

A Call for Knowledge-Based Planning

David E. Wilkins and Marie desJardins

■ We are interested in solving real-world planning problems and, to that end, argue for the use of domain knowledge in planning. We believe that the field must develop methods capable of using rich knowledge models to make planning tools useful for complex problems. We discuss the suitability of current planning paradigms for solving these problems. In particular, we compare knowledge-rich approaches such as hierarchical task network planning to minimal-knowledge methods such as STRIPS-based planners and disjunctive planners. We argue that the former methods have advantages such as scalability, expressiveness, continuous plan modification during execution, and the ability to interact with humans. However, these planners also have limitations, such as requiring complete domain models and failing to model uncertainty, that often make them inadequate for real-world problems.

In this article, we define the terms *knowledge-based* and *primitive-action planning* and argue for the use of knowledge-based planning as a paradigm for solving real-world problems. We next summarize some of the characteristics of real-world problems that we are interested in addressing. Several current real-world planning applications are described, focusing on the ways in which knowledge is brought to bear on the planning problem. We describe some existing knowledge-based approaches and then discuss additional capabilities, beyond those available in existing systems, that are needed. Finally, we draw an analogy from the current focus of the planning community on disjunctive planners to the experiences of the machine learning community over the past decade.

We are interested in solving real-world planning problems and believe that doing so will require techniques that are more expressive and provide a wider range of capabilities than current planning systems. Real-world problems have been found to

require more expressive representations and capabilities than are needed for the standard set of benchmark planning problems (blocks world, towers of Hanoi, simplified logistics, and the like) or for the problems used in the 1998 and 2000 Artificial Intelligence Planning and Scheduling (AIPS) Conference planning competitions (Bacchus et al. 2000; Long 2000; McDermott 2000).

Past research in AI planning can roughly be divided into two camps: (1) systems that take a minimalist approach to domain knowledge and (2) systems that focus on leveraging as much domain knowledge as possible. Techniques in the first set generally restrict themselves to domain models consisting of STRIPS-style descriptions of primitive actions. We refer to these methods in the first set as *primitive-action planning techniques* because they construct plans from descriptions of the actions that can appear in the final plan.¹ A currently popular form of primitive-action planning is *disjunctive planning*, which uses primitive action descriptions to encode and solve propositional representations of planning problems. (A recent survey of current directions in AI planning [Weld 1999] focuses almost exclusively on this style of planning.)

Techniques in the second set are based on a philosophy of using whatever domain knowledge is available to solve the planning problem. These systems are characterized by the use of multiple types of domain knowledge and complex domain models to support their reasoning processes. This knowledge can include task and goal structures, additional constraints, search-control techniques, and interaction with humans when necessary to make use of their expertise. We refer to these techniques as *knowledge-based planning methods*.

*... to scale up
to complex
problems,
multiple types
of knowledge
must be
explicitly
encoded in
under-
standable
structures ...*

In this article, we argue that to scale up to complex problems, multiple types of knowledge must be explicitly encoded in understandable structures, and planning algorithms must be able to use this explicit knowledge effectively.

One possible approach to achieving this knowledge encoding would be to augment primitive-action planners with additional knowledge. Indeed, researchers have begun to investigate how these systems can be augmented with domain-specific, hand-encoded control rules (see sidebar). However, only certain types of knowledge can be captured in these rules, and they are difficult and time consuming to construct. The result is that this approach to encoding knowledge is not scalable to large, complex problems. It seems unlikely that hand-coded control rules will be a sufficient approach to knowledge modeling for the types of problem that arise in the real world.

Knowledge-based planning methods have limitations as well. On the one hand, current knowledge-based methods are not knowledge based enough: As we discuss later in the article, they do not incorporate many types of knowledge that are important for real-world problems. On the other hand, they are too knowledge based for some problems: The inference entailed by modeling many aspects of the planning problem can be expensive computationally, with the result that knowledge-based methods might not be the most efficient approach to solving certain types of constraint-satisfaction subproblem within a larger planning problem. Therefore, we believe that solving large, complex problems will require both the development of new knowledge-based methods and the integration of primitive-action methods with these new techniques.

Knowledge-based planning and primitive-action planning represent two ends of a continuum: A planner might be more or less knowledge based depending on what the range of knowledge is that it uses and how effectively it uses it. Integrated systems might ultimately provide the best of both worlds. However, we argue that knowledge-based methods can solve problems that primitive-action methods cannot because of the greater expressivity and more natural representations of knowledge-based planning.

The remainder of this article is organized as follows: We first describe some common characteristics of real-world planning problems that are not solvable by current primitive-action methods, and we argue that these methods are unlikely to extend to these problems.

We describe some real-world planning problems that have been addressed by current knowledge-based planners. We then discuss how multiple types of knowledge and capability are exploited in existing knowledge-based planners. However, current knowledge-based techniques represent only a small step in the direction of the level of knowledge-based planning that we envision. We next argue that achieving goal-directed behavior in a complex, dynamic world will require reasoning about the consequences of future actions, which, in turn, will entail the use of much more knowledge and richer knowledge models than those used in today's knowledge-based planners. We discuss forms of knowledge that current knowledge-based planners do not use and give some examples of problems that today's planners are not able to solve. Finally, we draw some lessons from the history of the machine learning research community that are analogous to the current trends in the planning community.

Characteristics of Real-World Planning Problems

Real-world problems have been found by many researchers to require more expressive representations and capabilities than those provided by current AI planning systems. Chien et al. (1996) conclude from their experience with multiple National Aeronautics and Space Administration (NASA) applications that current plan representations are impoverished. They discuss the requirements of an operational context in which users must interact with the system and must be able to understand and modify the plans produced by the planner. Our experience with military and oil-spill-planning applications supports these conclusions. Here, we describe some of the specific capabilities that are needed to solve real-world problems: numeric reasoning, concurrent actions, context-dependent effects, interaction with users, execution monitoring, replanning, and scalability.

Reasoning with numbers is essential in every realistic domain that we have studied. Common needs for numbers are time, sharable resources having a specific capacity, continuous resources available in limited quantities, and goals of accumulation. An example of goals of accumulation is obtaining a certain quantity of resource that must be assembled from smaller aggregations, such as getting enough boom from several warehouses to contain an oil spill or enough soldiers or equipment for a military operation. In practice, disjunctive planners have difficulty handling

problems involving reasoning about numbers.² In most existing non-hierarchical task network (HTN) AI planners, the need for numeric reasoning is reduced by assuming that sharable resources have infinite capacity and that continuous resources are unlimited (Srivastava and Kambhampati 1999).

Realistic domains can have dozens of (perhaps necessarily) parallel activities as activities of various agents are coordinated. Parallelism can cause computational problems for disjunctive planners, and some systems produce only sequential plans.

Realistic domains often have numerous context-dependent effects, which can cause an exponential explosion in the number of STRIPS operators needed. This problem is being addressed to some extent in disjunctive planners. Extensions to the GRAPHPLAN algorithm to handle conditional effects are given in Kambhampati, Parker, and Lambrecht (1997) and Guere and Alami (1999), but applications of GRAPHPLAN are already limited by computational efficiency, and neither paper discusses the time or space complexity of the algorithms. Other approaches have also been tried, perhaps the most promising being *factored expansion*, in which an action with conditional effects is split into new actions called *components*, one for each conditional effect. This approach appears to outperform a straightforward splitting of each action into (a possibly exponential number of) STRIPS operators, at least on large problems. The cost is added complexity in the planning algorithms involving tricky extensions (Anderson, Smith, and Weld 1998).

Interacting with people is a critical aspect of real-world planning. Realistic problems are embedded in the world and generally do not have precisely defined boundaries or evaluation functions. Thus, most interesting planning problems will be difficult or impossible to model fully. For example, criteria for plan evaluation often cannot be quantified, such as when the political consequences of a military or media action are crucial. It is also hard to specify when a situation warrants breaking rules or ignoring certain information, yet such situations are common in real life. In such cases, a human user must be able to guide the planner and evaluate the plans produced, allowing the planning system to take advantage of the user's expertise.

In the real world, the goal of planning is not simply to build the plan but to use it to control actions in the world. Therefore, realistic planning systems must support execution monitoring and continuous plan modification during execution. Figure 1 shows a possible architec-

ture for a system incorporating both planning and execution. The intention is not to promote this architecture but to show some of the added complication introduced by executing a plan. The planner and executor have specialized knowledge bases to support their respective roles and must also share common ontologies, models of actions, and knowledge about the world. The executor requires a rich model of the relationships among tasks in the plan and of possible outcomes and contingencies. The executor generally operates at a faster tempo than the planner and attempts to provide effective responses to changing conditions, which requires rapid replanning and possibly responding to time-critical situations without invoking the planner (perhaps using procedures from the monitoring knowledge).

The *plan initializer-synchronizer* in figure 1 prepares plans for execution. This module can generate plan-specific monitors to efficiently monitor conditions over certain intervals. It must also synchronize the transfer of control from the original plan to a newly modified version of the plan, even though the original plan has continued execution during the replanning process. In many domains, large quantities of information about the world state (for example, from sensor networks) are constantly arriving, but only small portions might affect the current plan. The information manager filters the incoming data, passing the relevant parts to the execution manager.

Because there is no dependency structure in the disjunctive planner's plans, monitoring them is difficult. In addition, the disjunctive planning approach is brittle in the face of changing problem requirements, and any change in the environment can result in the planning system having to start from scratch. Because knowledge-based planners record dependencies and goal structure, knowledge-based replanning techniques can often modify an executing plan in the face of new requirements (Myers 1999; Wilkins et al. 1995).

Finally, realistic problems involve enormous search spaces, so scalability is essential. Vast strides have been made in the size of problems solved by disjunctive planners, such as solving instances with 10^{16} to 10^{19} configurations. However, these large disjunctive planning problems are still representations of toy problems, such as a logistics problem with 9 packages, 5 trucks, 2 airplanes, and 15 locations (Kautz and Selman 1998). Simply increasing the number of locations to a realistic number will make even these toy problems unsolvable. In contrast, HTN planners can generate plans in domains with thousands of objects and hun-

Interacting with people is a critical aspect of real-world planning.

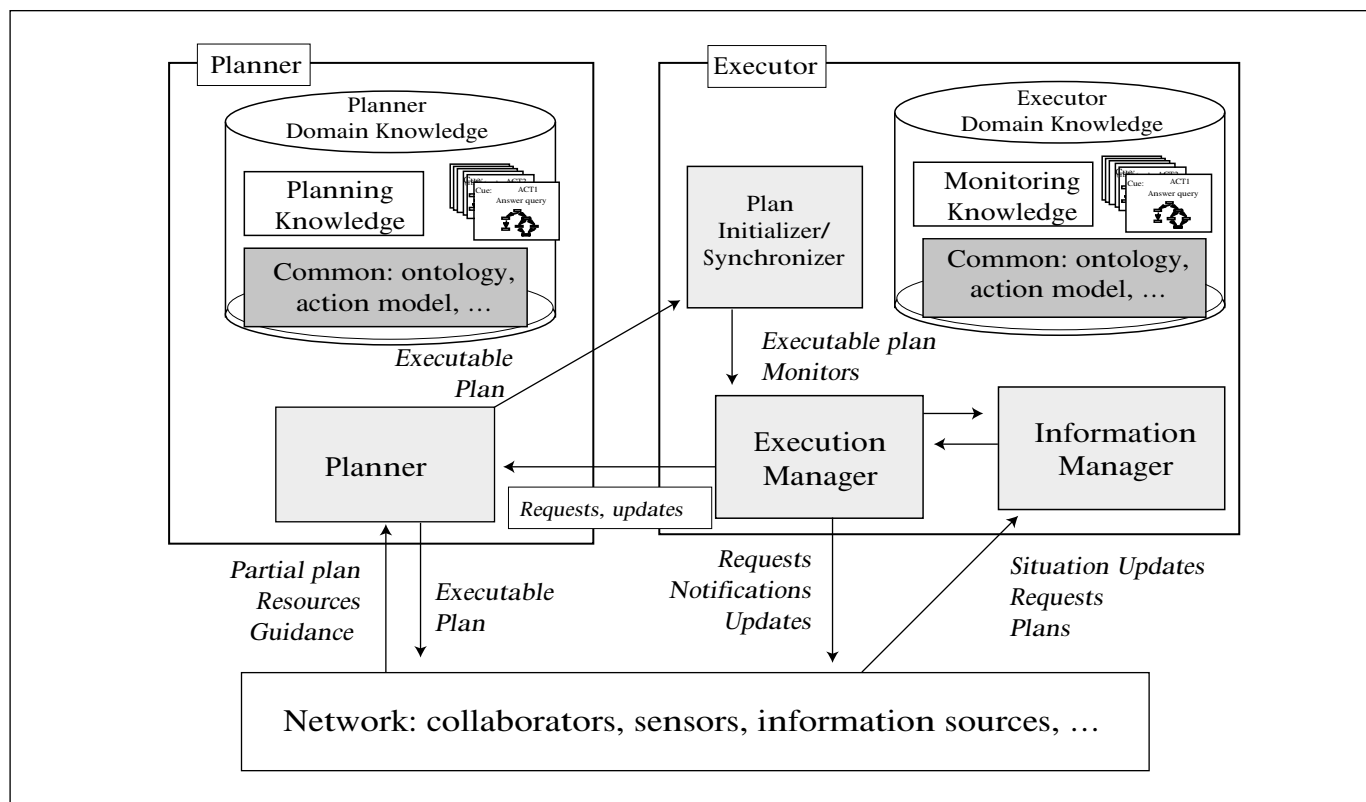


Figure 1. Possible Architecture for a System That Incorporates Both Planning and Execution.

The planner sends plans to the executor, which monitors the executing plan and sends requests and situation updates to the planner.

dreds or thousands of actions (Wilkins and Myers 1998; Wilkins and Desimone 1994).

Real-World Applications

We describe several existing planning applications that include many of the characteristics of hard real-world problems that we are interested in exploring. These applications use a variety of knowledge-based planning techniques and represent a starting point for research into real-world planning methods. They also highlight the need for further work in many of the areas we have discussed. Another survey of real-world planning applications can be found in Knoblock (1996).

Unlike toy problems, real-world problems generally cannot be modeled completely, particularly when plans are executed in the real world. Therefore, validating the correctness of a planner raises many challenging issues. The NASA domains described in this section have been fielded, and the planners have been subjected to extensive validation procedures. Generally, validation involves empirical tests for carefully selected test cases (against simulators, test harnesses, or the real world). Results can be checked automatically against a set of correct-

ness requirements that have carefully been defined and validated by human experts. A description of such testing and the issues involved can be found in Smith, Feather, and Muscettola (2000).

Oil-Spill Crisis Response Planning

Oil-spill incident response is a race against time to contain or remove oil before it damages the shore. Planning begins by entering the specifics of a spill incident—location, time of day, spill rate, and so on—and then forecasting (using legacy systems) the spill trajectory, considering the uncertainty in its spreading caused by wind and waves. This forecast determines which environmentally sensitive shore sectors that the oil will hit and when.

The spill response configuration system (SRCS) helps the United States Coast Guard (USCG) estimate the adequacy of the amounts and locations of cleanup equipment in its coastal oil-spill incident response plans (Agosta and Wilkins 1996). SRCS determines adequacy by building plans that meet a range of spill scenarios and then evaluating the plans. Previous approaches used approximate rules to estimate equipment needs. By automating the planning process, SRCS enables users to plan and evalu-

ate a range of detailed responses to a range of spill scenarios, enabling the USCG to more accurately estimate its needs.

The planner works from the spill-trajectory forecast together with geographic information such as the sectors into which the region is divided and the USCG requirements for protection of these areas. In addition, the planner works with a database of the quantities and capabilities of available equipment and resources, which are often located over a large geographic area with varying time and transportation costs.

SRCS integrates simulation, evaluation, map display, and scheduling tools with SIPE-2. The planner, which uses a knowledge base of oil-spill-response HTN operators, and the scheduler work interactively with the user to generate a plan consisting of equipment deployment and employment actions. The actions in this plan must satisfy constraints determined by the projected oil-dispersal pattern, equipment cleanup capabilities and transport times, and environmental protection requirements. SRCS is intended to be used for configuration planning—advance planning to prepare for likely incidents—rather than for real-time planning as an incident unfolds.

This application requires extensive use of metric (numeric) goals, primarily resource and temporal reasoning. The temporal reasoning involves deadlines and concurrency. One important use of resource reasoning is the accumulation of a certain level of some resource at a certain place and time. An example is a goal to provide several thousand feet of oil-containment boom to protect a sensitive area. This goal must typically be met by transporting several shipments of boom from different locations around the state or country.

Most of the user's interaction with SRCS is mediated by the map interface, implemented in a commercial geographic information system. Thus, the user can immediately see both what the extent of the spill is and where resources are used at various times.

Plans are evaluated on the degree to which they achieve the overall objective of cleaning up the spilled oil. In many spills, much of the oil will escape, no matter how much equipment is available, because of the difficulty of operations and speed of spread because of the weather. Furthermore, for any spill, SRCS can generate many possible plans, and users can partially or completely sacrifice a sector cleanup goal if they believe equipment that would have been assigned to a sector better serves the overall goals by being used elsewhere.



Figure 2. The Space Shuttle Discovery Lands at Kennedy Space Center after Successfully Completing Mission STS-85.

Photo by National Aeronautics and Space Administration (taken by Bionetics). Used with permission.

Because there are many feasible plans that vary widely in their degree of success, SRCS includes an evaluation model for finding good plans. The plan is used as an input to this model, along with the projected flows determined by the trajectory model. The evaluation model accounts for the quantities of oil contained and removed in each sector for each period. From this accounting, it can calculate measures of plan merit, such as the final fraction of oil removed by the plan.

Space Applications

Several planning domain models have been developed by NASA researchers for studying a range of space applications of AI planning systems.³ Three of these domains are described here; they highlight the need for planning systems that can represent and reason about complex activities, resources, and interactions.

DATA-CHASER

DATA-CHASER flew aboard the Space Shuttle *Discovery* on mission STS-85 as a hitchhiker payload with the international extreme ultraviolet hitchhiker bridge (IEH-2) in August 1997 (figure 2). This mission used automated planning and scheduling techniques to decrease mission commanding effort by 80 percent and increase science return (that is, efficiency of instrument utilization) by 40 percent (compared to manual sequence generation) (Chien et al. 1999).

DATA-CHASER consisted of three coaligned instruments that collected data in the far and extreme ultraviolet wavelengths. These instruments obtained images of the sun that correlated solar activity with radiation flux, associating




Figure 3. The CHASER Module of the DATA-CHASER Planning and Scheduling System That Flew aboard the Space Shuttle Discovery in August 1997.

Photo by National Aeronautics and Space Administration. Used with permission.

this flux with individual active regions of the sun. The irradiance data could be sent to the ground system using low-rate (available 90 percent of the time, at 1200 bits a second) or medium-rate (available when scheduled, at 200 kilobits a second) transmission. The payload was capable of receiving commands sent from the ground system when uplink was available. The DATA module contained the science instruments themselves. The CHASER module (figure 3) contained the planning and scheduling system that managed the shuttle resources to accomplish the mission successfully.

Shuttle resources are shared by multiple missions, and their availability is subject to change every 12 hours (the frequency at which NASA changes shuttle flight plans). These resources include access to uplink and downlink channels and time windows when the instruments are allowed to operate. In addition, DATA-CHASER had thermal constraints that limited the duration of payload exposure to the sun and environmental constraints that restricted the state and activities of the payload when shuttle contamination events occurred. Therefore, DATA-CHASER's planner needed to interoperate with the shuttle flight plan to enforce numerous resource constraints.

DATA-CHASER posed a challenge for automated scheduling techniques because of its complex resource- and power-management requirements. The scheduler needed to identify an optimal data-collection schedule while it adhered to the resource constraints. In addition, scientists wanted to be able to perform dynamic scheduling during the mission. For example, the summary data might indicate the presence of a solar flare. If a solar flare occurred, scientists could change their requirements and goals, for example, raising the priority on certain instruments or providing longer integration times. These new goals could require a different schedule of activities.

Citizen Explorer

The *Citizen Explorer* (CX-1) satellite project is a small satellite built and operated by the Colorado Space Grant Consortium at the University of Colorado at Boulder. CX-1 was scheduled to launch as a secondary payload aboard a *Delta-II* launch vehicle in November 2000 but has been postponed. The science mission of CX-1 will focus on obtaining geographic coverage of ozone, aerosols, and ultraviolet radiation measurements using both on-board and ground-based science instruments. CX-1 mission operations will include ground-based automated planning using the ASPEN knowledge-based planning system (Willis, Rabideau, and Wilklow 1999).

CX-1 operations require managing resources such as the spectrophotometer (Speck) science instrument, battery power, solar array power, and the solid-state disk. There are also several data-collection modes that must be scheduled based on the spacecraft's orbital location. A typical daily CX-1 operations scenario includes Speck data collection, engineering health and status data collection, data downlink to operations ground stations and participating schools, command-set uplink, updates to the on-board executive control database, and updates to ephemeris data. Interactions between limited power availability and limited downlink opportunities (because of ground-station placement, orbital constraints, on-board memory limitations, and the transmission power costs) make mission operations a complex optimization problem.

Antarctic Mapping Mission

The Modified Antarctic Mapping Mission (MASS) used *RadarSAT*, a synthetic aperture radar (SAR) satellite operated by the Canadian Space Agency, to gather interferometry information covering the Antarctic continent from September to December 2000. The ASPEN automated planner (Chien et al. 2000) was used to develop and verify the MAMM mission plan, resulting in a reduction from one work year of planning effort in the first Antarctic Mapping Mission to around eight work weeks for MAMM. The mission plan was executed flawlessly aboard *RadarSAT* during the operation.

RadarSAT is in a 100-minute polar orbit around the earth. One mapping cycle, consisting of 356 orbits, takes about 24 days and must end by positioning *RadarSAT* in the same position and trajectory as at the start of the cycle. During each of its three mapping cycles, *RadarSAT* collected images of the entire continent of Antarctica, with significant redundancy around interesting regions, such as dynamic areas around the coastline and fast-moving ice flows. These three sets of images were used to construct interferometry data, which will allow scientists to determine the surface flow of the continent.

The downlink scheduling problem was complex and highly constrained. Five ground stations were used to downlink information from *RadarSAT*. There are two methods of capturing an image: (1) *real time*, where the image is captured and simultaneously downlinked to a receiving station, and (2) *recorded*, where the image is captured on the on-board recorder (OBR). The OBR has a capacity of approximately 916 seconds. Data on the OBR must be downlinked to a receiving station when *RadarSAT* is visible from the station. Each

ground station is only available for downlinking for as many as 15 minutes an orbit. The only station capable of receiving real-time downlinks from MAMM was MacMurdo Ground Station, which is located in Antarctica. Furthermore, the availability of each downlinking station could change; for example, if MacMurdo were shut down by weather, all imaging had to be recorded and downlinked to other stations.

Many additional operational constraints complicated the scheduling problem. For example, real-time imaging and OBR downlink can occur simultaneously, but real-time imaging and recording images to the OBR cannot. The SAR imager can be on for at most 32 minutes an orbit. Each imaging activity has to be at least one minute long, including the eight-second intervals before and after the imaging. To use the OBR, there is a 13-second spin-up time and a 2.5-second spin-down time, although if two images occur less than 30 seconds apart, the OBR continues to record. The OBR cannot play back until it has recorded all 916 seconds, and then it must play the entire tape back with no pause. There are also delays associated with switching from recording to playback mode, connecting to and disconnecting from a ground receiver, and calibrating this connection between transmissions. In the final schedule, there were approximately 819 imaging activities a cycle, of which one-third were recorded, and two-thirds were real time.

The most challenging planning issue for MAMM was to ensure that all the images were captured within the operational constraints and that all the data were downlinked successfully within the downlink constraints. Because the availability of resources could change during the cycle, rapid replanning in the event of such changes was critical.

Military Air-Campaign Planning

In air-campaign planning, a human planner is typically given a set of high-level political and military goals (for example, protect U.S. citizens and forces from hostile attack) and refines the goals that are attainable (wholly or in part) by the employment of air power into more specific goals. This process iterates until each goal is directly attainable by the execution of a mission. A group of identical aircraft acting in concert performs a mission. Each mission consists of a mission type, a time and a place, a type of aircraft, munitions, and the number of sorties required to execute the mission. Thus, a mission might be expressed as four F-15Cs to escort strike package *P* to target *T* on day *D* + 1. Mission planning details such as flight path

and altitude profile are at a lower level of granularity and can be left until later in the planning process. Support missions, which include refueling and reconnaissance, must be planned because they must compete with combat missions for resources such as aircraft and fuel.

There are often multiple ways to refine goals into subgoals. These refinements reflect the different strategies and tactics that are available. Available options are constrained by the situation, which include local geography; the enemy's characteristics and capabilities; restrictions imposed by political authority; and the availability of aircraft, fuel, crews, and other resources.

The multiagent planning architecture (MPA) was used to demonstrate automated planning capability within the air-campaign planning domain (Wilkins and Myers 1998). This application starts with a high-level military objective—achieve air superiority—and expands the plan down to the level of individual missions and their support missions. The planner works from a knowledge base of planning operators (encoded specifically for the planner), which encode air-campaign tactics and strategy for goals from achieving air superiority down to mission-level goals. There are often multiple ways to refine goals into subgoals. Thus, the HTN operators at multiple abstraction levels encode what it means to achieve air superiority, a concept that would be difficult to express in primitive-action planners.

The planner has access to an extensive knowledge base of available assets (including aircraft and munitions), which was downloaded from existing military databases. Each aircraft has a dozen or more properties that affect its suitability for missions, such as speed, range, crew requirements, and munitions. Constraints on these properties appear in the HTN operators. The planner also has access to the results of the (human-conducted) intelligence analysis of the situation, which the planning operators use to focus the planner on enemy strengths, weaknesses, and other salient aspects of the situation.

The air-campaign planning application, like oil-spill response, requires extensive use of metric goals, such as deadlines, resource use, and resource accumulation. Combat missions and their support missions must compete for use of pooled resources such as fuel, aircraft, and munitions. Concurrency is important because dozens or hundreds of missions must all take place at the same time. Capacity analysis is used to determine the number of missions that a given pool of resources can support.

Evaluation of plans in this domain is com-

plex and has not been automated. Some simple measures can be computed such as whether deadlines are met and what percentage of desired targets are attacked given available resources. In MPA, the air campaign simulator (Cohen, Anderson, and Westbrook 1996) from the University of Massachusetts provided many Monte Carlo simulations of plans, and the results were presented to the user through visualization tools. Many variables could be viewed, including levels of destruction of the targets and attrition of assets for both friend and foe. However, only humans can evaluate some of the more complex effects such as political costs and benefits.

Using Knowledge in Knowledge-Based Planning

We describe some of the uses made of domain knowledge in current knowledge-based planners. These features can be candidates for extending non-knowledge-based planners, and in some cases, such extensions are currently being explored. Kautz and Selman (1998) identify three kinds of planning knowledge: (1) knowledge about the domain, (2) knowledge about good plans, and (3) explicit search-control knowledge. Knowledge-based planners are also concerned with other types of knowledge, such as knowledge about interacting with the user, knowledge about a user's preferences, and knowledge about plan repair during execution (see the discussion of expressiveness later).

Knowledge-Based Planning

Our intent is not to provide a comprehensive survey of knowledge-based planning approaches in this article, although they are sometimes ignored in other planning surveys (Weld 1999). Instead, we mention several examples of knowledge-based planners and draw our examples of knowledge use from them. HTN planning is the most studied and well understood of the knowledge-based methods. The best-known knowledge-intensive applications of HTN are SIPE-2 and O-PLAN.

Smith, Frank, and Jonsson (2000) identified a common framework that is emerging from the NASA work: the use of interval representations for actions and propositions and constraint-satisfaction techniques for reasoning about these intervals. They refer to this framework as the constraint-based interval approach. Recent knowledge-based approaches include OZONE, remote-agent experiment planner-scheduler (RAX-PS), and ASPEN. Each of these is described briefly here.

OZONE (Becker and Smith 2000; Smith, Lassi-

la, and Becker 1996), a planning and scheduling tool kit, is centered on a knowledge-intensive modeling of the problem domain. A model is specified in terms of basic types of entity, operation, resource, demand, and product. OZONE provides knowledge-structuring primitives for each of these, including several specialized resource classes. Operations can be organized hierarchically to model processes at different levels of detail.

ASPEN (Chien et al. 2000) automates planning and scheduling for space mission operations. It provides, among other capabilities, an expressive constraint modeling language, a language for representing plan preferences, constraint-reasoning systems, and a graphic interface for visualizing plans in mixed-initiative systems. These capabilities are used to model many forms of knowledge, including spacecraft operability constraints, flight rules, spacecraft hardware models, science experiment goals, and operations procedures. High-level activities can be decomposed into lower-level activities using ASPEN's activity hierarchies.

RAX-PS (Jonsson et al. 2000) generates plans that could safely be executed on the *Deep Space One* spacecraft. The plans achieve high-level goals that are provided as input to the planner while they satisfy resource constraints and complex flight safety rules. As in ASPEN, large amounts of knowledge are encoded about spacecraft resources, constraints, and procedures. RAX-PS also provides a rule language for the search controller that can be used to help avoid inefficient searches. RAX-PS incorporates specialized knowledge about the development of plan fragments from planning experts, which are generally legacy software systems or other specialized software.

SIPE-2 (Wilkins 1990; Wilkins et al. 1995) is a domain-independent HTN planner that models various types of domain knowledge. For example, SIPE-2 includes languages to represent activities at multiple levels of abstraction (HTN operators, also known as methods or schemas), knowledge about a user's preferences (Myers 1996) (which are expressed as advice to the planner), search-control knowledge, and knowledge about plan repair during execution. Example applications include containing oil spills (Agosta and Wilkins 1996), planning air campaigns for the United States Air Force (Wilkins and Myers 1998; Lee and Wilkins 1996), and joint military operations planning (Wilkins and Desimone 1994). In the latter applications, the domain knowledge includes 100 to 200 operators, around 500 objects with 15 to 20 properties to an object (which are mentioned in constraints), and a few thousand ini-

tial predicate instances. Plans can include as many as several hundred actions—several thousand if all abstraction levels are counted—usually having numerous parallel activities.

O-PLAN (Tate, Drabble, and Kirby 1994) is a domain-independent HTN planner with the ability to encode extensive domain knowledge, including temporal constraints, object-variable constraints, resource constraints, goal structure, and condition types. Plug-in constraint managers can be used to extend or modify system capabilities. O-PLAN's agenda mechanism provides flexible control of the planning and execution process. Applications include space station assembly and the control of a simple satellite.

In the following subsections, we mention features of these knowledge-based planning systems that are good candidates for extending primitive-action planners, particularly disjunctive planning techniques.

Expressiveness

It has been known for some time that HTN formalisms are more expressive than the STRIPS formalism used by most primitive-action planners, roughly analogous to the additional expressivity of context-free grammars over right-linear (regular) grammars (Erol, Hendler, and Nau 1994a).

In practice, the gap in expressiveness is wide. In the problems addressed by current knowledge-based planners, actions can occur concurrently and have different durations. Goals can include temporal deadlines and constraints, maintenance conditions, and accumulation of metric quantities of some entity. Goals and actions can be at multiple levels of abstraction. Metric resource constraints must be satisfied. All these aspects are problematic in the STRIPS formalism.

The knowledge-based approaches mentioned provide languages for expressing the types of goal and constraint mentioned earlier, making them suitable for complex domains. Figure 4 shows a hypothetical knowledge-based planning system, illustrating the range of domain knowledge, input, and output that can be required for planning in real-world domains. General domain knowledge includes knowledge about actions, tactics, and strategies, at multiple abstraction levels, as well as situation assessment information, knowledge about resources, world knowledge, and so on. Problem-specific input for a particular planning session can include goals and assumptions, constraints, additional resources that can be brought to bear, and advice or guidance from the user.

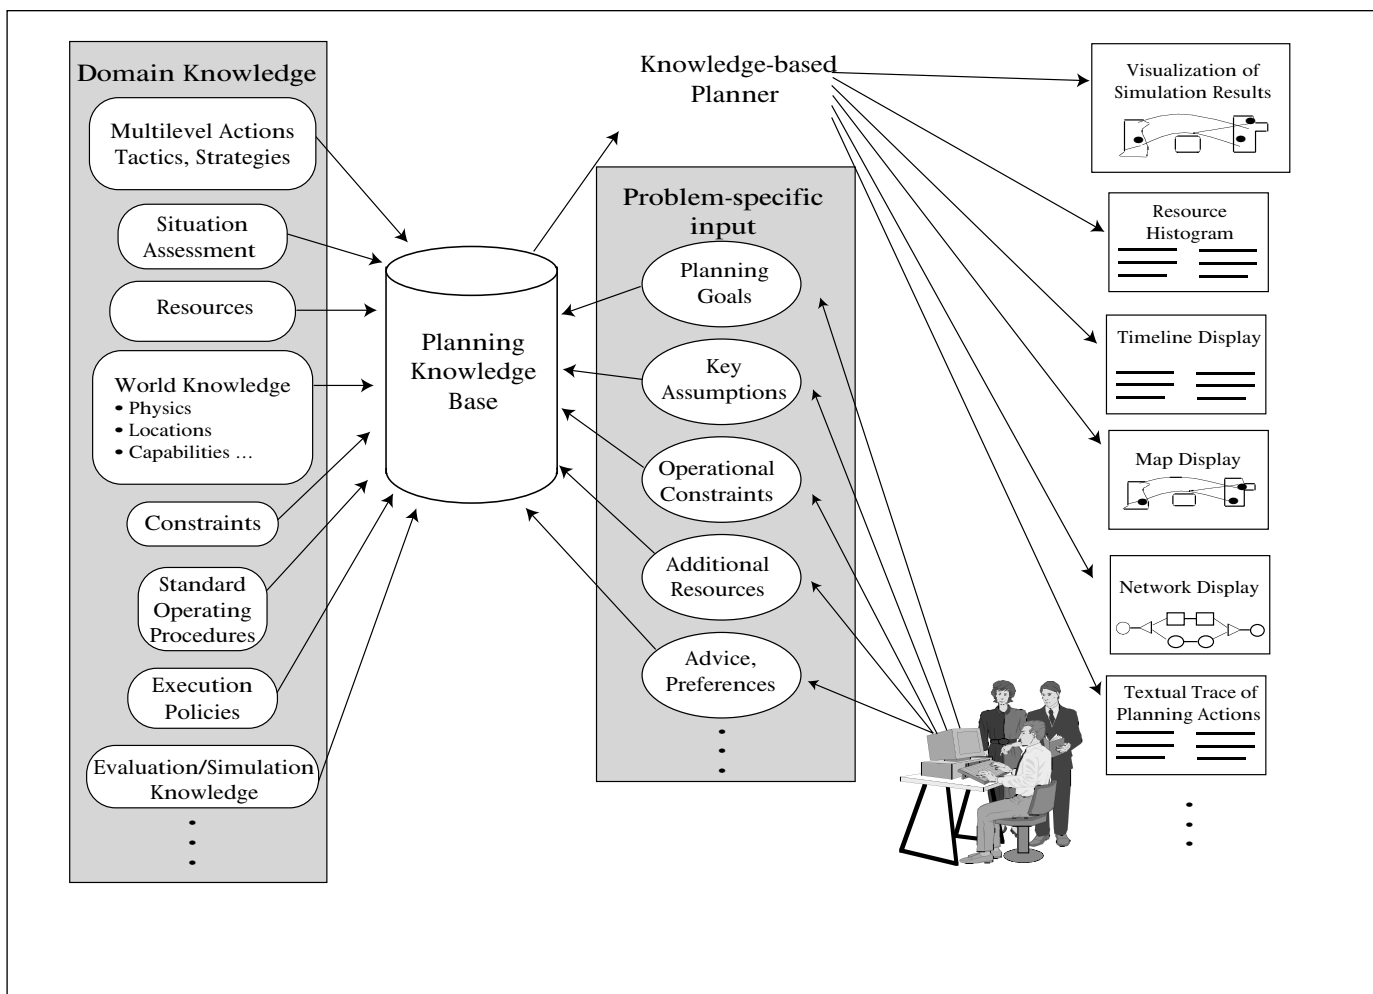


Figure 4. Potential Sources of Knowledge and Modes of Interaction for a Hypothetical Knowledge-Based Planning System.

Encoding activities at multiple abstraction levels is crucial in many complex problems. The higher levels can model various solution methods and constraints on the goal and plan structure, which can be required by the domain or desired for efficient search. The high-level goals in knowledge-based planners can only be expressed in terms of primitive actions as disjunctions of thousand or millions of possible final states (corresponding to all possible plans that achieve the high-level goal). Multiple levels can be necessary for user interaction or to support different planning interactions for different levels of management. In time-critical domains, multiple abstraction levels might be required to quickly produce a plan within the available time. ASPEN, OZONE, O-PLAN, and SIPE-2 all support hierarchical descriptions at multiple abstraction levels.

SIPE-2 operators can dynamically generate a set of goals at planning time, a capability that has been extensively used. For example, a

defend goal can be generated for every currently known threat.

All the knowledge-based planners mentioned can reason about numbers, a capability that is crucial in nearly all their applications. For example, a planning variable can be constrained to refer to a runway with length greater than 9000 feet, multiple-capacity resources that have a specific capacity, and continuous resources that are available in limited quantities. In many application domains, it is necessary to accumulate a certain quantity of some resource or achieve a certain level of effect, such as obtaining a sufficient length of boom to surround an oil spill. Such goals are not accomplished by a single action; rather, several (often concurrent) actions contribute to the accumulation. For example, SIPE-2 determines when a set of actions (that individually produce some amount of the resource in question) together achieve an accumulation goal.

Temporal reasoning is important in nearly

all complex problems. All the knowledge-based planners mentioned here have temporal constraint reasoners. In the HTN planners, they are plug-in modules that can be replaced by external temporal reasoners. For example, SIPE-2 has two different modes for reasoning about time. The most general allows specification of any of the 13 Allen relations between any two nodes. The temporal constraints are solved separately from the other constraints by passing them to TACHYON (Arthur and Stillman 1992).

Situation-dependent effects of actions are deduced by a causal theory in SIPE-2 but not supported by O-PLAN. Such effects have proven their use in practice—without them, the number of operators can grow exponentially in complex domains. As mentioned earlier, Smith and Weld have developed factored expansion to address this problem. Correcting an error in previous publications, SIPE-2 does recalculate these deductions (and always has) when new actions or ordering links are added to the plan before the action in question.

Finally, calls can be made to planning experts, which are specialized software modules, including legacy software. RAX-PS uses such experts in the development of plan fragments. In SIPE-2, functions on planning variables can compute an instantiation (for example, the duration of a flight), and procedural attachment on predicates can compute whether a condition is true. These techniques allow encoding of knowledge in arbitrary domain-specific Lisp code for knowledge that cannot easily be modeled in the planner's formalism and for sophisticated numeric calculations.

Correctness

Erol, Handler, and Nau (1994b) gave a formal definition of HTN planning and also analyzed its complexity (1994a). Since then, many HTN planners have been proven correct and complete (for example, SHOP [Nau et al. 1999]). Planning and execution in real-world domains generally cannot be modeled completely. Although particular properties of a system can formally be verified, system validation must often rely on empirical methods, which raises many challenging issues (Smith, Feather, and Muscettola 2000). Defining evaluation criteria and correctness requirements for empirical tests is another type of knowledge that must be specified.

User Interaction

In many real-world domains, plans and actions have far-reaching effects, not all of which are modeled within the planner's formalism. For example, political consequences of actions can

be important in choosing a plan but difficult or impossible to model formally. Thus, it is often necessary to have an interactive planner that allows a human expert to guide plan development. In addition, experienced human planners can guide the search effectively and are often reluctant to give control to an automated system in any case.

Hierarchical knowledge like that used by knowledge-based planners often models the world in the same way that human users do, using the same abstractions (generally provided by the human experts themselves). For example, in the air-campaign planning domain, higher-level goals include achieving air superiority and breaching the enemy's air defenses. Human users use these same abstractions and ontology, so they can naturally, for example, advise the system on how to breach air defenses or drill down to see how the air defenses were breached in an effort to understand the plan. This modeling approach enables users to control and understand the planning process and the resulting plans. (Note, however, that current HTN planners still leave much to be desired in terms of interactive planning.) In contrast, disjunctive planning approaches model the planning problem as millions of conjunctive normal form expressions, making it difficult for users even to understand the planner's reasoning process, much less intervene to modify or guide it.

Because HTN plans and domain knowledge can be complex, a powerful graphic user interface (GUI) is essential. Without natural pictorial representations of the knowledge and plans, it would be nearly impossible for a human to understand them. ASPEN, OZONE, O-PLAN, and SIPE-2 all provide a GUI to aid in generating plans, viewing complex plans, and following and controlling the planning process. The GUI can also be used to view information relevant to planning decisions.

SIPE-2 provides hyperlinked descriptions of plans and plan objects in a web server. Particularly useful for visualizing the plan derivation and structure is the ability to view a tree rooted at the most abstract goals and selectively drill down through abstraction levels for selected goals (such as breaching air defenses). When the plan contains thousands of nodes, selective drill-down is often a user's preferred method for understanding it.

Several of the knowledge-based techniques mentioned provide flexible and powerful interactive planners. For example, the user might be able to interact with the planning process at many levels of detail and can direct the planner to solve certain parts of the problem auto-

Adding Knowledge to Disjunctive Planners

Much of the effort of the planning community is currently focused on improving the performance of disjunctive planners, which embody a form of primitive-action planning. Kambhampati (1997) defines disjunctive planners as planners that retain the current plan set without splitting its components into different search branches. This family of planners includes GRAPHPLAN (Blum and Furst 1995), SATPLAN (Kautz and Selman 1996), and their derivatives. These systems all use STRIPS-style planning knowledge to represent a planning problem and then transform the problem into a propositional form that can be solved using efficient graph manipulation or constraint-satisfaction techniques.

To impact realistic problems, we predict that disjunctive planners will have to incorporate the types of knowledge used by hierarchical task network (HTN) planners as well as knowledge to overcome the limitations of HTN approaches that we have discussed. It is encouraging that this knowledge incorporation is already starting to occur. For example, there is initial work on adding knowledge about the temporal extent of actions to SATPLAN encodings (Smith and Weld 1999) and encoding HTN method knowledge for satisfiability solvers (Mali and Kambhampati 1998).

However, although HTN planners can generally make effective use of additional knowledge, the same is not necessarily true of disjunctive planners. Additional knowledge encoded as axioms can increase the size of the problem with redundant axioms and make the problem harder to solve. Initial experiments indicate that whether added knowledge helps or hurts can depend on the particular combination of knowledge, problem, and algorithm (Kautz and Selman 1998). For example, the point of diminishing returns from the addition of axioms would be sooner reached for stochastic search than for systematic search (Kautz and Selman 1998). Thus, the knowledge added to a disjunctive planner might have to be chosen carefully for the problem being solved and the algorithm being used.

Knowledge-based approaches are rightly criticized for the expense of modeling a new domain. However, we conjecture that building computationally efficient encodings for disjunctive planners of complex planning domains is no easier than building HTN models. Many encoding issues are still under study, even for toy domains (Brafman 1999; Kautz and Selman 1999; Mali and Kambhampati 1998). Disjunctive planners often start with STRIPS-based encodings, further restricting the types of action that can be encoded and causing a possible exponential explosion in the number of operators when context-dependent effects are not permitted.

Under interactive control in SIPE-2, the user can determine (among other things) when and how resources are allocated, which operators to select, which goal to expand next, how to instantiate planning variables, and how to resolve conflicts (Wilkins 1999). The user can also control or influence the plan-development process using the ADVISABLE PLANNER (Myers 1996), which allows users to direct the planning process by providing high-level guidance that influences the nature of the solutions generated. Advice consists of task-specific properties of both the desired solution and the problem-solving process to be used for a particular problem.

Constraints and Efficiency

It is sometimes argued that the knowledge used by HTN planners is simply search-control knowledge rather than part of the problem statement. (We argued earlier that knowledge-based planners encode much more than just search-control knowledge.) However, if the goal is to solve realistic planning problems, then intelligent, principled search control that takes advantage of knowledge about the domain is precisely what is needed. This knowledge can often be captured naturally and efficiently in HTN operators, where much of the context is implicit and therefore need not be expressed or checked during each attempted application.

OZONE, SIPE-2, and other systems represent invariant object properties in a hierarchical ontology, which describes the classes to which an object belongs and allows for inheritance of properties. The ontology encodes a large amount of knowledge, and the planner can reason more efficiently about this knowledge because it knows that the relationships cannot change as actions are performed.

O-PLAN and SIPE-2 have, for a long time, separated various classes of constraints for efficiency, for example, solving temporal constraints and constraints on static attributes of objects separately. O-PLAN pioneered the use of modular, specialized constraint solvers used at certain intervals. This separation of constraints is an example of an HTN idea that has migrated to disjunctive planners in recent work, showing that the performance of disjunctive planners can be improved by separating out resource reasoning to prevent thrashing (Srivastava and Kambhampati 1999).

How soon to make commitments in the search depends on the search strategy being used and the problem being solved. SIPE-2 and O-PLAN have developed techniques to exploit the least-commitment approach (Myers and

Wilkins 1998; Tate, Drabble, and Kirby 1994; Wilkins 1990). For example, in SIPE-2, constraints are placed on variables by domain knowledge in the operators (for example, a particular truck must have a capacity greater than 100). Instantiations are not chosen until sufficient constraints accumulate to identify a unique acceptable value. (Early instantiation might result in a poor choice and failure to find a solution leading to a possible exploration of a large search space.) Because uninstantiated variables increase computational complexity, domain-specific knowledge can be used to require early instantiation of variables (by software planning experts) at particular points in the planning process. Human experts often know when certain instantiations can be done without adversely affecting solution quality (Myers and Wilkins 1998). Thus, this knowledge can improve performance without sacrificing quality.

In some knowledge-based systems, predicates can be declared as functional in certain arguments, allowing a dramatic speedup, which has been documented experimentally (Myers and Wilkins 1998). Functional predicates are of particular importance to reasoning about locations in planning systems and have proven valuable in nearly every application of SIPE-2 as well as in procedural reasoning systems (Georgeff and Ingrand 1989).

HTN systems often rely on plan critics that find conflicts or flaws in a plan. Plan critics can be invoked after some number of plan modifications, rather than after every modification, thus reducing computational costs during plan expansion. Examples of plan critics include the finding of resource conflicts, failed preconditions, or unsatisfiable constraints. In SIPE-2, domain knowledge can be used to increase or decrease the frequency of plan-critic application (Wilkins 1990).

Finally, OZONE, ASPEN, and O-PLAN include specialized resource classes. The system can efficiently implement specialized reasoners for these classes, instead of trying to represent, for example, multiple-capacity resources in the underlying planning formalism.

Knowledge beyond Hierarchical Task Networks

Despite the power of HTN planning systems, and their demonstrated ability to address real-world planning problems, they have limitations that make them inadequate for many problems of interest. In particular, HTN planners (1) require complete (except for anticipated incompleteness) and certain knowledge

about the world; (2) model the effects of actions as deterministic, fully understood outcomes; (3) assume that the planner controls all agents that cause changes in the world state; (4) require significant effort in domain modeling and knowledge acquisition for complex problems; (5) cannot perform or incorporate complex or decision-theoretic evaluations of plan quality; (6) ignore the qualification problem; (7) use simplistic frame-problem solutions that prevent drawing the most appropriate conclusions when contradictory (perceptual) information arrives (Pollack 1998); (8) do not consider risks and utilities; (9) do not use knowledge and probabilities to handle uncertainty; and (10) are brittle (might not work if the problem changes slightly). These limitations are shared by primitive-action planners, although some of them, such as handling uncertainty, are the subject of ongoing research (Boutilier, Dean, and Hanks 1999; Onder and Pollack 1999; Majercik and Littman 1998; Smith and Weld 1998; Weld, Anderson, and Smith 1998; Kushmerick, Hanks, and Weld 1994).

In addition to these limitations, planning systems that could solve interesting problems in a complex, dynamic world will need capabilities that represent a fundamental shift in how we think about planning problems. An ideal system would be able to behave like humans do in these sorts of environments; in particular, it would have to exhibit creativity, devising new actions that can solve a problem or shorten a plan; use analogy to transfer solutions from other problems; effectively interact with humans to use their knowledge in decisions; and behave intelligently in the face of conflicting or incomplete information.

We believe that these capabilities will require more knowledge, including background knowledge of what other domains are and how the world works.

Erol, Handler, and Nau (1994b) showed that if a domain is completely modeled, then an HTN planner can provide a guarantee that the plans it produces are correct with respect to this domain model. However, for realistic domains, evaluation criteria other than correctness and plan length will have to be factored in explicitly. Interacting effectively with humans will be essential because we will never model every possible issue that might affect a planning decision. Humans often have evaluation criteria that cannot be captured precisely and have commonsense knowledge that allows them to determine appropriate actions in unusual situations that were unforeseen when the domain was modeled. Plans that are pro-

An ideal system would exhibit creativity, use analogy, [and] interact with humans to use their knowledge in decisions

duced should be evaluated in the same way that plans produced by humans are evaluated, for example, scoring performance against a simulator or the real world or having human experts evaluate the plans by hand.

Of course, not all interesting problems have these characteristics, and in any given case, it might be possible to formulate the problem in such a way as to remove the need for these capabilities. For example, in developing the Burton planner, Williams and Nayak (1996) used a purely propositional representation. However, it seems unlikely that most interesting problems will be amenable to such an approach, and other NASA applications have required richer representations (Chien et al. 1996).

All system types mentioned earlier (including both disjunctive planners and the knowledge-based planners mentioned) ignore the qualification problem and have simplistic solutions to the frame problem (Pollack 1998). These issues must be addressed by future knowledge-based planners. Among other problems with traditional planners, Pollack shows that new perceptions that contradict old assumptions cause difficulty. Pollack's system handles these problems more robustly, although it does not currently appear scalable to real-world problems. Not surprisingly, significantly more knowledge must be encoded, such as knowledge about causation, defeasibility, and when causations can be "undercut."

Lessons from Machine Learning

In every research community, there is an ongoing tension between well-defined and more ambitious problems. On the one hand, if a field focuses on small, well-understood problems, with well-defined algorithmic properties and evaluation metrics, then a set of benchmark problems can be formulated to facilitate formal and empirical analysis and comparison of competing methods. On the other hand, many of the interesting challenges posed by realistic applications have broader implications and less well-understood properties, and the problems are more difficult to define crisply and to evaluate.

Several years back, the machine learning community established a repository of benchmark problems to evaluate machine learning systems. Naturally, these problems all had commonalities: Most used an attribute-vector representation; most consisted of sets of instances with no background knowledge. In practice, they could be used only to evaluate predictive accuracy on propositional, super-

vised learning algorithms. Despite these limitations, however, it became the de facto standard that papers submitted to the International Conference on Machine Learning had to include an evaluation on these benchmark problems.

In some ways, these benchmarks, and the emphasis on evaluation, were good for the community: On the one hand, they forced researchers to think about metrics and the comparison of their systems to other systems, and they provided a baseline of performance against which researchers could test new ideas. On the other hand, they tended to stifle research that did not fit neatly into the problem space defined by the benchmark problems. Application-oriented researchers reported that it was difficult or impossible to get their papers accepted to the leading machine learning conferences (Provost and Kohavi 1998). Meanwhile, more and more papers appeared showing minor tweaks and incremental improvements to existing algorithms (but they showed statistically significant improvements on the benchmark algorithms!).

As a result, there are now many well-understood and effective methods for propositional supervised learning—and there has been much less progress in other areas of machine learning, such as the incorporation of background knowledge, feature engineering, relational learning, interactive learning techniques, visualization of learned knowledge, and complex evaluation criteria.

Most recently, there has been an explosion of interest in learning Bayesian networks. Bayesian network learning and inference techniques have appealing computational properties that are analogous to those of disjunctive planning approaches: They efficiently capture certain types of problem structure and significantly speed up certain types of inference over previous methods. However, like disjunctive planning approaches, they use a propositional representation and do not address many of the other challenges posed by realistic learning problems. As with disjunctive planning approaches, the rush of enthusiasm over Bayesian network techniques has threatened to overshadow the fact that despite their computationally attractive properties, they still solve only a small subproblem within the overall field of machine learning.

Similarly, in the planning community, there is a danger that by focusing too much attention and effort on disjunctive planning methods and the problems they solve, we risk losing the ability to recognize other contributions and advances. In particular, if the benchmark

of performance becomes solely how many blocks our planners can stack, and how fast they can do it, then it will become increasingly difficult to recognize and learn from research that performs well along other dimensions—or that addresses problems that disjunctive planning systems overlook completely. As we discussed earlier, disjunctive planning researchers within the planning community are starting to look toward extending their systems to incorporate richer forms of knowledge. We applaud this trend and hope that it will continue, but it is not enough to simply broaden the uses of disjunctive planning systems: We need to be open to completely different approaches and paradigms as well.

Although improving the speed of solving problems we know how to formulate precisely is a valuable research activity, so is continuing to investigate problems that we do not yet have a good handle on formulating or solving. Results might be more difficult to achieve or quantify for these latter problems, but it does not mean we should not be working on them.

Conclusion

Ginsberg (1996) pointed out that the SATPLAN approach is successful because it solves the puzzle part of a problem and overlooks any commonsense reasoning aspects of the true problem. In Ginsberg's view, commonsense reasoning is the heuristic process by which we reduce extremely complex problems to NP-hard or simpler problems for which search is feasible. The aspects of a problem to pay attention to, frame and context assumptions, and default strategies for organizing complex activities are all aspects of commonsense reasoning. As Ginsberg (1996) puts it, "It is Kautz and Selmán who are solving the commonsense aspects of the problem; their 'planner' is solving the puzzle-mode kernel of the problem instead of the problem itself" (p. 624). Indeed, the problems solved by primitive-action approaches are almost exclusively puzzle-style problems (or real-world problems that have been reformulated as puzzles).

We favor using primitive-action methods to solve puzzle-style subproblems that can be handled by constraint-satisfaction engines in acceptable time. However, AI planners also need to provide support for the commonsense reasoning aspects of the problem so that plans can be used to guide behavior while they are embedded in a complex, dynamic environment. We have argued that incorporating knowledge, encoded in understandable struc-

tures, into the planning process is the most promising way to provide these abilities. HTN planning methods are better suited than disjunctive planners for such problems because HTN systems can interact with humans effectively, use more expressive representations, and can make use of domain knowledge to scale up to complex problems. However, HTN methods still have significant limitations, and we have argued that one must use still more knowledge (both in quantity and in quality) than HTN planners do to solve the hardest problems.

Although primitive-action methods are clearly useful approaches for solving certain subproblems, it is important for the field as a whole to continue to look at a wider range of problems. There is a danger of allowing the current popularity of disjunctive planning approaches, and the associated evaluation techniques and puzzle-style problem suite, to overly influence the field, making it more difficult for advanced knowledge-based planning methods to find an audience.

Acknowledgments

This research was supported by contract F30602-95-C-0235 with the Defense Advanced Research Projects Agency under the supervision of Air Force Research Lab-Rome. Thanks to Dana Nau, Steve Chien, and Foster Provost for contributing their ideas to this article and to Steve Chien for parts of the section on space applications.

Notes

1. The term *domain-independent planning* has sometimes been used in the literature to describe these systems, but this term is a misnomer because the action descriptions do in fact constitute a form of domain knowledge. In this article, we use the term *domain independent* to refer to any planner that is designed to be applicable generically to any domain encoded in the proper form. Using this definition, knowledge-based planning methods that operate on explicit domain descriptions—as opposed to systems that are hard wired for particular domains—are as domain independent as primitive-action planning systems.
2. Although simple, finite arithmetic could be added to disjunctive planners, the combinatorics would generally explode. Another approach is given by Wolfman and Weld (1999), who describe a system that combines SATPLAN with an incremental SIMPLEX algorithm for solving linear inequalities, a useful extension, but combinatorics allow the solution of only toy problems.
3. These descriptions were generously provided by Steve Chien of NASA's Jet Propulsion Lab; however, the authors take responsibility for the final text.

References

- Agosta, J. M., and Wilkins, D. E. 1996. Using SIPE-2 to Plan Emergency Response to Marine Oil Spills. *IEEE Expert* 11(6): 6–8.
- Anderson, C.; Smith, D. E.; and Weld, D. 1998. Conditional Effects in GRAPHPLAN. In *Proceedings of the 1998 International Conference on AI Planning Systems*, 44–53. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Arthur, R., and Stillman, J. 1992. TACHYON: A Model and Environment for Temporal Reasoning. Technical report, GE Corporate Research and Development Center, Schenectady, New York.
- Bacchus, F.; Kautz, H.; Smith, D. E.; Long, D.; Geffner, H.; and Koehler, J. 2000. The Fifth International Conference on Artificial Intelligence Planning and Scheduling: Planning Competition. Available at www.cs.toronto.edu/aips00.
- Becker, M. A., and Smith, S. F. 2000. Mixed-Initiative Resource Management: The AMC Barrel Allocator. In *Proceedings of the 2000 International Conference on AI Planning and Scheduling*, 32–41. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Blum, A., and Furst, M. 1995. Fast Planning through Planning Graph Analysis. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 1636–1642. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Boutilier, C.; Dean, D.; and Hanks, S. 1999. Decision-Theoretic Planning: Structural Assumptions and Computational Leverage. *Journal of Artificial Intelligence Research* 11:1–94.
- Brafman, R. I. 1999. Reachability, Relevance, Resolution, and the Planning as Satisfiability Approach. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, 976–981. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Chien, S.; Rabideau, G.; Willis, J.; and Mann, T. 1999. Automating Planning and Scheduling of Shuttle Payload Operations. *Artificial Intelligence* 114(1): 239–255.
- Chien, S.; Hill-Jr, R.; Wang, X.; Estlin, T.; Fayyad, K.; and Mortensen, H. 1996. Why Real-World Planning Is Difficult: A Tale of Two Applications. In *Advances in AI Planning*, eds. M. Ghallab and A. Milani, 287–298. Amsterdam: IOS.
- Chien, S.; Rabideau, G.; Knight, R.; Sherwood, R.; Engelhardt, B.; Mutz, D.; Estlin, T.; Smith, B.; Fisher, F.; Barrett, T.; Stebbins, G.; and Tran, D. 2000. ASPEN Automating Space Mission Operations Using Automated Planning and Scheduling. Paper presented at SpaceOps 2000, 19–23 June, Toulouse, France.
- Cohen, P.; Anderson, S.; and Westbrook, D. 1996. Simulation for ARPI and the Air Campaign Simulator. In *Advanced Planning Technology: Technological Achievements of the ARPA/Rome Laboratory Planning Initiative*, ed. A. Tate, 113–118. Menlo Park, Calif.: AAAI Press.
- Erol, K.; Hendler, J.; and Nau, D. S. 1994a. HTN Planning: Complexity and Expressivity. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1123–1128. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Erol, K.; Hendler, J.; and Nau, D. S. 1994b. UMCP: A Sound and Complete Procedure for Hierarchical Task-Network Planning. Paper presented at the 1994 International Conference on AI Planning Systems, 13–15 June, Chicago, Illinois.
- Georgeff, M. P., and Ingrand, F. F. 1989. Decision-Making in an Embedded Reasoning System. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 972–978. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Ginsberg, M. L. 1996. Do Computers Need Common Sense? Paper presented at the Fifth International Conference on Knowledge Representation and Reasoning, 5–8 November, Cambridge, Massachusetts.
- Guere, E., and Alami, R. 1999. A Possibilistic Planner That Deals with Non-Determinism and Contingency. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, 996–1001. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Jonsson, A.; Morris, P.; Muscettola, N.; and Rajan, K. 2000. Planning in Interplanetary Space: Theory and Practice. In *Proceedings of the 2000 International Conference on AI Planning and Scheduling*, 177–186. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Kambhampati, S. 1997. Challenges in Bridging Plan Synthesis Paradigms. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, 44–49. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Kambhampati, S.; Parker, E.; and Lambrecht, E. 1997. Understanding and Extending GRAPHPLAN. Paper presented at the Fourth European Conference on Planning, 24–26 September, Toulouse, France.
- Kautz, H., and Selman, B. 1999. Unifying SAT-Based and Graph-Based Planning. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, 318–325. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Kautz, H., and Selman, B. 1998. The Role of Domain-Specific Knowledge in the Planning as Satisfiability Approach. In *Proceedings of the 1998 International Conference on AI Planning Systems*, 181–189. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Kautz, H., and Selman, B. 1996. Planning as Satisfiability. In *Proceedings of the Tenth European Conference on Artificial Intelligence*, 359–363. New York: Wiley.
- Knoblock, C. 1996. AI Planning Systems in the Real World. *IEEE Expert* 11(6): 4–12.
- Kushmerick, N.; Hanks, S.; and Weld, D. 1994. An Algorithm for Probabilistic Least Commitment Planning. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1073–1078. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Lee, T. J., and Wilkins, D. E. 1996. Using SIPE-2 to Integrate Planning for Military Air Campaigns. *IEEE Expert* 11(6): 11–12.
- Long, D. 2000. The AIPS-98 Planning Systems Competition. *AI Magazine* 21(2): 13–33.
- McDermott, D. 2000. The 1998 AI Planning Systems Competition. *AI Magazine* 21(2): 35–55.
- Majercik, S., and Littman, M. 1998. MAXPLAN: A New Approach to Probabilistic Planning. In *Proceedings of the 1998 International Conference on AI Planning Systems*, 86–93. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Mali, A. D., and Kambhampati, S. 1998. Encoding HTN Planning in Propositional Logic. In *Proceedings of the 1998 International Conference on AI Planning Systems*, 190–198. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Myers, K. L. 1999. CPEF: A Continuous Planning and Execution Framework. *AI Magazine* 20(4): 63–70.
- Myers, K. L. 1996. Strategic Advice for Hierarchical Planners. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifth International Conference (KR '96)*, eds. L. C. Aiello, J. Doyle, and S. C. Shapiro, 112–123. San Francisco, Calif.: Morgan Kaufmann.
- Myers, K. L., and Wilkins, D. E. 1998. Reasoning about Locations in Theory and Practice. *Computational Intelligence* 14(2): 151–187.
- Nau, D.; Cao, Y.; Lotem, A.; and Munoz-Avila, H. 1999. SHOP: Simple Hierarchical Ordered Planner. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, 968–983. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

- Onder, N., and Pollack, M. E. 1999. Conditional, Probabilistic Planning: A Unifying Algorithm and Effective Search Control Mechanisms. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, 577–584. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Pollock, J. L. 1998. Perceiving and Reasoning about a Changing World. *Computational Intelligence* 14(4): 498–562.
- Provost, F., and Kohavi, R. 1998. Guest Editors' Introduction: On Applied Research in Machine Learning. *Machine Learning* 30(2–3): 127–132.
- Smith, B.; Feather, M.; and Muscettola, N. 2000. Challenges and Methods in Validating the Remote Agent Planner. In *Proceedings of the 2000 International Conference on AI Planning and Scheduling*, 254–263. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Smith, D., and Weld, D. 1998. Conformant GRAPHPLAN. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, 889–896. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Smith, D. E., and Weld, D. S. 1999. Temporal Planning with Mutual Exclusion Reasoning. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, 326–333. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Smith, D. E.; Frank, J.; and Jonsson, A. K. 2000. Bridging the Gap between Planning and Scheduling. *Knowledge Engineering Review* 15(1): 61–94.
- Smith, S.; Lassila, O.; and Becker, M. 1996. Configurable, Mixed-Initiative Systems for Planning and Scheduling. In *Advanced Planning Technology: Technological Achievements of the ARP/ARPA Laboratory Planning Initiative*, 235–241. Menlo Park, Calif.: AAAI Press.
- Srivastava, B., and Kambhampati, S. 1999. Scaling Up Planning by Teasing Out Resource Scheduling. In *Proceedings of the Fifth European Conference on Planning*, 8–10 September, Durham, United Kingdom.
- Tate, A.; Drabble, B.; and Kirby, R. B. 1994. O-PLAN2: An Open Architecture for Command, Planning, and Control. In *Intelligent Scheduling*, eds. M. Fox and M. Zweben, 213–239. San Francisco, Calif.: Morgan Kaufmann.
- Weld, D. S. 1999. Recent Advances in AI Planning. *AI Magazine* 20(2): 93–123.
- Weld, D. S.; Anderson, C. R.; and Smith, D. E. 1998. Extending GRAPHPLAN to Handle Uncertainty and Sensing Actions. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, 897–904. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Wilkins, D. E. 1999. Using the SIPE-2 Planning System: A Manual for Version 6.0. SRI International Artificial Intelligence Center, Menlo Park, California.
- Wilkins, D. E. 1990. Can AI Planners Solve Practical Problems? *Computational Intelligence* 6(4): 232–246.
- Wilkins, D. E., and Desimone, R. V. 1994. Applying an AI Planner to Military Operations Planning. In *Intelligent Scheduling*, eds. M. Fox and M. Zweben, 685–709. San Francisco, Calif.: Morgan Kaufmann.
- Wilkins, D. E., and Myers, K. L. 1998. A Multiagent Planning Architecture. In *Proceedings of the 1998 International Conference on AI Planning Systems*, 154–162. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Wilkins, D. E.; Myers, K. L.; Lowrance, J. D.; and Wesley, L. P. 1995. Planning and Reacting in Uncertain and Dynamic Environments. *Journal of Experimental and Theoretical AI* 7(1): 197–227.
- Williams, B. C., and Nayak, P. P. 1996. A Model-Based Approach to Reactive Self-Configuring Systems. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 971–978. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Willis, J.; Rabideau, G.; and Wilklow, C. 1999. The CITIZEN EXPLORER Scheduling System. In *Proceedings of the IEEE Aerospace Conference*. Washington, D.C.: IEEE Computer Society.
- Wolfman, S. A., and Weld, D. S. 1999. The LPSAT Engine and Its Application to Resource Planning. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, 310–316. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.



David E. Wilkins is a senior computer scientist in the Artificial Intelligence Center at SRI International. A fellow of the American Association for Artificial Intelligence, he received his Ph.D. in computer science from Stanford University in 1979. His current research focuses on continuous

planning and execution monitoring, mixed-initiative planning, and multiagent planning. He is interested in the application of these technologies to real-world problems. His e-mail address is wilkins@ai.sri.com.



Marie desJardins is a senior computer scientist at SRI International. Her ongoing research projects are developing methods for multiagent planning and negotiation, mixed-initiative planning, and model visualization for machine learning. Other research interests include probabilistic reasoning,

decision theory, and knowledge representation. desJardins was awarded a Ph.D. in AI from the University of California at Berkeley in 1992, where her dissertation presented a model for autonomous machine learning in probabilistic domains. Her e-mail address is marie@ai.sri.com.