs, and do so in two ways. First, they dynamically modify doctrinal guidance by eliminating irrelevant planning tasks and adding others, depending on the operation's needs, resource availabilities, and relevant past experiences. For example, although NEO doctrine states that a forward command element must be inserted into the evacuation area with enough time to plan the insertion of the JTF's main body, this is not always feasible (e.g., in *Operation Eastern Exit*, combined elements of the JTF were inserted simultaneously due to the clear and imminent danger posed to the targeted evacuees (Siegel, 1991)). Second, they employ experiences from previous NEOs, which complement doctrine by suggesting tactical refinements suitable for the current NEO. For example, they could draw upon their previous experiences to identify whether it is appropriate to concentrate the evacuees in the embassy or to plan for multiple evacuation sites.

3 Knowledge Representation

Because HTNs are expressive representations for plans, we used a variant of them in HICAP. A HTN is a set of tasks and their ordering relations, denoted as $N = \langle \{T_1, \ldots, T_m\}, \prec \rangle$ $(m \geq 0)$. The relation $\prec$ has the form $T_i \prec T_j (i \neq j)$, and expresses temporal restrictions between tasks.

Problem solving with HTNs occurs by applying *methods* to decompose or reduce tasks into subtasks. Each method has the form $M = \langle l, T, N, P \rangle$, where l is a label, T is a task, N is a HTN, and $P = \langle p_1, \ldots, p_k \rangle$ is a set of preconditions for applying M. When P is satisfied, M can be applied to a task T to yield N.

HICAP's HTN consists of three task types. First, *non-decomposable* tasks are concrete actions and can occur only at a network's leaves. Next, *uniquely decomposable* tasks correspond to guideline tasks (e.g., doctrine), and are solved by unconditional methods ($k = 0$). Finally, *multi-decomposable* tasks must be solved in a specific problem-solving context.

There are two sources of knowledge for decomposing multi-decomposable tasks: standard operating procedures (SOPs) and recorded episodes. SOPs describe how to reduce a task in a typical situation. Recorded episodes describe how tasks were reduced in situations that are not covered by SOPs. In our representation, SOPs and recorded episodes are both represented as methods and we loosely refer to both as *cases*. However, there is an important difference in the way SOPs and recorded episodes are applied. To apply a SOP to reduce a task, all its preconditions must be matched because they are typically rigid in their use. In contrast, recorded episodes can be applied to reduce a task even if some of its preconditions are not satisfied.

When reducing a task T, HICAP retrieves all cases (i.e., standard procedures and recorded episodes) that can decompose T. If all the preconditions of a SOP are met, then it should be used to decompose T. Otherwise, a case corresponding to the most similar episode should be used. For example, standard NEO procedures state that the evacuees must be concentrated in the embassy prior to troop deployment, but this is not always possible: in Operation Eastern Exit, only some of the evacuees were concentrated in the embassy after the Joint Task Force was deployed. This occurred because escorted transports were not available to gather these evacuees, who were unable to reach the embassy due to the dangerous conditions in the surrounding areas (Siegel, 1991). Likewise, the evacuees of *Operation Sharp Edge* (Sachtleben, 1991) were concentrated in several places, forcing multiple separate evacuations.

4 HICAP: An Interactive Case-Based Planner

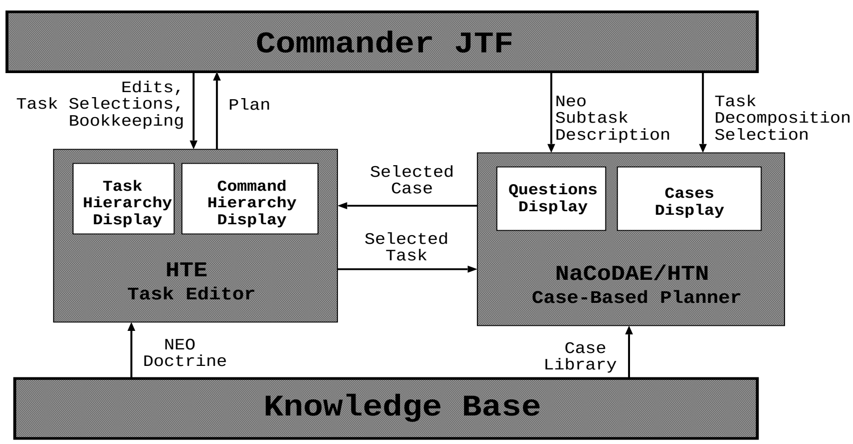

Fig. 1. The HICAP architecture.

HICAP (Figure 1), which integrates HTE with NaCoDAE/HTN, inputs a HTN describing the guidelines for an application along with a set of cases for each multi-decomposable subtask. It displays all uniquely decomposable tasks as expanded. Under user control, HICAP outputs an elaborated HTN whose leaves are concrete actions as specified by case applications and manual edits. In this way, HICAP satisfies the requirements stated in Section 1. First, all plans formulated using HICAP are in accordance with the guidelines or user modifications of them. Second, HICAP supports interactive task editing and triggers conversations for tasks that can be decomposed by case application. Third, it incorporates knowledge from previous problem solving episodes as cases, which serve as task decomposition alternatives. Finally, it allows users to visually check that all tasks are assigned to JTF elements, and to record/update their completion status.

4.1 Hierarchical Task Editor

In complex environments where dozens of tasks must be performed by many people, tracking the completion status for each task can be challenging. For example, during the NEO Operation Eastern Exit, the task to inspect evacuees prior to embarkation was not assigned (Siegel, 1991). One of the evacuees produced a weapon during a helicopter evacuation flight. Although it was immediately confiscated, this oversight could have resulted in tragedy and illustrates the difficulties with planning NEOs manually.

The Hierarchical Task Editor (HTE) (Muñoz-Avila et al., 1998) serves HICAP as a bookkeeping tool to track the status of each task. HTE inputs a knowledge base consisting of a HTN task agenda, its ordering relations, the organization's command hierarchy, and an assignment of tasks to command elements. It allows users to edit the knowledge base and select tasks to refine by invoking NaCoDAE/HTN, thus tailoring the plan to the particular circumstances of the current NEO.

For our NEO application, we encoded a HTN to capture critical planning doctrine (DoD, 1994), yielding 200+ tasks and their ordering relations. Next, we used this doctrine to elicit the JTF command hierarchy commonly used in NEO operations. Finally, we elicited relations between tasks and the JTF elements responsible for them. The mapping of tasks to command elements is many-to-one. Figure 2 displays (left) the top level tasks that, according to doctrine, must be performed during a NEO and (right) the elements in the JTF responsible for them.

4.2 Conversational Task Decomposer

NaCoDAE/HTN, an extension of the NaCoDAE conversational case retrieval tool (Aha & Breslow, 1997; Breslow & Aha, 1997), supports HTN planning by allowing users to refine selected tasks into concrete actions. When given a task T to refine by HTE, NaCoDAE/HTN uses T as an index for initial case retrieval and conducts an interactive *conversation*, which ends when the user selects a case $C = \langle l, T, N, P \rangle$. Network N is then used to decompose T (i.e., into a set of subtasks represented as T's child nodes). Subtasks of N might themselves be decomposable, but non-decomposable tasks corresponding to concrete actions will eventually be reached. Task expansions are displayed by HTE.

During conversations, NaCoDAE/HTN displays the labels of the top-ranked cases that can decompose the selected node and the top-ranked questions from

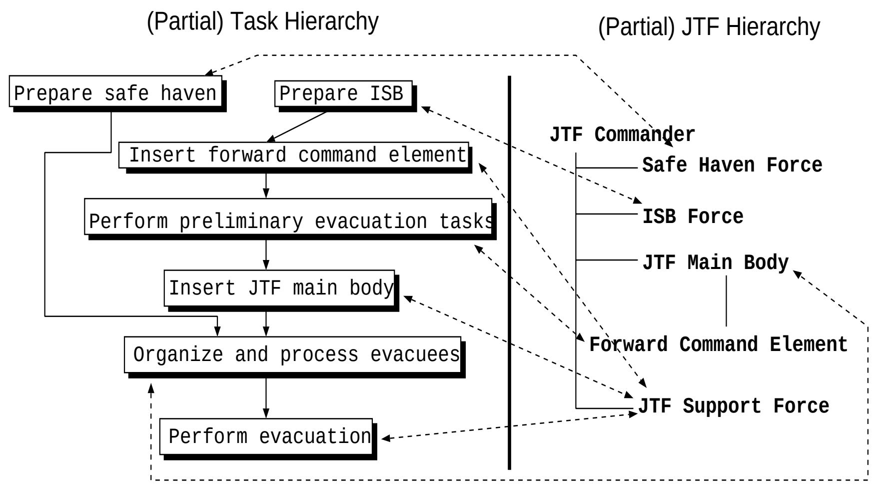

Fig. 2. Top level NEO tasks and their assignment to JTF command elements (double arrows denote assignments; arrows denote task orderings; ISB = intermediate stage base).

these cases whose answers are not yet known for the current situation. The user can select and answer any displayed question; question-answer pairs are used to compute the similarity of the current task to its potential decomposition methods (cases). Cases are ranked according to their similarity to the current situation (Aha & Breslow, 1997), while questions are ranked according to their frequency among the top-ranked cases. Answering a question modifies the case and question rankings. A conversation ends when the user selects a case for decomposing the current task.

Some of the displayed cases are standard procedures; they can only be selected to decompose a task after all of their questions have been answered and match the current planning scenario. That is, preconditions of the standard procedures must match before they can be applied. In contrast, cases based on previous experiences can be selected even if some of their questions have not been answered, or if the user's answers differ. Thus, they support partial matching between their preconditions and the current planning scenario.

5 Example: NEO Planning

During NEO planning, users are first shown the tasks corresponding to doctrine, and revise them as needed. They can expand any task and view its decomposition. In Figure 3, the user has selected the *Select assembly areas for evacuation & Evacuation Control Center sites* task, which is highlighted together with the command element responsible for it.

Standard procedure dictates that the embassy is the ideal assembly area. However, it is not always possible to concentrate the evacuees in the embassy.

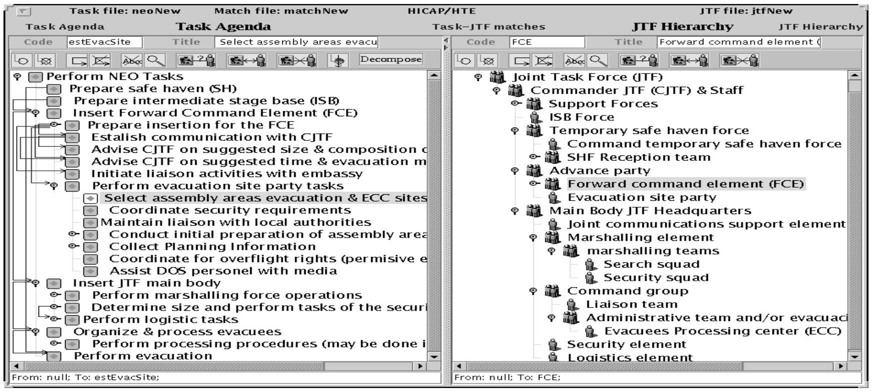

Fig. 3. HTE: Task agenda (left) and command hierarchy (right) displays (arrows denote ordering constraints).

Alternative methods can be considered for decomposing this task. When the military planner selects this task, HICAP displays the alternatives and initiates a NaCoDAE/HTN conversation (see Figure 4 (top)).

If the user answers *Are there any hostiles between the embassy and the evacuees?* with *uncertain*, a perfect match occurs with the case labeled "Handle situation in which it is unknown whether hostiles are present," which now becomes the top-ranked case (Figure 4 (bottom)). Figure 5 (left) shows the decomposition when selecting this case to decompose this task in HTE; two new subtasks are displayed, corresponding to this case's decomposition network. *Send unmanned air vehicle to* ... is a non-decomposable concrete action. If the user tells HICAP to decompose *Determine if hostiles are present*, HICAP will initiate a new NaCoDAE/HTN dialogue (Figure 5, right).

The user can again prompt a dialogue by selecting the *The UAV detects hostiles* alternative and decomposing its subtasks. This cycle, in which HICAP displays alternatives and the user answers questions and selects an alternative, continues until non-decomposable tasks (i.e., concrete actions) are reached, which form part of the final plan.

6 The Case-Based Planning Cycle in HICAP

The case-based planning component of HICAP, Nacodae/HTN, typically performs three steps: retrieval, revise, and retain. As illustrated in Section 5, the adaptation process can be viewed as embedded in the conversational retrieval process.

6.1 Case Retrieval

We previously explained that, during a conversation, cases are ranked according to the proportion of their question-answer pairs that match the current scenario.

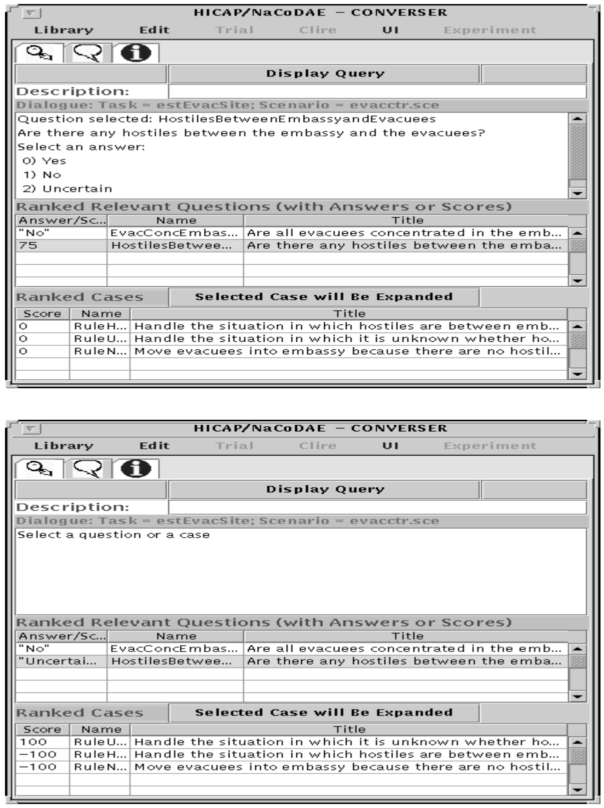

Fig. 4. NaCoDAE/HTN: Before (top) and after (bottom) answering a question. The top window lists possible answers to selected questions, while the lower windows display the ranked questions and cases.

More specifically, a case c's similarity score is computed with a query q using

$$\text{case\_score}(q, c) = \frac{\text{num\_matches}(q, c) - \text{num\_mismatches}(q, c)}{\text{size}(c)} \tag{1}$$

where num_matches(q, c) (num_mismatches(q, c)) is the number of matches (mismatches) between the states (i.e., $\langle q, a\rangle$ pairs) of q and c, and size(c) yields the number of ⟨q,a⟩ pairs in c's state.[3]

6.2 Case Revision

The user can revise the current solution by editing the task hierarchy (in HTE) and by selecting alternative cases during a NaCoDAE/HTN conversation. In

[3] Matching for numeric-valued questions is implemented using a suitable partial matching routine, but we focus on symbolic and boolean questions here.

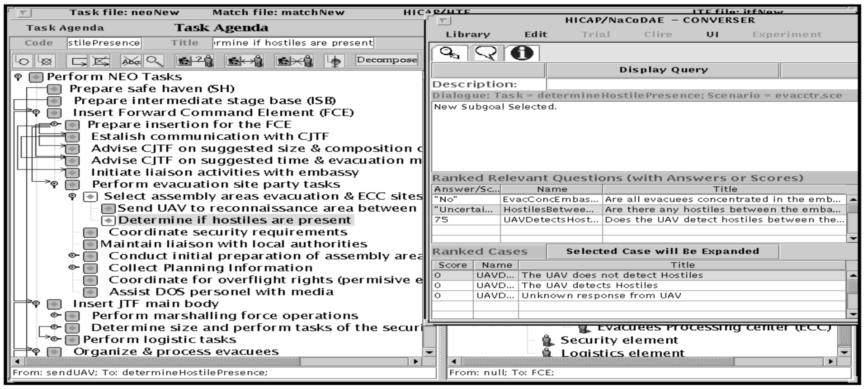

Fig. 5. HICAP's interface after selecting the *Determine hostile presence* task.

addition, the user can revise their answers to previously selected questions, which can modify case rankings. Although, revising an answer does not alter the plan automatically, the new ranks may prompt the user to change their case selection, which in turn may prompt additional edits to the task hierarchy.

This ability to explore alternatives (i.e., "what-if" analyses) is particularly important in NEO planning for two reasons. First, military planners typically plan for a main course of actions and for contingency alternatives should certain key events occur. These events may trigger changes to answers and case rankings, thus helping the user formulate these alternatives. Second, NEO planning is dynamic in nature and the user must be able to replan due to unforeseen contingencies.

6.3 Case Retention

NaCoDAE incorporates an approach introduced by Racine and Yang (1997) for maintaining case libraries. It evaluates whether any case "subsumes" another case (i.e., whether its question-answer pairs are a proper subset of the question-answer pairs of another case). If so, the subsuming case will block the subsumed case from being retrieved. A case library evaluation function alerts the user to all such pairs of cases in the case library. The user can then decide which of the two cases to revise and/or delete.

7 Empirical Validation

An experiment was run to test HICAP's effectiveness in choosing successful plans for an example NEO subtask. In particular, we showed the importance of considering episodic records over standard procedures. A larger experiment, demonstrating the capability of HICAP to generate a complete NEO plan, is currently under development.

Two researchers performed the experiment: one operated a military simulator while the other operated HICAP. A strict blind was imposed to ensure that the HICAP user had no advance knowledge concerning the simulated hostile forces, and had to take appropriate, realistic ations to acquire this knowledge. This tests HICAP's utility for planning under realistic situations where decision makers have uncertain information about the state of the world. We hypothesized that HICAP would allow users to choose a relatively successful plan from among known tactical options. HICAP's strategy was evaluated versus three other planning strategies: *random choice*, *heuristic choice*, the *most frequently used plan* used in previous NEOs. Because their definitions require explaining the scenario, we define them in Section 7.3.

7.1 The ModSAF Simulation System

We used Marine Corps SAF (MCSF), a variant of ModSAF (Modular Semi-Automated Forces), to evaluate the quality of NEO plans elicited using HICAP. ModSAF, developed by the U.S.A. Army to inject simulated auxiliary forces into training exercises, has been deployed to simulate real-world military scenarios (Ceranowicz, 1994). It is a finite state simulation with modular components that represent individual entities and parts of entities. For example, a simulated tank would have physical components such as a turret. It would also have behavioral components that represent its nominal tasks such as move, attack, target, and react to fire. Certain 3D aspects are also represented (e.g., terrain elevation, trees and vegetation, rivers, oceans, atmospheric conditions) that can affect sensory and movement behavior. The realism of ModSAF/MCSF simulations is sufficient for training exercises.

Figure 6's MCSF snapshot displays a simulated American embassy, a host country government compound, and some simulated objects. For example, a simulated transport helicopter is positioned at the heliport within the embassy site.

MCSF is a non-deterministic simulator that models several sources of stochastic variation. Some events are determined by a random number generator; others are highly sensitive to the initial startup conditions. MCSF simulates the behavior of military units in context as they follow given tactical orders. Therefore, MCSF can simulate simplified NEO subtasks in which a single planning decision determines tactical orders.

7.2 Experimental Setup

We created a NEO subtask scenario for this evaluation concerning how to move 64 evacuees from a meeting site to an embassy. The meeting site was at a crossroads in an uninhabited area outside but nearby the embassy's city. Evacuees had to be transported (8 per vehicle) through this undeveloped area, which had heavy tree cover, and out through the city to the embassy. Evacuees had to pass near a local government complex to enter the embassy grounds. This NEO context requires only a single tactical plan decision with four distinct choices:

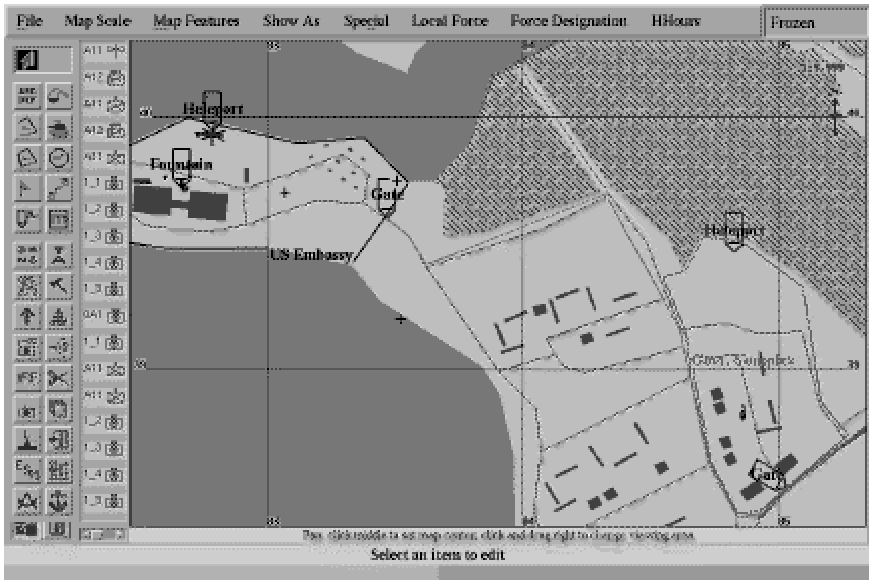

Fig. 6. A MCSF snapshot.

1. Land evacuation using 8 armored trucks
2. Land evacuation using 8 armored trucks with an escort of 8 tanks
3. Air evacuation using 8 transport helicopters
4. Air evacuation using 8 transport helicopters with an escort of 8 attack helicopters

The kind of military units used in the simulation are typical of those available to the Marine Expeditionary Units that frequently perform NEO's. A detailed terrain database of the Camp Lejeune (North Carolina, U.S.A.) area was chosen to simulate the environment. We chose this location because Marine Expeditionary Units train there for NEOs.

Two scenarios were defined that were identical except for the type of hostile forces. All hostiles were two-person dismounted infantry teams. Hostile teams in both scenarios were armed with two automatic rifles and a portable missile launcher. Each scenario included only one type of missile for hostile teams (i.e., either anti-tank missiles or anti-air missiles, but not both). These types of infantry teams, positioned in an urban environment, are typical of the kinds of hostile forces encountered in real NEO's. The positions of the hostile teams were the same for both scenarios and selected to ensure that the opposing forces will meet.

All four plan options were simulated ten times for each of the two scenarios. This resulted in 80 (2 scenarios $\times$ 4 plan choices $\times$ 10 simulations) total MCSF runs. Each of the eight plan-and-scenario combinations was repeated ten times because MCSF is non-deterministic. For example, slight differences produced

Table 1. Summaries of casualties, to individual evacuees and military teams (mean & standard deviation), averaged over 80 MCSF simulations.

	Scenario 1 (anti-tank)						**Scenario 2 (anti-air)**					
Tactical Plans	**Evacuees**		**Friends**		**Hostiles**		**Evacuees**		**Friends**		**Hostiles**	
Land	6.4	5.1	0.8	0.6	5.5	1.3	0		0		4.2	0.8
Land/Escort	3.2	10.1	7.4	1.5	6.5	1.8	0		0		7.6	0.6
Air	56.0	9.2	7.0	1.2	0		64.0	0.0	8.0	0.0	0	
Air/Escort	0		0.8	1.5	8.0	0.0	20.0	18.6	6.3	4.4	5.7	2.9

by MCSF's stochastic movement models yield strikingly different formations of friendly units when they first encountered the hostile teams. These differences can often yield drastically different simulated battle outcomes.

The HICAP user had no knowledge of the scenarios being tested; scenario information was gradually extracted through the questions prompted by NaCoDAE/HTN. That is, case-based planning was done with incomplete information about the world. Furthermore, the effects of actions were uncertain; the only way to learn the effects of an action was to actually execute it. This contrasts with traditional planning approaches that assume an action's effects are known a priori (Fikes and Nilsson, 1971).

7.3 Alternative Planning Strategies

HICAP's decision-making performance was compared with three baseline strategies. First, *random choice* simply averaged the results of all four planning choices. Second, *heuristic choice* always sent an escort, and its results were the average of the choices that include escorts. Finally, the *most frequently used* plan strategy for this subtask in recent NEOs (i.e., conducted during the past decade) was to move evacuees using escorted land vehicles.

7.4 Results

Table 7.4 summarizes the casualty results for the 80 total simulations, which each required approximately 15 minutes to run. The success measures were taken from the U.S.A. Navy's Measures of Effectiveness (MOE's) published in the Universal Naval Task List. Recommended MOEs are specified for evaluating each kind of military operation. There are several MOE's for the tactical aspects of NEO's, but only three were chosen as most important for evaluating the results of this experiment: (1) the number of evacuees safely moved, (2) the number of casualties to friendly forces, and (3) the number of casualties to hostile forces.

HICAP did not choose the same tactical plan for both scenarios. For the first (anti-tank) scenario, it chose to move the evacuees by helicopter with an attack helicopter escort. For the second (anti-air) scenario, it chose to move evacuees by armored truck with a tank escort.

HICAP's conversational case-based planning method was evaluated by comparing the success of its chosen plans to plans chosen by the other three plan

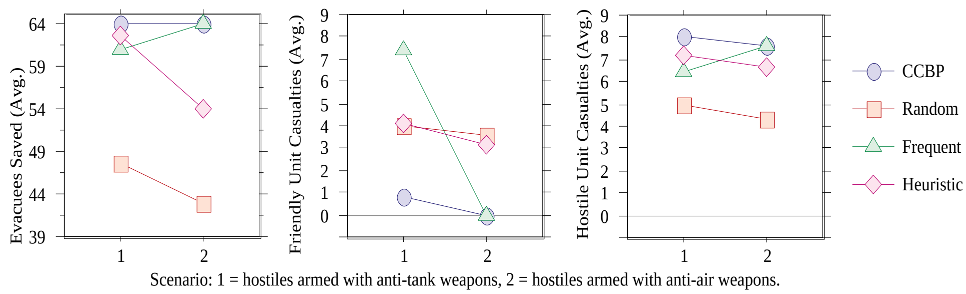

Fig. 7. Comparison of plan selection strategies using Navy MOEs for NEOs.

selection strategies. Figure 7 compares the effectiveness of these four strategies. Overall, HICAP selected plans of higher quality than the other strategies because its plan selection decisions are tailored to the characteristics of each scenario.

8 Related Research

Case-based planning (CBP) has been extensively researched (Bergmann et al., 1998). Our research is closely related to studies on hierarchical CBP (e.g., Kambhampati, 1993; Bergmann & Wilke, 1995; Branting & Aha, 1995). HICAP differs from these other approaches in that it includes the user in its problem solving loop. This is particularly important for applications like NEO planning, where completely automated tools are unacceptable. MI-CBP (Veloso et al., 1997) uses rationale-directed CBP to suggest plan modifications in a mixed-initiative setting, but does not perform doctrine-driven task decomposition.

Some researchers have used CBP with HTNs for military tasks. For example, Mitchell (1997) used integrated CBP to select tasks for a tactical response planner. NEO planning requires that each task be addressed - no choice is involved - and we use CBP to instead choose *how* to perform a task. HICAP's interactions instead focus on retrieval rather than plan adaptation and learning.

9 Conclusion and Future Work

The HICAP case-based planner helps users to formulate a course of action for hierarchical tasks. It is the first tool to combine a task guideline decomposition process with CBR to support interactive plan formulation. It yields plans that benefit from previous experiences and conform to predefined guidelines. HICAP also supports experience sharing, thus allowing planners to exploit knowledge from other planning experts. These design characteristics enhance HICAP's acceptance by military planning personnel.

We are currently integrating HICAP with a generative HTN planner that can evaluate numeric expressions (Nau et. al., 1999), which is particularly important for NEOs because decisions often depend on resource capability and availability (i.e., determining whether a helicopter requires in-flight refueling for a given mission). HICAP will serve as the plan formulation component for the Space and Naval Warfare Systems Command's Interactive Decision Support (IDS) system. When completed, IDS will perform distributed NEO plan formulation, execution, monitoring, and replanning.

Our collaborative research with IDS partners will focus on associating temporal durations with tasks, developing a resource tracking module (i.e., to solve resource conflicts), implementing a strategy for justifying case rankings, integrating HICAP with a powerful dynamic planner (i.e., SIPE-2 (Wilkins, 1998)), and integrating existing GUIs for plan authoring. We will also investigate methods for performing information gathering in HICAP using a planning approach (e.g., Carrick et al., 1999).

Acknowledgements

Thanks to ONR Program Managers Michael Shneier and Paul Quinn, and Program Officer Lt. Cdr. Dave Jakubek, for their encouragement throughout this project. This research was supported by grants from the Office of Naval Research, the Naval Research Laboratory, and the Army Research Laboratory. Many thanks to members of the Center for Naval Analyses and ONR's Naval Science Assistance Program for their guidance and support. And thanks to our ICCBR-99 reviewers for their thoughtful suggestions, which improved this paper.

References

Aha, D. W., & Breslow, L. A. (1997). Refining conversational case libraries. *Proceedings of the Second International Conference on CBR* (pp. 267–278). Providence, RI: Springer.

Bergmann, R., Muñoz-Avila, H., Veloso, M., Melis, E. (1998). Case-based reasoning applied to planning tasks. In M. Lenz, B. Bartsch-Spoerl, H.-D. Burkhard, & S. Wess (Eds.) *CBR Technology: From Foundations to Applications.* Berlin: Springer.

Bergmann, R. & Wilke, W. (1995). Building and refining abstract planning cases by change of representation language. *Journal of AI Research*, *3*, 53–118.

Branting, L. K., & Aha, D. W. (1995). Stratified case-based reasoning: Reusing hierarchical problem solving episodes. *Proceedings of the Fourteenth International Joint Conference on AI* (pp. 384–390). Montreal, Canada: Morgan Kaufmann.

Breslow, L., & Aha, D. W. (1997). *NaCoDAE: Navy Conversational Decision Aids Environment* (TR AIC-97-018). Washington, DC: Naval Research Laboratory, Navy Center for Applied Research in Artificial Intelligence.

Carrick, C., Yang, Q., Abi-Zeid, I., & Lamontagne, L. (1999). Activating CBR systems through autonomous information gathering. To appear in *Proceedings of the Third International Conference on Case-Based Reasoning.* Munich, Germany: Springer.

Ceranowicz, A. (1994). *Modular Semi-Automated Forces. Proceedings of the Winter Simulation Conference of the ACM* (pp. 755–761). New York, NY: IEEE.

DoD (1994). *Joint tactics, techniques and procedures for noncombat evacuation operations* (Joint Report 3-07.51, Second Draft). Washington, DC: Department of Defense.

Erol, K., Nau, D., & Hendler, J. (1994). HTN planning: Complexity and expressivity. *Proceedings of the Twelfth National Conference on Artificial Intelligence* (pp. 1123–1128). Seattle, WA: AAAI Press.

Fikes, R.E., & Nilsson, N.J. (1971). Strips: A new approach to the application of theorem proving in problem solving. *Artificial Intelligence*, *2*, 189–208.

Kambhampati, S. (1994). Exploiting causal structure to control retrieval and refitting during plan reuse. *Computational Intelligence*, *10*, 213–244.

Lambert, Kirk S. (1992). *Noncombatant evacuation operations: Plan now or pay later* (Technical Report). Newport, RI: Naval War College.

Mitchell, S.W. (1997). A hybrid architecture for real-time mixed-initiative planning and control. *Proceedings of the Ninth Conference on Innovative Applications of AI* (pp. 1032–1037). Providence, RI: AAAI Press.

Muñoz-Avila, H., Breslow, L.A., Aha, D.W., & Nau, D. (1998). *Description and functionality of HTE* (TR AIC-98-022). Washington, DC: NRL, NCARAI.

Muñoz-Avila, H., Aha, D.W., Breslow, L. & Nau, D. (1999). HICAP: An interactive case-based planning architecture and its application to noncombatant evacuation operations. To appear in *Proceedings of the Ninth National Conference on Innovative Applications of Artificial Intelligence*. Orlando, FL: AAAI Press.

Nau, D. S., Cao, Y., Lotem, A., & Muñoz-Avila, H. (1999). SHOP: Simple Hierarchical Ordered Planner. To appear in *Proceedings of the Sixteenth National Conference on Artificial Intelligence*. Stockholm, Sweden: Morgan Kaufmann.

Racine, K., & Yang, Q. (1997). Maintaining unstructured case bases. *Proceedings of the Second International Conference on CBR* (pp. 553–564). Providence, RI: Springer.

Sachtleben, G.R. (1991). Operation Sharp Edge: The Corps MEU (SOC) program in action. *Marine Corps Gazette*, *11*, 76–86.

Siegel, A.B. (1991). *Eastern Exit: The noncombatant evacuation operation (NEO) from Mogadishu, Somalia, in January 1991* (TR CRM 91-221). Arlington, VA: Center for Naval Analyses.

Siegel, A.B. (1995). *Requirements for humanitarian assistance and peace operations: Insights from seven case studies* (TR CRM 94-74). Arlington, VA: CNA.

Veloso, M., Mulvehill, A.M., & Cox, M.T. (1997). Rationale-supported mixed-initiative case-based planning. *Proceedings of the Ninth Conference on Innovative Applications of Artificial Intelligence* (pp. 1072–1077). Providence, RI: AAAI Press.

Wilkins, D.E. (1998). *Using the SIPE-2 planning system: A manual for Version 5.0* (Working Document). Menlo Park, CA: Stanford Research International, Artificial Intelligence Center.