Acquiring Criteria for Plan Quality Control [1]

Brian Drabble
AI Applications Institute
University of Edinburgh
80 South Bridge
Edinburgh EH1 1HN, Scotland
B.Drabble@ed.ac.uk

Yolanda Gil
Information Sciences Institute
University of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292
Gil@isi.edu

Austin Tate
AI Applications Institute
University of Edinburgh
80 South Bridge
Edinburgh EH1 1HN, Scotland
A.Tate@ed.ac.uk

Extended Abstract

October 28, 1994

1 Problem addressed

In producing plans, human planners take into account a variety of criteria that guide their decisions. Besides constraints imposed by the domain itself, these criteria often express preferences among alternative plans that achieve the given goals. Human planners can use these criteria for two important purposes:

- when asked to generate one plan, human planners are able to discern between an ordinary solution and a better quality one and propose the latter.
- when asked to generate several alternative plans, human planners are able to discern between similar alternative solutions and qualitatively different ones relaxing different criteria to explore tradeoffs.

Current AI planners are good at generating a solution that satisfies the goals that they are given. AI planners often have a facility to use a plan evaluation function for the immediate needs of selecting a good solution during the search for a single plan. However, they do not usually integrate quality considerations across several plans. The quality criteria that human expert planners consider:

- are highly dependent on the situation and the scenario at hand. Some criteria may be more important if there is a certain deadline, or new criteria may need to be considered if new considerations come up.
- include complex factors and tradeoffs that are often not represented by an automatic planner

[1] We would like to thank Mark Hoffman, Bing Leng, and Bill Swartout for their help with various aspects of this work. The EXPECT project is supported by the Advanced Research Projects Agency under contract DABT63-91-C-0025. O-Plan work is supported by the US Advanced Research Projects Agency (ARPA) and the US Air Force Rome Laboratory acting through the Air Force Office of Scientific Research (AFSC) under contract F49620-92-C-0042. The view and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of ARPA or the U.S. Government.

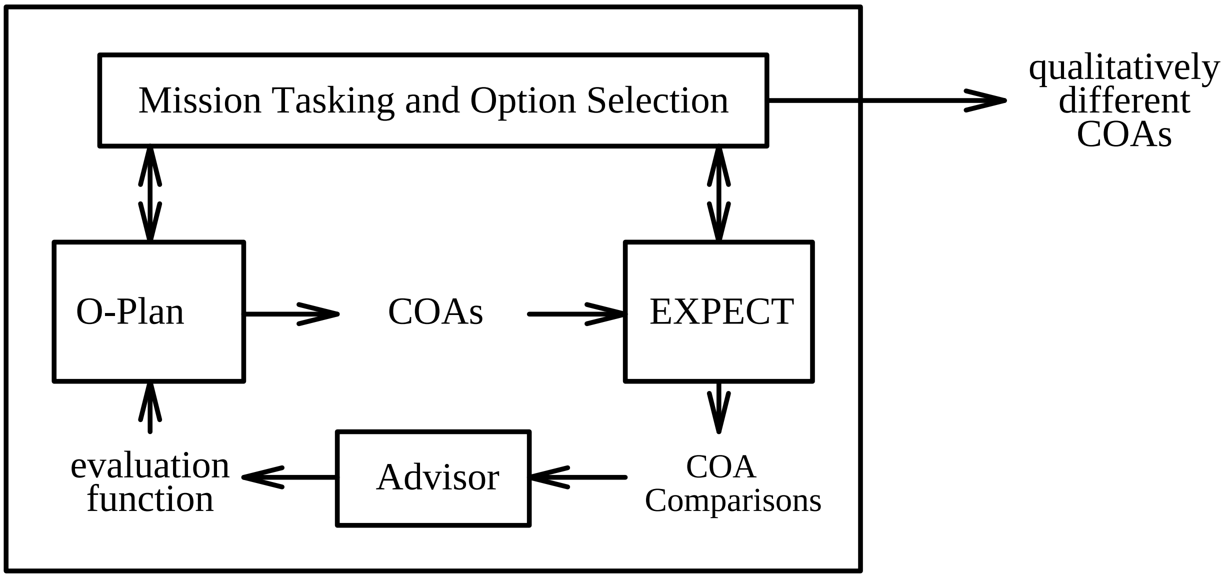

Figure 1: A generative planner and a knowledge-based system cooperate to produce better plans.

Thus, evaluating plan quality requires both complex reasoning abilities and sophisticated knowledge acquisition tools that current planning technology lacks. The goals of this work are twofold:

- to provide tools that allow expert planners to define criteria for plan quality and for preferences among alternative plans
- to operationalize these criteria to guide generative planners in proposing better quality plans.

2 Approach

Our approach combines a generative planner with a knowledge-based system that reasons about plan evaluation. Knowledge-based system technology enables us to build an interface between the planner and the user that provides the following functionality:

- support the user in defining criteria for evaluating plan quality through a knowledge acquisition tool
- evaluate the quality of plans proposed by the planner
- provide justifications for good and bad plan quality

Figure 1 shows an overview of the approach. The knowledge-based system evaluates the solutions generated by the planner according to the criteria for plan quality that are currently known to the system. The knowledge acquisition tool communicates with a domain specialist to acquire the plan evaluation criteria that are relevant to the planning domain in general or to a specific problem instance. The result of the knowledge acquisition process is an updated set of criteria that may include new criteria or more precise definitions of already existing ones. The advice generation module takes plan evaluations and their justifications and analyzes them together with the user-defined criteria to produce planner-independent advice that can be used to guide planning choices. The advice operationalization module integrates the advice generation module with a particular generative planner. This module takes generic advice and produces advice in a language specific to the planner that will use it.

This approach will result on a closed-loop integration of plan generation and plan quality evaluation that will let the user guide a planner in finding the desired kind of solutions.

3 The Domain: Transportation Planning

Our work is motivated by the transportation planning domain that is the focus of the ARPA/Rome Laboratory Planning Initiative. This domain involves the movement of materials and forces with a mixture of aircraft and ships. The task is to have the materials in place by a designated starting date usually referred to as D-DAY. This is simple to state but often difficult to achieve in practice. The main problem is that materials move through a number of ports and airfields which have finite capacities in terms of warehouse and parking space. In addition a number of support personnel are required to monitor and operate these facilities. The forces and materials to be moved are identified and a fixed number of transport assets are provided by the US Transportation Command. The number and make up of the forces can vary and as such a number of alternative plans or courses of actions (COAs) can be generated. These COAs are plans that specify at a high level the sequences of actions for movement and employment of forces. The commander of the operation is presented with several alternative COAs and an evaluation of the tradeoffs among them. These options are explored and different aspects/variables altered to identify potential new COAs. A decision is finally made on the scale of the mission and the chosen COA needed to support it. This COA is refined to a more detailed level with improved plan feasibility estimators.

A version of this domain is available upon request in a fictitious but realistic scenario called PRECiS, and is described in detail in [4, 3].

4 Current Status

We are exploring this approach using O-Plan [1, 6] as the plan generation system and the EXPECT system [5, 2] as the knowledge-based framework.

The O-Plan Project is exploring a practical computer based environment to provide for specification, generation, interaction with, and execution of activity plans. O-Plan is intended to be a domain-independent general planning and control framework with the ability to embed detailed knowledge of the domain. In O-Plan, a user specifies a task that is to be performed through some suitable interface. We call this process *task assignment*. Then, a *planner* plans and (if requested) arranges to execute the plan to perform the task specified. The *execution system* seeks to carry out the detailed activities specified by the planner while working with a more detailed model of the execution environment. The Edinburgh O-Plan prototype is currently being demonstrated generating transportation plans of the kind shown in Figure 2.

EXPECT is an architecture for the development of knowledge-based systems that includes a run-time environment, a natural language explanation facility, and a knowledge acquisition tool. We have developed with EXPECT a prototype system that takes an assessment of the situation and evaluates relevant factors for the alternative courses of action from the logistics perspective. The system has a map-based interface that displays force deployment, and allows the user to analyze factor evaluations. The user can correct the system's knowledge about how to compute these evaluations if a knowledge deficiency is detected. The user can also correct the system's knowledge base to add new relevant factors or to expand the level of detail at which the evaluations are computed.

At present, we have defined the interface between both systems so that EXPECT can evaluate the plans generated by O-Plan. Given a mission statement, O-Plan generates a plan that accomplishes the mission and includes force deployment data. EXPECT takes this plan and evaluates it from a logistics perspective, producing estimates of the relevant factors such as support personnel required and deployment closure date. This would be done at the Task Assignment level of the O-Plan architecture. When there are several alternative plans generated, the result is a comparison matrix that is useful to human planners in identifying the best alternative. Figure 2 shows a sample plan generated by O-Plan. A plan is composed, among other things, of a set of movements that specify when and where to transport a force, the ports of embarkation and debarkation, the lift resources available, and the amount of passengers (PAX) and cargo of different categories in the force to be moved (bulk, oversize, outsize, and non aerotransportable). Figure 3 shows EXPECT's evaluations of several alternative plans provided by O-Plan. The figure shows five main evaluation criteria: the airport and seaport facilities, the closure date of the operation (i.e., the day when all the forces have arrived to their destinations), the amount of logistics personnel needed to support the operation, and the lines of communications. In this example, COA 1 the closure date of the operation would be the latest one, but it uses few logistics personnel to support the operation. COA 2 and COA3 are more or less equivalent, i.e., not qualitatively different. COA 4 requires the most support personnel, but it closes early and offers better seaport and airport facilities than the others. Based on this kind of tradeoff analysis, one of the four courses of action is chosen by the commander.

POE/POD	WHO	WHEN	LIFT	PAX	Bulk	Over	Out	NAT
Andrews-afb/ Tunis-Carthage	107th-ACR	21	3 C-141	5492	1362	0	0	0
Andrews-afb/ Tunis-Carthage	107th-ACR	21	3 C-141	0	0	2670	2581	0
Wilmington-port/ Tunis-port	107th-ACR	21	2 FSS	0	0	10678	10324	0
Ft-Meade-port/ Tunis-port	107th-ACR	21	3 FSS	0	0	0	0	83250
Rota-port/ Palermo-port	BB62-SAG	15	10 FSS	3748	68	0	0	154
Shaw-afb/ Tunis-Carthage	16th-F15E	25	4 C-141	794	199	0	0	0
Shaw-afb/ Tunis-Carthage	16th-F15E	25	4 C-141	0	0	59	0	0
Charleston-port/ Tunis-port	16th-F15E	25	10 FSS	0	0	0	0	2553

Figure 2: Summary of a sample O-Plan solution to a transportation problem.

	COA 1	COA 2	COA 3	COA4
A-PORTS:				
- airports	1	1	1	2
- sorties/hr	315	315	315	480
- sq ft ac parking	2M	2M	2M	3M
S-PORTS:				
- seaports	1	1	1	2
- piers	6	6	6	15
- berths	6	6	6	16
- max vessel size in ft	600	600	600	765
- oil facilities	1	1	1	3
CLOSURE DATE	C + 29	C + 22	C + 23	C + 23
LOG PERS	1154	5360	5396	7362
LOCs:				
- number locations	1	5	7	6
- miles max distance	20	99	140	120
- air and sea?	yes	yes	yes	yes

Figure 3: EXPECT'S evaluation of several alternative plans generated by O-Plan.

5 Future Work

We plan to extend EXPECT and O-Plan to strengthen the ability to support a user in specifying, comparing and refining the constraints on qualitatively different plans at the task assignment level of a planning support environment. EXPECT needs to be provided with the ability to look at how the evaluations are derived and extract justifications that record the dependencies between plan features and the evaluations. These justifications can serve as a basis to interact with the planning experts to acquire advice on which feature values produce better results in the evaluations. O-Plan needs to be extended to effectively operationalize that advice, i.e., to transform that advice into information that is useful to its evaluation function. Once the nature of that advice is better understood, we will use it as a basis to define the functionality and to develop the advice generation and operationalization modules.

References

[1] Currie, K.W. and Tate, A., *O-Plan: the Open Planning Architecture*, Artificial Intelligence, Vol. 51, No. 1, North-Holland, Autumn 1991.

[2] Gil, Y. *Knowledge Refinement in a Reflective Architecture* Proceedings of the Twelfth National Conference on Artificial Intelligence, Seattle, WA, USA. August 1994. Published by AAAI Press/ The MIT Press Menlo Park, CA, USA.

[3] Gil, Y., Hoffman, M., and Tate, A. *Domain-Specific Criteria to Direct and Evaluate Planning Systems*, presented at the 1994 Arpa/Rome Laboratory Planning Initiative Workshop, Tucson, Arizona, USA. Published by Morgan Kaufman, San Francisco, CA, USA. Also available as Information Sciences Institute Technical Report ISI-RR-93-365, University of Southern California.

[4] Reece, G., Tate, A., Brown, D., and Hoffman, M. *The PRECiS Environment*, papers of the ARPA/Rome Laboratory Planning Initiative Workshop at the National Conference on Artificial Intelligence (AAAI-93), Washington D.C., USA. ARPI Report ARPA-RL/CPE/Version 1, August 1993. Also available as Artificial Intelligence and Applications Institute Technical Report AIAI-TR-140, University of Edinburgh.

[5] Swartout, W. R., Paris, C. L., and Moore, J. D. *Design for explainable expert systems.* IEEE Expert 6(3):58–64. 1991.

[6] Tate, A., Drabble, B. and Kirby, R., *O-Plan2: an Open Architecture for Command, Planning and Control*, in Intelligent Scheduling, (eds. M.Fox, M. and M.Zweben), Morgan Kaufmann, 1994.