HICAP: An Interactive Case-Based Planning Architecture and its Application to Noncombatant Evacuation Operations *

Héctor Muñoz-Avila[†‡], David W. Aha[‡], Len Breslow[‡], & Dana Nau[†]
† Dept. of Computer Science, University of Maryland, College Park, MD 20742-3255
‡ Navy Center for Applied Research in AI, Naval Research Laboratory, Washington, DC 20375
† *lastname*@cs.umd.edu ‡ *lastname*@aic.nrl.navy.mil

Abstract

This paper describes *HICAP*, a general purpose planning architecture that we have developed and applied to assist military commanders and their staff with planning *NEOs* (Noncombatant Evacuation Operations). HICAP integrates a hierarchical task editor, HTE, with a conversational case-based planner, NaCoDAE/HTN. In this application, HTE maintains an agenda of tactical planning tasks that, according to military doctrine, must be addressed in a NEO. Military planning personnel select a task to decompose from HTE and then use NaCoDAE/HTN to interactively refine it into an operational plan by selecting and applying cases, which represent task decompositions from previous NEO operations. Thus, HICAP helps commanders by using previous experience to formulate operational plans that are in accordance with NEO doctrine.

Introduction

NEOs (Noncombatant Evacuation Operations) are military operations that require performing hundreds of subtasks and whose primary goal is to minimize loss of life. Formulating NEO plans is a complex task because it involves considering many factors (e.g., military resources, political issues, meteorological predictions) and uncertainties (e.g., hostility levels and locations).

Flawed NEO plans could yield dire consequences. For example, doctrine states that evacuees must be inspected prior to military transport. However, Siegel (1991) reported that this task was never assigned to any ground team in Operation Eastern Exit, and one evacuee produced his 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.

Our thesis is that commanders and their staff can greatly benefit from the assistance of an intelligent NEO plan formulation tool. After analyzing NEO doctrine, reviewing case study analyses, and consulting with NEO planning experts, we concluded that commanders will not accept any assistant tool unless it exhibits the following characteristics:

- *Doctrine-driven*: Uses a doctrine task analysis to guide plan formulation.
- *Interactive*: Users control how the abstract doctrine is instantiated for a specific NEO.
- *Provide Case Access*: Indexes plan segments from previous NEOs, and retrieve them for users if warranted by the current operational environment.

Also, this tool must perform several bookkeeping tasks. Although several systems have been proposed for NEO planning, none have been deployed because they do not exhibit all of these characteristics.

This paper introduces HICAP (Hierarchical Interactive Case-based Architecture for Planning), a general-purpose plan formulation tool designed to meet these characteristics, and its application to assist NEO planners.[1] HICAP integrates a task decomposition editor, HTE, with a conversational case-based planner, NaCoDAE/HTN. The former allows users to edit doctrine tasks and select tasks to operationalize, while the latter allows users to interactively refine HTN plans. Their integration in HICAP ensures that operational plans are framed within planning doctrine or within the changes made by commanders through three mechanisms. First, it uses doctrine to guide plan formulation, as it is done in practice. Second, it supports interactive task editing. Finally, it incorporates knowledge of previous operations, represented as cases, that can be used to augment or replace doctrine-standard operational procedures.

Noncombatant Evacuation Operations

NEOs are conducted to assist the USA Department of State in evacuating endangered noncombatants (e.g., nonessential military personnel) from locations in a host foreign nation to an appropriate safe haven. They usually involve a swift insertion of a force, temporary occupation of an objective (e.g., a USA Embassy), and a planned withdrawal. NEOs are usually planned and operated by a Joint Task Force (JTF) and conducted

*IAAI-99, NCARAI TR AIC-99-002

[1] A Java 2 Applet version of HICAP can be found at www.aic.nrl.navy.mil/~munoz/hicap.

under an Ambassador's authority. Force sizes can range into the hundreds with all branches of armed services involved, while the evacuees can number into the thousands. At least ten NEOs were conducted within the past decade. Unclassified publications exist that describe NEO doctrine (e.g., DoD, 1994), case studies (e.g., Stewart et al., 1994), and more general analyses (Lambert, 1992).[2]

The NEO decision making process is made at three increasingly-specific levels. First, the *strategic* level involves global and political considerations (e.g., whether to perform the NEO). Then the *tactical* level involves considerations such as determining the size and composition of the NEO force. Finally, the *operational* level is the concrete level, which assigns specific resources to specific tasks.

JTF commanders plan NEOs in the context of doctrine (DoD, 1994), which establishes a framework for designing strategic and tactical plans, but only partly addresses operational considerations. Doctrine describes general aspects that must be considered when planning a military operation (e.g., chain of command, task agenda). However, doctrine is idealized and cannot account for characteristics of specific NEOs. Thus, the JTF commander must always adapt doctrine, and does so in two ways. First, he must eliminate irrelevant planning tasks and add others, depending on the operation's needs, resource availabilities, and his experience. For example, although NEO doctrine states that a forward command element must be inserted into the evacuation area prior to the primary evacuation elements, temporal constraints sometimes prevent this insertion (e.g., Siegel, 1991). Second, he must employ experiences from previous NEOs, which complement doctrine by suggesting operational refinements suitable for the current environment. For example, previous relevant experience might dictate whether it is appropriate to concentrate the evacuees in the embassy or to plan for multiple evacuation sites. In summary, commanders and their staff use guidance from both doctrine and their operational experiences to plan NEOs.

We will describe how HICAP can assist NEO planners by interactively refining operational plans, which is the focus of the JTF's efforts. We ignore strategic issues because they involve political aspects that are challenging to model and simulate.

Knowledge Representation

We use a variant of hierarchical task networks (HTNs) (Erol et al., 1994) to represent plans in HICAP because they are expressive. We define a HTN as 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.

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

Problem solving with HTNs occurs by applying *methods* to decompose 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.

Methods in HICAP are either *non-decomposable*, *uniquely decomposable*, or *decomposable by multiple methods*. Non-decomposable tasks are concrete actions; they can occur only at leaves of the network. Uniquely decomposable tasks correspond to those specified by doctrine, and are solved by unconditional methods (i.e., $k = 0$). In contrast, tasks that can be decomposed in multiple ways correspond to those that must be solved in a specific problem-solving context.

Methods for problem-specific tasks are represented as *cases*, which encode preconditions as a set of question-answer pairs. Cases are obtained from either operational manuals (i.e., standard operational procedures (SOP)) or reports detailing previous problem-solving episodes. When solving a task T, HICAP accesses all cases that can decompose T. If all the preconditions of a SOP case are met, then it should be used to decompose T. Otherwise, a case corresponding to the most similar problem-solving episode should instead be used. For example, one standard procedure is to concentrate all evacuees in the embassy prior to troop deployment. This is not always possible; escorted transports were organized *after* JTF deployment in Eastern Exit (Siegel, 1991) while Sharp Edge required *multiple* separate evacuations (Stewart et al., 1994).

HICAP: Interactive Case-Based Planner

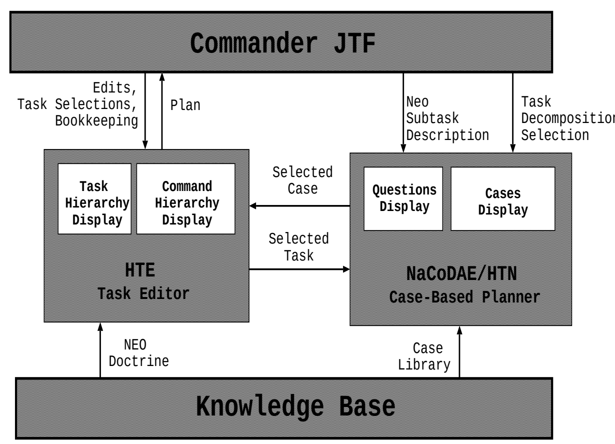

Figure 1: The HICAP architecture.

HICAP (Figure 1) integrates HTE with NaCoDAE/HTN, which are described below.

Hierarchical Task Editor

It is difficult for JTF planners to keep track of the status of NEO subtasks and JTF elements. The *Hierarchical Task Editor* (HTE) was conceived to facilitate the NEO

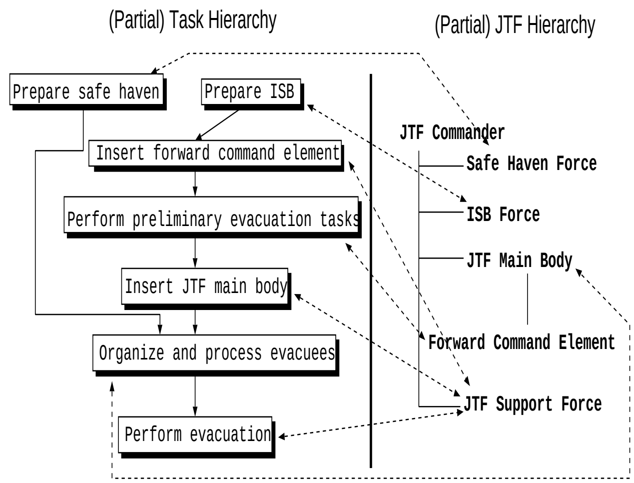

Figure 2: Top level NEO tasks and their assignment to JTF command elements (double arrows denote assignments; arrows denote task orderings; ISB = intermediate stage base).

planning process. Given domain-specific knowledge for operational planning, HTE can be used to:

1. browse and edit this knowledge base's components,
2. select tasks for further decomposition,
3. edit assignments of military personnel to tasks, and
4. investigate the status of tasks.

HTE's knowledge base consists of a HTN, a command hierarchy, and an assignment of tasks to commands. For this NEO application, we encoded a HTN to capture critical planning knowledge (i.e., more than 200 tasks and their ordering relations) corresponding to NEO doctrine (DoD, 1994). Next, we elicited the JTF command hierarchy that is commonly used in NEO operations. Finally, we elicited many-to-one relations between tasks and the JTF elements responsible for them.

In addition to providing users with a visual aid for NEO doctrine, HTE can be used to edit the HTN, its ordering relations, the command hierarchy, and the task-command assignments. Thus, military planners can tailor HTE's knowledge base for a particular NEO. Figure 2 displays (left) the top level tasks that, according to doctrine, must be performed during a NEO and (right) the JTF elements responsible for them.

Conversational Task Decomposer

Doctrine describes decision-making procedures for the operational level but does not provide sufficient detail to formulate tactical plans. NaCoDAE/HTN, an extension of the NaCoDAE case retrieval tool (Breslow & Aha, 1997), helps planners to refine selected operational tasks into subtasks. Given a task T to refine, NaCoDAE/HTN conducts an interactive *conversation* that ends when the user applies a method to T.

NaCoDAE/HTN displays the labels of cases that can decompose the selected node and the questions from these cases whose answers are not yet known for the current NEO. The user can select and answer any displayed question; $\langle q, a\rangle$ pairs are used to compute the current task's similarity with its potential decomposition methods. Cases are ranked by similarity, while questions are ranked by their frequency among the displayed cases. Answering a question modifies the case rankings and the displays. A conversation ends when the user selects a case to decompose the current task.

Some cases are SOPs; they can only be selected to decompose a task after all of their questions have been answered and match the current planning scenario (i.e., their preconditions must all match). In contrast, cases based on previous experiences can be selected even when 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.

Integration

HICAP integrates HTE with NaCoDAE/HTN to formulate plans that are in accordance with both doctrine and cases. It inputs a HTN describing the doctrine for an application, along with a set of cases for each of the subtasks that can be decomposed in multiple ways. Under user control, it outputs an edited HTN whose leaves are concrete actions as specified by case applications.

Using HTE, military planners can edit the NEO knowledge base by deleting (or replacing) any task subtree, by editing the command hierarchy, and by reassigning tasks and/or command elements. HTE also allows users to select a (leaf) task T to be decomposed. This initializes a NaCoDAE/HTN conversation, using T as an index for case retrieval. A conversation yields a selected case $C = \langle l, T, N, P\rangle$, whose network N is used to decompose T. This expansion can be recursive. Eventually, non-decomposable tasks corresponding to concrete actions will be reached. Task expansions are immediately displayed by HTE.

NEO planners formulate both a main course of action (i.e., a plan) and contingency plans, such as when an element cannot accomplish a key task. Users can generate contingency plans by performing alternative task decompositions and edits to the knowledge base. Handling unforeseen contingencies that may occur during plan execution are beyond the scope of this paper.

In summary, HICAP satisfies the requirements stated in the Introduction. First, all plans obtained using HICAP are clearly circumscribed by the doctrine or by any modification introduced by the JTF commander. Second, HICAP interactively supports task editing and triggers conversations for decomposable tasks. Third, it incorporates knowledge from previous operations as cases, which serve as task decomposition alternatives. Finally, it allows the user to visually check that all tasks are assigned to JTF elements.

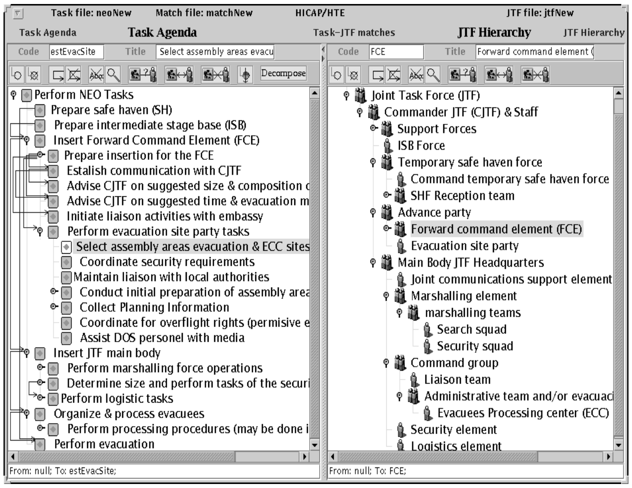

Figure 3: HICAP/HTE's snapshot displaying the doctrine tasks (left) and the JTF's hierarchical organization (right). Arrows denote ordering constraints.

Example: NEO Planning

NEO planners view the top level tasks first, revising them or their assignments if necessary. They may decompose any task and view its decomposition. Figure 3 shows an intermediate stage during this process. The user has selected the task *Select assembly areas for evacuation & ECC (Evacuation Control Center) sites*, which is highlighted together with the command element responsible for it (i.e., the FCE).

Standard procedure dictates that the embassy is the ideal assembly area. However, it is not always possible to concentrate the evacuees in the embassy. 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 (left), on the following page).

If a user answers *Are there any hostiles between the embassy and the evacuees?* with "uncertain," the second displayed case (Figure 4 (right)) will become a perfect match. Figure 5 shows the decomposition when the user selects this case for decomposition; two new subtasks are displayed that correspond to its decomposition network. The *Send UAV (Unmaned Air Vehicle) to* ... task is non-decomposable; it corresponds to a concrete action. If the user tells HICAP to decompose *Determine hostile presence*, HICAP will initiate a new NaCoDAE/HTN dialogue (also in Figure 5).

Figure 6 shows a snapshot of HICAP when the user selects *The UAV detects hostiles* alternative and decides to decompose the *Handle hostile presence* subtask, which prompts a new NaCoDAE/HTN dialogue. Assuming that the user answers *Can the hostile forces* ... with "yes," this matches the situation in Operation Eastern Exit in which the evacuees were dispersed into multiple locations in Mogadishu. An escorted transport gathered all evacuees into the embassy. If the user

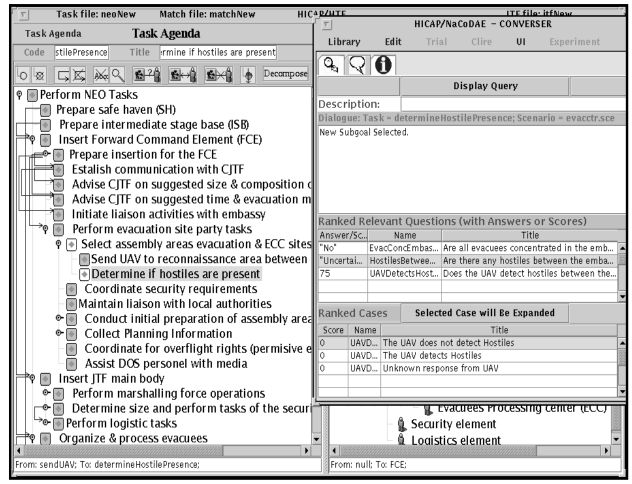

Figure 5: HICAP's interface after selecting the *Determine hostile presence* task.

selects this case, then its two non-decomposable subtasks, *Assign dissuasive escort* and *Escort evacuees to embassy*, will be displayed.

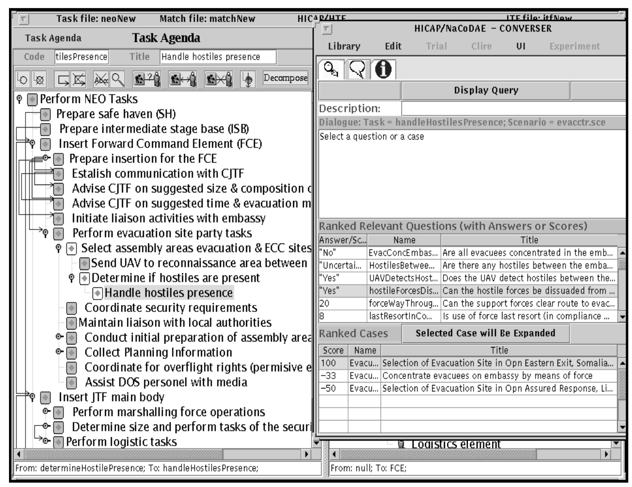

Figure 6: Advising the commander on how to handle a hostile presence.

Related Research

At this time, there are no intelligent deployed NEO planning tools. Kostek (1988) proposed a conceptual design for predicting the force size and type required for a NEO. Chavez and Henrion (1994) described a decision-theoretic approach for instantiating a general NEO plan with specific parameters for locations, forces, and destinations, and used it to assess alternative plans. Gil et al. (1994) presented a system for predicting manning estimates for certain NEO tasks. None of these systems formulate NEO plans, although desJardins et

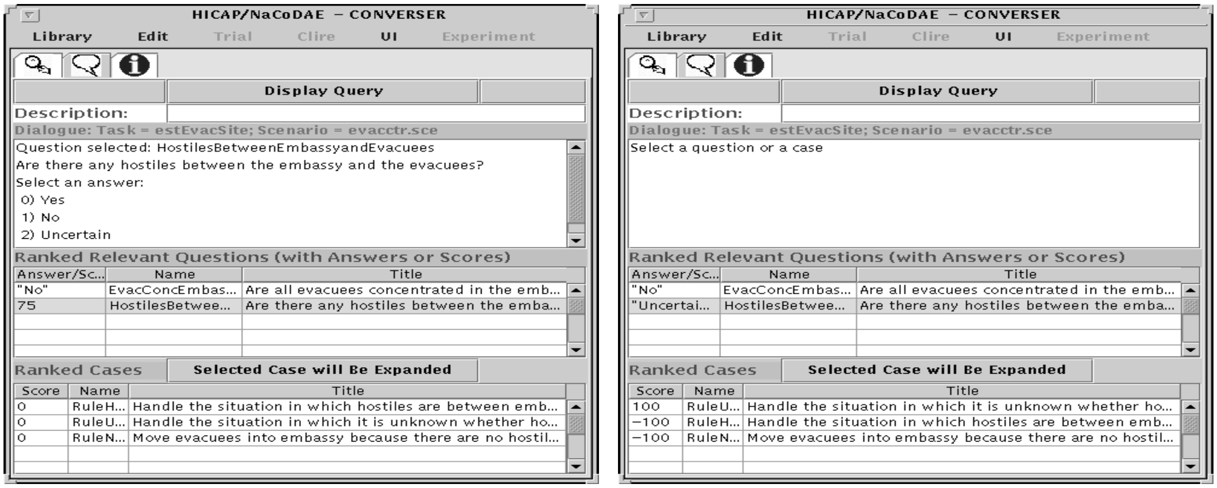

Figure 4: NaCoDAE/HTN's interface, before (left) and after (right) answering a question. The top window directs the user and, when the user answers a question, lists the possible answers. The lower windows display the questions and cases, respectively.

al. (1998) recently proposed a distributed hierarchical planning approach.

Although DARPA and other agencies have sponsored several projects related to NEO planning (e.g., ARPI, Tate (1994)), HICAP is the *first* system to use a *conversational* case-based approach for plan formulation. HICAP's advantage is that it allows users to incrementally elaborate a planning scenario, provides a focus of attention that guides this elaboration, and provides access to stored plan fragments for use in new NEO plans. Our approach was inspired by research on and applications of conversational case retrieval systems (Aha & Breslow, 1997), but extends them to apply to hierarchical planning tasks.

Some researchers have used case-based approaches for HTN planning tasks on military domains. For example, Mitchell (1997) used an integrated CBP (case-based planner) 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. MI-CBP (Veloso et al., 1997) uses rationale-directed CBP to suggest plan modifications, but does not perform doctrine-driven task decomposition. HICAP's interactions instead focus on retrieval rather than plan adaptation and learning. IFD4 (Bienkowski & Hoebel, 1998) automatically generates plans as guided by an editable objectives hierarchy. In contrast, HICAP's objectives are fixed, and user interaction focuses on task formulation.

Other researchers have described related crisis response systems. Ferguson and Allen (1998) described an interactive planner for military crisis response, but their system does not use cases during plan formulation and does not perform doctrine-driven task decomposition. Likewise, Wolverton and desJardins' (1998) distributed generative planner also does not use cases. Gervasio et al. (1998) described an interactive hierarchical case-based scheduler for crisis response that does not perform interactive plan formulation. Avesani et al. (1998) described a CBP for fighting forest fires that supports interactive plan adaptation, but does not use hierarchical guidelines to formulate plans (ala HICAP). Finally, Leake et al. (1996) described a CBP applied to disaster response that focuses on learning case adaptation knowledge, but it is not doctrine-driven and focuses interaction on knowledge acquisition rather than problem elicitation.

Conclusions and Transitions

HICAP is a case-based tool that assists the military commander with formulating a plan. It is the first tool to combine a doctrine-guided task decomposition process with a case-based reasoning approach to support interactive plan formulation. Thus, it yields a plan that benefits from previous experiences and is sound according to doctrine. Furthermore, HICAP supports experience sharing, which allows planners to exploit knowledge from other planning experts. These design characteristics were chosen so as to enhance HICAP's acceptance by military planning personnel; previous approaches did not include all of these capabilities.

There is a great potential for combining HTN and case-based planning techniques in real-world applications. HICAP illustrates this by using a unified framework for a rather complex domain, noncombatant evacuation operations, that has multiple sources of information (i.e., doctrine, previous operations and standard operating procedures).

HICAP will serve as the plan formulation component for the Interactive Decision Support (IDS) system being developed at the Space and Naval Warfare Systems Command. When completed, IDS will perform distributed plan formulation, execution, monitoring, and replanning for NEO planning efforts.

We are currently integrating HICAP with SHOP (Nau et al., 1999), a simple generative HTN planner that can process numeric computations, which is particularly important for NEO planning in decisions made concerning resource capability and availability (i.e., determining whether a helicopter requires in-flight refueling for a given mission).

In an initial evaluation of HICAP using ModSAF simulations, HICAP outperformed three default planning strategies for a single planning subtask under two planning scenarios (Muñoz-Avila et al., 1999). After incorporation in IDS, HICAP will be evaluated in controlled studies by military mission planners.

Acknowledgements

Thanks to Pat Langley and Frank Weberskirch for their comments on an earlier draft. This research was supported by grants from the Naval Research Laboratory (N00173-98-1-G007), the Army Research Laboratory (DAAL01-97-K0135), and the Office of Naval Research.

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.

Aha, D.W., & Daniels, J.J. (Eds.) *CBR Integrations: Papers from the 1998 Workshop* (TR WS-98-15). Menlo Park, CA: AAAI Press.

Avesani, P., Perini, A., & Ricci, F. (1998). The twofold integration of CBR in decision support systems. In (Aha & Daniels, 1998).

Bienkowski, M.A., & Hoebel, L.J. (1998). Integrating AI components for a military planning application. *Proceedings of the Fifthteenth National Conference on AI* (pp. 561–566). Madison, WI: AAAI Press.

Breslow, L., & Aha, D. W. (1997). *NaCoDAE: Navy Conversational Decision Aids Environment* (TR AIC-97-018). Washington, DC: NRL, NCARAI.

Chavez, T., & Henrion, M. (1994). Focusing on what matters in plan evaluation: Efficiently estimating the value of information. *Proceedings of the ARPA/Rome Laboratory Knowledge-Based Planing and Scheduling Initiative* (pp. 387–399). Tuscon, AR: Morgan Kaufmann.

desJardins, M., Francis, A., & Wolverton, M. (1998). Hybrid planning: An approach to integrating generative and case-based planning. In (Aha & Daniels, 1998).

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

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

Ferguson, G., & Allen, J.F. (1998). TRIPS: An integrated intelligent problem-solving assistant. *Proceedings of the Fifthteenth National Conference on AI* (pp. 567–572). Madison, WI: AAAI Press.

Gervasio, M.T., Iba, W., & Langley, P. (1998). Case-based seeding for an interactive crisis response assistant. In (Aha & Daniels, 1998).

Gil, Y., Hoffman, M. & Tate, A. (1994). Domain-specific criteria to direct and evaluate planning systems. *Proceedings of the ARPA/Rome Laboratory Knowledge-Based Planing and Scheduling Initiative* (pp. 433–444). Tuscon, AR: Morgan Kaufmann.

Kostek, S.R. (1988). *A User's Design of a Decision Support System for Noncombatant Evacuation Operations for United States Central Command.* Master's thesis, School of Engineering, Air Force Institute of Technology, Dayton, Ohio.

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

Leake, D. B., Kinley, A., & Wilson, D. (1996). Acquiring case adaptation knowledge: A hybrid approach. *Proceedings of the Thirteenth National Conference on AI* (pp. 684–689). Portland, OR: AAAI Press.

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., McFarlane, D., Aha, D.W., Ballas, J., Breslow, L.A., & Nau, D. (1999). Using guidelines to constrain interactive case-based HTN planning. To appear in *Proceedings of the Third International Conference on CBR.* Munich: Springer.

Nau, D., Cao, Y., Lotem, A., & Muñoz-Avila, H. (1999). To appear in *Proceedings of the Sixteenth International Joint Conference on AI.* Stockholm: AAAI Press.

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.

Stewart, G., Fabbri, S.M., & Siegel, A.B. (1994). *JTF operations since 1983* (TR CRM 94-42). Arlington, VA: CNA.

Tate, A. (1994). Mixed initiative planning in O-Plan2. *Proceedings of the ARPA/Rome Laboratory Knowledge-Based Planing and Scheduling Initiative* (pp. 512–516). Tuscon, AR: Morgan Kaufmann.

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 AI* (pp. 1072–1077). Providence, RI: AAAI Press.

Wolverton, M., & desJardins, M. (1998). Controlling communication in distributed planning using irrelevance reasoning. *Proceedings of the Fifthteenth National Conference on AI* (pp. 868–874). Madison, WI: AAAI Press.