Yes, But Why is That Plan Better? [1]

Dr Brian Drabble
AI Applications Institute
University of Edinburgh
80 South Bridge
Edinburgh EH1 1HN, Scotland
Tel: +44 (131) 650 2755
Fax: +44 (131) 650 6513
B.Drabble@ed.ac.uk

Dr Yolanda Gil
Information Sciences Institute
University of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292
Tel: +1 310 822-1511
Fax: +1 310 823-6714
Gil@isi.edu

Prof Austin Tate
AI Applications Institute
University of Edinburgh
80 South Bridge
Edinburgh EH1 1HN, Scotland
Tel: +44 (131) 650 2732
Fax: +44 (131) 650 6513
A.Tate@ed.ac.uk

Keywords: Prototype, Application, Methods and Techniques

Abstract

1 Introduction

Generating sets of qualitatively different plans is crucial in the decision making support systems within any organisation or business. Current business and project planning tools are tasked such that all the alternative plans or Courses of Action (COAs) generated are produced under some fixed set of assumptions. For example, the COAs developed for an oil spill recovery task may include one that uses many resources but can be deployed fast, another that uses less resources and the deployment takes longer, another is somewhere in the middle and another is a bit more extreme. Generating qualitatively different plans would allow more variety and better quality solutions.

Our approach combines a generative planner with a knowledge-based system that reasons about plan evaluation. 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.

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:

[1] We would like to thank David Brown, 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.

- 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. Some planners provide facilities to control the quality of the solution to be returned, such as evaluation functions or search control rules. However, they do not usually integrate quality considerations across several plans. In addition, there is not enough data on the adequacy of these representations to reflect the plan quality criteria that are necessary in practice. Often, 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

What we see as the architecture for a system capable of generating qualitatively different plans is an outer "strategic/task assignment" layer of which performs some task analysis and sets direction. This would be used to set up definite targets and constraints for a "tactical" planner to develop further. The tactical planner would thus establish that a plan was possible within the constraints specified i.e. keeping various quality criteria at favourable levels. The tactical planner would be tasked with different requirements and constraints to produce alternative COAs which are qualitatively different. This may be simple to state but difficult to achieve in practice.

2 Example of the Use of Plan Quality Criteria

The sections describes two areas would could benefit from the use of the technology described in this paper.

2.1 The Oil Industry

The generation of activity plans of Courses of Action (COAs) is required within a number of areas of the oil and gas industry. Projects such as the construction of oil platforms, refineries, etc, the recovery plan for an identified field, oil spill recovery, exploration, etc, all need to be coordinated, controlled and monitored. However, in many cases there is no single solution for the task as a number of different factors, constraints and tradeoffs need to be taken into consideration. It is therefore important during the concept development phase of a project to develop careful estimates of the situation and alternative COAs. This leads to a situation in which:

- A broad spectrum of possible COAs is considered.
- The uncertainties in each of the COAs are analysed and estimated to reduce unknowns.
- The analysis can be used as the basis for the business decision and subsequent selection of the appropriate COA.

Taking the exploitation plan for a field as an example several possible COAs could be generated which have different assumptions and tradeoffs. For example, the simplest and cheapest COA may be to moor a tanker near a subsea template as a floating barge. Other COAs range from making use of an undersea pipeline which the company either owns, can lease capacity in or build through to othe commissioning of a platform to be built for the field. There are a number of criteria which could be used to analyse and choose between the different options in the plan and these include cost, technology, lead time to production, initial investment, operating costs, etc. These criteria will be specific to the problem domain and could be incrementally added to or removed as COA generation or analysis progresses. Choosing to generate a COA from each of these catargories would result in a course grain of solution.

Alternatively it becomes possible to identify potential spin off benefits. For example, a COA to build an undersea pipeline would result in higher start up costs but would significantly increase the companies technology base. This might prove useful when producing fields in new areas such as Vietnam.

2.2 Military 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 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 has been created based around a fictitious but realistic scenario called PRECiS, and is described in detail in [4, 3].

3 Approach

By having an intelligent planning system identify a number of candidate COAs it becomes possible for human planners to identify those parts of a COA they find acceptable and to task the planner to alter those parts of the COA which are not. For example, a COA to connect to an existing pipeline may be acceptable but the installation costs of the interface equipment should be reduced, i.e. choose a cheaper installation method. Any criteria which have the same value across all candidate COAs e.g all recovery plans have low startup costs can be ignored in the analysis because they do not become part of the justification for the business case.

Knowledge-based system technology enables us to build an interface between the planner and the user that provides the following functionality:

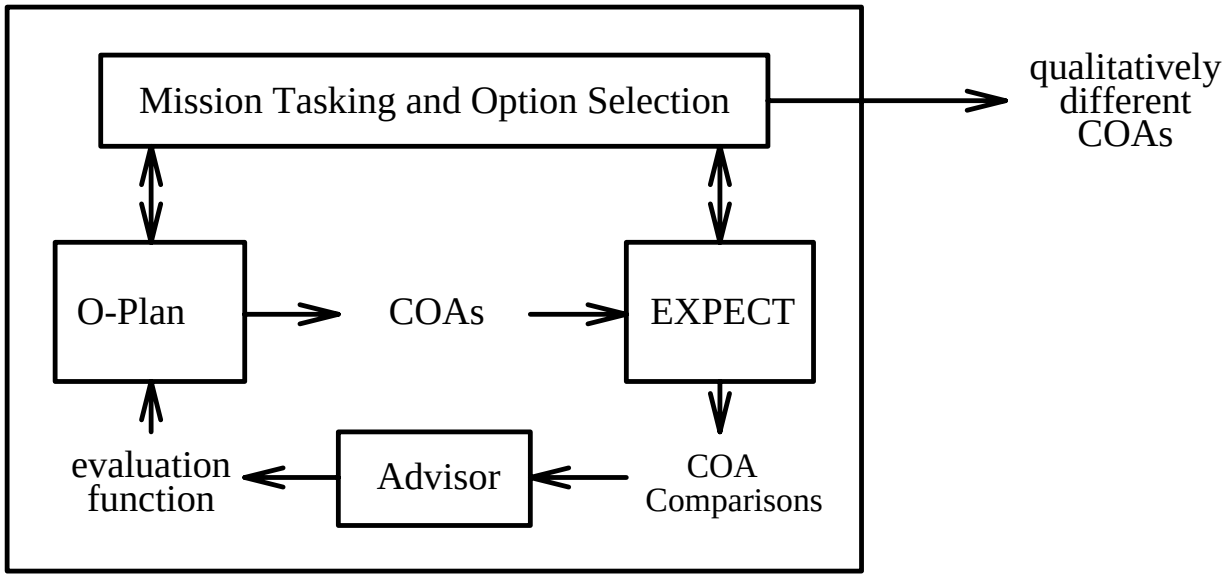

Figure 1: A generative planner and a knowledge-based system cooperate to produce better plans.

- support the user in defining criteria for evaluating plan quality
- 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 criteria system evaluates the solutions generated by the planner according to the criteria for plan quality that are currently known to it. In addition the criteria system interacts with a domain specialist i.e. a human planner and with the systems overall Tasking and Option Selection system to acquire further plan evaluation criteria that are relevant to the planning domain in general or to a specific problem instance. The domain specialist and the Tasking and Option Selection systems provide not only the criteria to be used in the evaluation, but also the procedural knowledge necessary to compute the evaluations automatically. The result of these interactions is an updated set of criteria that may include new criteria or more precise definitions of already existing ones. The function of the Advisor module is to take plan evaluations and their justifications and to analyse them together with the defined criteria to produce planner-independent advice that can be used to by the generative planner to guide its planning choices. The Advice Operationalization module integrates the Advisor 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.

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 has sought to identify modular components within an AI command, planning and control system and to provide clearly defined interfaces to these components and modules. The background to this work is provided in [?]. The various components plug into "sockets" within the architectural framework. The sockets are specialised to ease the integration of particular types of component. See figure 2.

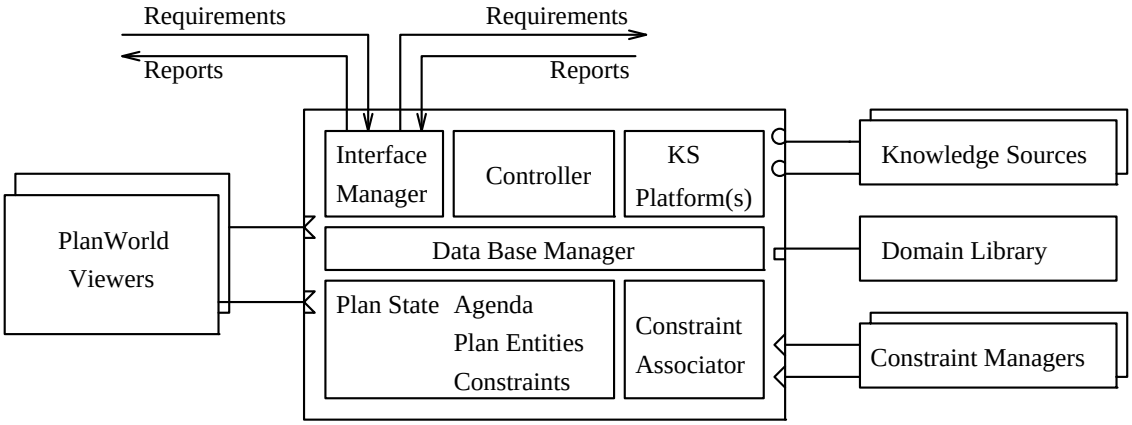

Figure 2: O-Plan Agent Architecture

The various components of the agent architecture are:

PlanWorld Viewers – User interface, visualisation and presentation viewers for the plan – usually differentiated into technical *plan* views (charts, structure diagrams, etc.) and *world* views (simulations, animations, etc.).

Knowledge Sources – Functional components which can analyse, synthesise or modify plans.

Domain Library – A description of the domain and a library of possible actions.

Constraint Managers – Support modules which manage detailed constraints within a plan and seek to maintain as accurate a picture as possible of the feasibility of the current plan state with respect to the domain.

These plug-in components are orchestrated by an O-Plan agent kernel which carries out the tasks assigned to it via appropriate use of the Knowledge Sources and manages options being maintained within the agent's *Plan State*. The central control flow is as follows:

Interface Manager – Handles external events (requirements or reports) and, if they can be processed by the agent, posts them on the agent *Agenda*.

Controller – Chooses Agenda entries for processing by suitable Knowledge Sources

Knowledge Source Platform(s) – Chosen Knowledge Sources are run on an available and suitable Knowledge Source Platform.

Data Base Manager – Maintains the Plan State being manipulated by the agent and provides services to the Interface Manager, Controller and Knowledge Sources running on KS Platforms to allow this.

Constraint Associator Acts as a mediator between the Plan State maintained by the data base manager and the various Constraint Managers that are installed in the agent. It eases the management of interrelationships between entities and detailed constraints.

EXPECT is an architecture for the development of knowledge-based systems that includes a runtime 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 COAs 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. EXPECT's evaluation prototype generates plan evaluations for different factors shown in Figure 4. One of the factors shown in the figure is the S-PORTS, i.e. the characteristics of the seaports used in the plan. There are several different aspects that need to be evaluated, including the number of seaports used in the plan with their piers and berths (which is important for loading and unloading ships), the petroleum, oil, and lubricants facilities available to refuel the ships, and the limitations imposed by the seaports on the size of the ships that can be used for the operation. The evaluation criteria are computed based on the features of the input plan. For example, to compute the limitations on ship size, EXPECT looks at the berth types in each seaport and finds the maximum length of the berths, which determines the length of the ships that the seaports can accomodate. Users can change the method to evaluate the criteria using EXPECT's knowledge acquisition tool. For example, the user can extend these criteria and ask the system to check the maximum draft of ships that the seaports accomadate based on their depth, as described in [2].

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. 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 3 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). This would be done at the Task Assignment level of the O-Plan architecture. Figure 4 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.

The evaluations of the alternative plans are useful to the domain experts to characterize the tradeoffs among alternatives. 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 COA 3 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 com-

POE/POD	WHO	WHEN	LIFT	PAX	Bulk	Over	Out	NAT
Andrews-afb/ Delta-Pacifica	107th-ACR	21	3 C-141	5492	1362	0	0	0
Andrews-afb/ Delta-Pacifica	107th-ACR	21	3 C-141	0	0	2670	2581	0
Wilmington-port/ Delta-port	107th-ACR	21	2 FSS	0	0	10678	10324	0
Ft-Meade-port/ Delta-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/ Delta-Pacifica	16th-F15E	25	4 C-141	794	199	0	0	0
Shaw-afb/ Delta-Pacifica	16th-F15E	25	4 C-141	0	0	59	0	0
Charleston-port/ Delta-port	16th-F15E	25	10 FSS	0	0	0	0	2553

Figure 3: 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 4: EXPECT's evaluation of several alternative plans generated by O-Plan.

mander. The tradeoff analysis often suggests parts of the space of alternatives that needs to be explored more. For example, a plan that closes earlier than COA 4 and uses a similar amount of personnel would be preferrable to any of the current four alternatives. One of our goals is to generate such alternative plans automatically. Further work is needed to understand how to describe the parts of the space of alternatives that the user would like a solution from, and how to use that description to task O-Plan appropriately so that it generates an alternative from that subspace. Nevertheless, the plans and evaluations of the kind that are currently produced by our system are valuable in themselves, since they automate part of the process and provide information that is useful for human planners in making these decisions.

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, and M. Zweben.), Morgan Kaufmann, 1994.

[7] Stillman, J. and Bonissone, P., *Developing New Technology for the ARPA-Rome Laboratory Planning Initiative*, IEEE Expert Intelligent Systems and their Applications, pp10-17, February, 1995.