Intelligible AI Planning - Generating Plans Represented as a Set of Constraints

Austin Tate
Artificial Intelligence Applications Institute
University of Edinburgh, UK

Abstract

[1]Realistic planning systems must allow users and computer systems to co-operate and work together using a "mixed initiative" style. Black box or fully automated solutions are not acceptable in many situations. Studies of expert human problem solvers in stressful or critical situations show that they share many of the problem solving methods employed by hirearchical planning methods studied in Artificial Intelligence. But powerful solvers and constraint reasoners can also be of great help in tparts of the planning process. A new more intelligible approach to using AI planning is needed which can use the best "open" styles of planning based on shared plan representations and hierarchical task networks (HTN) and which still allow the use of powerful constraint representations and solvers.

I-Plan is a design for a new planning system based on these principles. It is part of the I-X suite of intelligent tools. I-Plan is modular and can be extended via plug-ins of various types. It is intended to be a "lightweight" planning system which can be embedded in other applications. In its simplest form it can provide a small personal planning aid that can be deployed in portable devices and other user-orientated systems to add planning facilities into them. In its more developed forms it will approach the power of generative AI planners such as O-Plan. It provides a framework for including powerful constraint solvers in a framework that is intelligible to the users.

I-Plan is grounded in the <I-N-OVA>[2] (*Issues – Nodes – Orderings / Variables / Auxiliary*) constraints model used to represent plans and processes. <I-N-OVA> is intended to support a number of different uses:

- for automatic and mixed-initiative generation and manipulation of plans and to act as an ontology to underpin such use;
- as a common basis for human and system communication about plans;
- as a target for principled and reliable acquisition of plans, process models and process product information;
- to support formal reasoning about plans.

The I-Plan design and the <I-N-OVA> ontology provide an extensible framework for adding detailed constraint representations and reasoners into planners. These can be based on powerful automated methods. But this can be done in a context which provides overall human intelligibility.

[1]This paper is partly based on a technical note to the AAAI-2000 Workshop on Representational Issues for Real-World Planning Systems, AAAI-2000 (Tate, 2000).

[2]<I-N-OVA> is pronounced as in "Innovate".

1 Introduction

Planning is about much more than solving specifically stated problems as efficiently as possible. It is also about modelling domains in which planning takes place, understanding the roles of the various human and system agents involved in the planning process and in the domain in which plans are executed, and it is about communicating tasks, plans, intentions and effects between those agents. Realistic planning systems must allow users and computer systems to cooperate and work together using a "mixed initiative" style. Black box or fully automated solutions are not acceptable in many situations. Studies of expert human problem solvers in stressful or critical situations (Klein, 1998) show that they share many of the problem solving methods employed by some of the methods studied in AI planning to address these issues.

This paper argues that a Hierarchical Task Network (HTN) least commitment planning approach - as used for many years in practical planning systems such as NOAH (Sacerdoti, 1975), Nonlin (Tate, 1977), SIPE (Wilkins, 1988) and O-Plan (Currie and Tate, 1991) - provides an intelligible framework for mixed-initiative multi-agent human/system planning environments. When joined with a strong underlying constraint-based ontology of plans it can provide a framework in which powerful problem solvers based on search and constraint reasoning methods can be employed and still retain human intelligibility of the overall planning process and the plan products that are created.

I-Plan is a design for a new "lightweight" planning system based on these principles. It is part of the I-X[3] suite of intelligent tools and is being designed to be embedded in other applications. I-Plan is modular and can be extended via plug-ins of various types. In its simplest form it can provide a small planning aid that can be deployed in portable devices and other user-orientated systems to add planning facilities into them. In its more developed forms it will approach the power of major generative AI planners such as O-Plan (Tate et. al, 1994; Tate et. al., 2000).

2 I-X

Work in Intelligent Planning and Activity Management at the University of Edinburgh[4] has led to a number of planning systems and approaches that are re-used on a number of projects. New work will drawn on this work, generalise it, and significantly extend the application of the core concepts and assets, leading to new re-usable components, and create opportunities for applications and further research.

This new programme is called I-X and the core components are a shared model representation called <I-N-CA> and a systems integration architecture. A variety of re-usable components and systems will be built on the new architecture and these will be collectively referred to as I-Technology and I-Tools.

[3]I-X is the successor project to O-Plan - see http://www.aiai.ed.ac.uk/project/ix/.
[4]See http://www.aiai.ed.ac.uk/project/plan/.

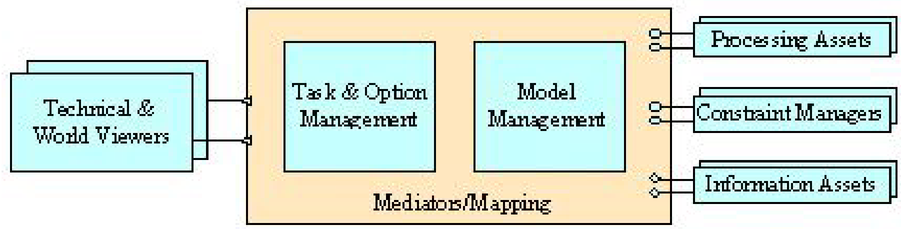

Figure 1: I-X Components

I-X provides a systems integration architecture. Its design is based on the O-Plan agent architecture. I-X incorporates components and interface specifications which account for simplifications, abstractions and clarifications in the O-Plan work. I-X provides an issue-handling workflow style of architecture, with reasoning and functional capabilities provided as plug-ins. Also via plug-ins it allows for sophisticated management and use of the internal model representations to reflect the application domain of the system being built in I-X. I-X agents may be recursively or fractally composed, and may interwork with other processing cells or architectures. This is a systems integration approach now being advocated by a number of groups concerned with large scale, long-lived, evolving and diverse systems integration issues.

The I-X approach has 5 aspects:

1. Systems Integration - A broad vision of an open architecture for the creation of intelligent systems for the synthesis of a result or "product" which is based on a "two cycle" approach which uses plug-in components to "handle issues" and to "manage and respect the domain model".

2. Representation - a core notion of the representation of a process or plan as a set of nodes making up the components of the process or plan model, along with constraints on the relationship between those nodes and a set of outstanding issues. This representation is termed <I-N-CA> - Issues, Nodes, Critical Constraints and Auxiliary Constraints.

3. Reasoning - the provision of reusable reasoning capabilities.

4. Viewers and User Interfaces - to understand user roles in performing activities and to provide generic modules which present the state of the process they are engaged in, their relationships to others and the status of the artifacts/products they are working with.

5. Applications - work in various application sectors which will seek to create generic approaches (I-Tools) for the various types of task in which users may engage. One important application is I-Plan for planning tasks.

We propose to bring together a number of threads of previous research and development, and use state-of-the-art understanding of the conceptual basis for flexible, incremental, mixed-initiative planning and activity management systems. We will incorporate these into an open, flexible, lightweight and embeddable system. This will be written in Java for portability and to maximise reuse potential. The core of the system will be an agenda-based issue handling system based on workflow principles. It will be specialised to any particular task by incorporating suitable issue-handling capabilities which could be supplied by human or system components. It will be designed to allow for very significant extension via an open capability plug-in interface and via an interface to allow for the use of constraint management methods, feasibility estimators, simulators, etc. The system will be able to inter-work with other workflow and cooperative working support systems, and will not make assumptions about the internal architecture of those other systems.

The components of the I-X systems integration architecture are shown diagrammatically in figure 1 and are as follows:

- Task and Option Management – The capability to support user tasks via appropriate use of the processing and information assets and to assist the user in managing options being used within the model.
- Model Management – coordination of the capabilities/assets to represent, store, retrieve, merge, translate, compare, correct, analyse, synthesise and modify models.
- Issue Handlers – Functional components (distinguished into those which can add to the model (synthesis) and those which analyse the model (to add information only).
- Constraint Managers – Components which assist in the maintenance of the consistency of the model.
- Information Assets – Information storage and retrieval components.
- Viewers – User interface, visualisation and presentation viewers for the model - sometimes differentiated into technical model views (charts, structure diagrams, etc.) and world model views (simulations, animations, etc.)
- Mediators – Intermediaries or converters between the features of the model and the interfaces of active components of the framework (such as viewers, processing assets, constraint managers and information assets).

A number of different types of "sockets" are available within the framework to reflect the protocols or interfaces into which the various components can fit. The necessity for specific sockets and the types of components vary across projects to some extent, but the separation into viewers, processing assets, constraint managers and information assets has been found to be useful in a number of AIAI projects. This also puts the I-X work on a convergent path with other Model/Viewer/Controller styles of systems framework.

3 <I-N-OVA> and <I-N-CA>

I-Plan is grounded in the <I-N-OVA> *Issues – Nodes – Auxiliary*) constraints model which is used to represent plans and processes. The more general <I-N-CA> (*Issues – Nodes – Critical/Auxiliary*) constraints model can be used for wider applications in design, configuration and other tasks which can be characterised as the synthesis and maintenance of an artifact or product.

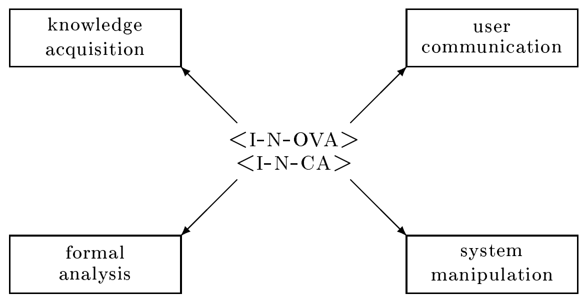

Figure 2: <I-N-OVA> and <I-N-CA> Support Various Requirements

As shown in figure 2, the <I-N-OVA> and <I-N-CA> constraint models are intended to support a number of different uses:

- for automatic and mixed-initiative generation and manipulation of plans and other synthesised artifacts and to act as an ontology to underpin such use;
- as a common basis for human and system communication about plans and other synthesised artifacts;
- as a target for principled and reliable acquisition of plans, process models and process product information;
- to support formal reasoning about plans and other synthesised artifacts.

These cover both formal and practical requirements and encompass the requirements for use by both human and computer-based planning and design systems.

The <I-N-OVA> (*Issues – Nodes – Orderings / Variables / Auxiliary*) Model is a means to represent plans and activity as a set of constraints. By having a clear description of the different components within a plan, the model allows for plans to be manipulated and used separately from the environments in which they are generated. The underlying thesis is that plans can be represented by a set of constraints on the behaviours possible in the domain being modelled and that plan communication can take place through the interchange of such constraint information.

<I-N-OVA>, when first designed (Tate, 1996), was intended to act as a bridge to improve dialogue between a number of communities working on formal planning theories, practical planning systems and systems engineering process management methodologies. It was intended to support new work then emerging on automatic manipulation of plans, human communication about plans, principled and reliable acquisition of plan information, and formal reasoning about plans. It has since been utilised as the basis for a number of research efforts, practical applications and emerging international standards for plan and process representations. For some of the history and relationships between earlier work in AI on plan representations, work from the process and design communities and the standards bodies, and the part that <I-N-OVA> played in this see Tate (1998).

In Tate (1996), the <I-N-OVA> model is used to characterise the plan representation used within O-Plan and is related to the plan refinement planning method used in O-Plan. The <I-N-OVA> work is related to emerging formal analyses of plans and planning. This synergy of practical and formal approaches can stretch the formal methods to cover realistic plan representations as needed for real problem solving, and can improve the analysis that is possible for practical planning systems.

We have generalised the <I-N-OVA> approach to design and configuration tasks with I, N, CA components - where C represents the "critical constraints" in any particular domain - much as certain O and V constraints do in a planning domain. We believe the approach is valid in design and synthesis tasks more generally - we consider planning to be a limited type of design activity. <I-N-CA> is used as an underlying ontology for the I-X project.

The <I-N-OVA> and <I-N-CA> work is intended to utilise a synergy of practical and formal approaches which are stretching the formal methods to cover realistic representations, as needed for real problem solving, and can improve the analysis that is possible for practical planning systems.

4 <I-N-OVA> - Representing Plans as a Set of Constraints on Behaviour

A plan is represented as a set of constraints which together limit the behaviour that is desired when the plan is executed. The set of constraints are of three principal types with a number of sub-types reflecting practical experience in a number of planning systems.

The node constraints (these are often of the form "include activity") in the <I-N-OVA> model set the space within which a plan may be further constrained. The I (issues) and OVA constraints restrict the plans within that space which are valid.

Planning is the taking of planning decisions (I) which select the activities to perform (N) which creates, modifies or uses the plan objects or products (V)

```
Plan Constraints
   I   - Issues (Implied Constraints)
   N   - Node Constraints (on Activities)
   OVA - Detailed Constraints
         O - Ordering Constraints
         V - Variable Constraints
         A - Auxiliary Constraints
             - Authority Constraints
             - Condition Constraints
             - Resource Constraints
             - Spatial Constraints
             - Miscellaneous Constraints
```

Figure 3: <I-N-OVA> Constraint Model of Activity

at the correct time (O) within the authority, resources and other constraints specified (A). The node constraints in the <I-N-OVA> model set the space within which a plan may be further constrained. The I (issues) and OVA constraints restrict the plans within that space which are valid. The Issues are the items on which selection of Plan Modification Operators is made in agenda based planners.

Others have recognised the special nature of the inclusion of activities into a plan compared to all the other constraints that may be described. Khambhampati and Srivastava (1996) differentiate Plan Modification operators into "progressive refinements" which can introduce new actions into the plan, and "non-progressive refinements" which just partitions the search space with existing sets of actions in the plan. They call the former genuine planning refinement operators, and think of the latter as providing the scheduling component.

If we consider the process of planning as a large constraint satisfaction task, we may try to model this as a Constraint Satisfaction Problem (CSP) represented by a set of variables to which we have to give a consistent assignment of values. In this case we can note that the addition of new nodes ("include activity" constraints in <I-N-OVA>) is the only constraint which can add variables dynamically to the CSP. The Issue (I) constraints may be separated into two kinds: those which may (directly or indirectly) add nodes to the plan and those which cannot. The I constraints which can lead to the inclusion of new nodes are of a different nature in the planning process to those which cannot.

Some ordering (temporal) and variable constraints are distinguished from all other constraints since these act as "critical" constraints, usually being involved in describing the others – such as in a resource constraint which will often refer to plan objects/variables and to relationships between time points or intervals.

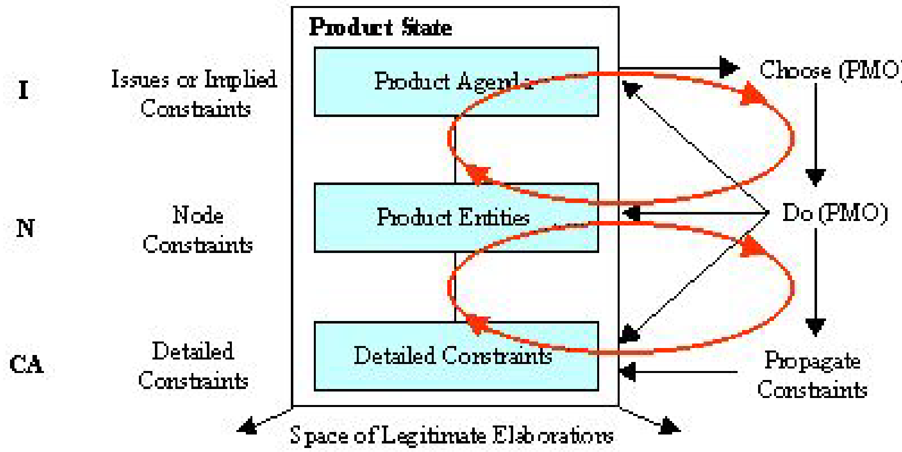

Figure 4: I-X and I-Plan Abstract Architecture: Two Cycles of Processing - Handle Issues, Respect Constraints. PMO=Product Modification Operator

5 I-Plan Abstract Design

The I-Plan design is based on two cycles of processing. The first addresses one or more "issues" from a task agenda, and the second ensures that constraints in the domain in which processing takes place is respected. So the processing cycles can be characterised as "handle issues, respect constraints". The emerging partial plan or schedule is analysed to produce a further list of issues or agenda entries. A choice of the issues to address is used to drive a workflow-style processing cycle of choosing 'Plan Modification Operators" and then executing them to modify the emerging plan state. Figure 4 shows this graphically for the more general case of designing or synthesising any product - where the issue handlers are labelled "PMO" - which then stands for the "Product Modification Operator".

This approach is taken in systems like O-Plan, OPIS (Smith, 1994), DIPART (Pollack, 1994), TOSCA (Beck, 1994), etc. The approach fits well with the concept of treating plans as a set of constraints which can be refined as planning progresses. Some such systems can also act in a non-monotonic fashion by relaxing constraints in certain ways.

Having the implied constraints or "agenda" as a formal part of the plan provides an ability to separate the plan that is being generated or manipulated from the planning system and process itself and this is used as a core part of the I-Plan design.

Mixed Initiative Planning approaches, for example in O-Plan (Tate, 1994), improve the coordination of planning with user interaction by employing a clearer shared model of the plan as a set of constraints at various levels that can be jointly and explicitly discussed between and manipulated by user or system in a cooperative fashion. I-Plan will adopt this approach.

6 Summary

The overall architecture of I-Plan has been described along with the <I-N-OVA> Constraint Model of Activity and the more general <I-N-CA> Constraint Model for Synthesised Artifacts. These are designed to draw on strengths from a number of different communities: the AI planning community with both its theoretical and practical system building interests; the issue-based design community, those interested in formal ontologies for processes and products; the standards community; those concerned with new opportunities in task achieving agents on the world wide web; etc.

<I-N-OVA> is intended to act as a bridge to improve dialogue between the communities working in these areas and potentially to support work on automatic manipulation of plans, human communication about plans, principled and reliable acquisition of plan information, and formal reasoning about plans. <I-N-CA> is designed as a more general underlying ontology which can be at the heart of a flexible and extensible systems integration architecture involving human and system agents.

The I-Plan planner and <I-N-OVA> ontology together provide an extensible framework for adding detailed constraint representations and reasoners which themselves can be based on powerful automated methods. But this can be done in a context which provides human intelligibility of the overall planning process[5].

Acknowledgements

The O-Plan and I-X projects are sponsored by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory Command and Control Directorate under grant number F30602-99-1-0024 ad the UK Defence Evaluation Research Agency (DERA). The U.S. Government, DERA and the University of Edinburgh are authorised to reproduce and distribute reprints for their purposes notwithstanding any copyright annotation hereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either express or implied, of DARPA, the Air Force Research Laboratory, the U.S. Government, DERA or the University of Edinburgh.

[5] The similarity of the AI planning techniques which can be employed within this framework to those observed in expert human problem solving in crisis situations (Klein, 1998) is described in the appendix.

References

Beck, H. (1993) TOSCA: A Novel Approach to the Management of Job-shop Scheduling Constraints, Realising CIM's Industrial Potential: Proceedings of the Ninth CIM-Europe Annual Conference, pages 138-149, (eds. Kooij, C., MacConaill, P.A., and Bastos, J.).

Currie, K.W. and Tate, A. (1991) O-Plan: the Open Planning Architecture, *Artificial Intelligence* 52(1), Autumn 1991, North-Holland.

Khambhampati, S. and Srivastava, B. (1996) Unifying Classical Planning Approaches, Arizona State University ASU CSE TR 96-006, July 1996.

Klein, G. (1998) Sources of Power - How People Make Decisions, MIT Press, 1998.

Pollack, M. (1994) DIPART Architecture, Technical Report, Department of Computer Science, University of Pittsburgh, PA 15213, USA.

Polyak, S. and Tate, A. (2000) A Common Process Ontology for Process-Centred Organisations, Knowledge Based Systems. Earlier version published as University of Edinburgh Department of Artificial Intelligence Research paper 930, 1998.

Sacerdoti, E. (1977) A structure for plans and behaviours. Artificial Intelligence series, publ. North Holland.

Smith, S. (1994) OPIS: A Methodology and Architecture for Reactive Scheduling, in Intelligent Scheduling, (eds, Zweben, M. and Fox, M.S.), Morgan Kaufmann, Palo Alto, CA., USA,

Tate, A. (1994) Mixed Initiative Planning in O-Plan2, Proceedings of the ARPA/Rome Laboratory Planning Initiative Workshop, (ed. Burstein, M.), Tucson, Arizona, USA, Morgan Kaufmann, Palo Alto.

Tate, A. (ed.) (1996a) Advanced Planning Technology - Technological Achievements of the ARPA/Rome Laboratory Planning Initiative (ARPI), AAAI Press.

Tate, A. (1996b) Representing Plans as a Set of Constraints – the <I-N-OVA> Model, Proceedings of the Third International Conference on Artificial Intelligence Planning Systems (AIPS-96), pp. 221-228, (Drabble, B., ed.) Edinburgh, Scotland, AAAI Press.

Tate, A. (1998) Roots of SPAR - Shared Planning and Activity Representation, Knowledge Engineering Review, Vol. 13, No. 1, March 1998. See also http://www.aiai.ed.ac.uk/project/spar/

Tate, A. (2000) <I-N-OVA> and <I-N-CA> - Representing Plans and other Synthesized Artifacts as a Set of Constraints, AAAI-2000 Workshop on Representational Issues for Real-World Planning Systems, at the National Conference of the American Association of Artificial Intelligence (AAAI-2000), Austin, Texas, USA, August 2000.

Tate, A., Drabble, B. and Kirby, R. (1994) O-Plan2: an Open Architecture for

Command, Planning and Control, in Intelligent Scheduling, (eds, Zweben, M. and Fox, M.S.), Morgan Kaufmann, Palo Alto, CA., USA.

Tate, A., Drabble, B. and Dalton, J. (1994) Reasoning with Constraints within O-Plan2, Proceedings of the ARPA/Rome Laboratory Planning Initiative Workshop, (ed. Burstein, M.), Tucson, Arizona, USA, Morgan Kaufmann, Palo Alto, CA, USA.

Tate, A., Levine, J., Jarvis, P. and Dalton, J. (2000) Using AI Planning techniques for Army Small Unit Operations, Poster Paper in the Proceedings of the Fifth International Conference on AI Planning and Scheduling Systems (AIPS-2000), Breckenridge, CO, USA, April 2000.

Wilkins, D. (1988) *Practical Planning*, Morgan Kaufmann, Palo Alto, 1988.

Appendix: Comparing the Intelligible Planning Approach to Studies of Expert Human Planners

This appendix describes some of the features of the O-Plan and I-Plan approaches and shows the similarity of these approaches with those observed in expert human problem solvers performing in stressful or unusual situations. These observations were made in studies over many years by Klein (1998) and he contrasts these with some automated "black box" AI and algorithmic techniques.

The following note was produced on the DARPA O-Plan Project for an US Army Small Unit Operations Application (Tate et. al., 2000) in June 1999 by Austin Tate

But I Don't Plan, I Just Know What to Do

There are different types of planning technology available from the AI community. This is not restricted to a simple kind of search from some known initial state to some final desired state seeking the best solution according to some predefined criteria. Gary Klein's book (Klein, 1998) on how people make decisions in situations such as military operations, fire fighting, or other life threatening environments provides a rich set of case studies to show that in relatively few situations were deliberative planning techniques in obvious use. People just seemed to be making the "right" choices - or a choice that worked which was all that was required. They attributed their rapid selection of a suitable course of action to training, experience, or even ESP! Where options were deliberated over and evaluated, the situation for those involved was novel or unusual to their previous experience.

Klein's studies show how people in stressful environments select a course of action and adapt it as circumstances alter. Many of the decisions made by the subjects relate to issues which AI planning researchers are addressing. However, they are far removed from the traditional search style of deliberative plan generation. So we need to establish for the outset that the techniques we are calling upon to address potential planning requirements also are much wider than these simple fully-automated search methods. We are seeking to use rich plan representations in a variety of ways. These are listed below, along with cross references to Klein's book, to show how we can address a variety of decision methods which he is advocating, and which are in use by real problem solvers and commanders . The hope is that the planning requirements we are identifying can be mapped to some of the AI concepts we are bringing to bear on practical planning problems.

- Overall management of the command, planning and control process steps to improve coordination.
- Expansion of a high level abstract plan into greater detail where necessary.

- High level "chunks" of procedural knowledge (Standard Operating Procedures, Best Practice Processes, Tactics Techniques and Procedures, etc.) at a human scale - typically 5-8 actions - can be manipulated within the system [Klein, p. 52 and p. 58].

- Ability to establish that a feasible plan exists, perhaps for a range of assumptions about the situation, while retaining a high level overview. [Klein, p.227, "Include only the detail necessary to establish a plan is possible - do not fall into the trap of choreographing each of their movements"].

- Analysis of potential interactions as plans are expanded or developed [Klein, p 53].

- Identification of problems, flaws and issues with the plan [Klein p. 63 and p. 71].

- Deliberative establishment of a space of alternative options perhaps based on different assumptions about the situation involved of especial use ahead of time, in training and rehearsal, and to those unfamiliar with the situation or utilising novel equipment [Klein p. 23].

- Monitoring of the execution of events as they are expected to happen within the plan, watching for deviations that indicate a necessity to re-plan (often ahead of this becoming a serious problem) [Klein p. 32-33].

- AI planning techniques represent the dynamic state of the world at points in the plan and can be used for "mental simulation" of the execution of the plan [Klein, p. 45].

- Pruning of choices according to given requirements or constraints [Klein, p. 94 "singular strategy"].

- Situation dependent option filtering (sometime reducing the choices normally open to one "obvious" one [Klein p.17-18].

- Satisficing search to find the first suitable plan that meets the essential criteria [Klein p. 20].

- Anytime algorithms which seek to improve on the best previous solution if time permits.

- Heuristic evaluation and prioritisation of multiple possible choices within the constraint search space [Klein, p. 94].

- Repair of plans while respecting plan structure and intentions.

- Uniform use of a common plan representation with embedded rationale to improve plan quality, shared understanding, etc. [Klein, p. 275 7 types of information in a plan].

Gary Klein was asked to comment upon this review of AI techniques as compared to his observations of natural problem solving and decision making in humans. He observed the following in this edited Personal Communication to Austin Tate on 24-Jun-1999 (quoted with permission):

1. I felt a strong kinship with what you are attempting. The effort to use satisficing criteria, the use of anytime algorithms to permit continual improvement, the shift from abstract to detailed plan when necessary, the analysis of interactions in a plan, the identification of flaws in a plan, the monitoring of execution, the use of mental simulation, the representation of a singular strategy, heuristic evaluation, plan repair, and so forth are all consistent with what I think needs to be done.

2. My primary concern is how you are going to do these things.... The discipline of AI can provide constraints that will help you understand any of these strategies in richer detail. But those constraints may also prevent you from harnessing these sources of power.

3. Your slogan "Search and you're dead" seems right. Unconstrained search is a mark of intellectual cowardice. And it is also not a useful strategy.

Edited version of Personal Communication from Austin Tate to Gary Klein on 25-Jun-1999:

I want to clarify my use of the slogan "Search and you're dead" over the last 20 years. This is the headline, but I then clarify what I mean as "(Unconstrained) search and you're dead".

I have found this to be a useful slogan to express my general approach, and it makes for good knock about fun on panels at conferences. The idea should be to richly describe the constraints known using whatever knowledge is available about the problem, and then to seek solutions in that constrained space. We seek to use knowledge of the domain to constrain the use of blind search or "black box" automated methods in ways which are intelligent and intelligible (to humans).

In reality all planning systems we build have sophisticated search and constraint management components, and it is an aim of our research to be able to utilise the best available in an appropriate context. Search can be a useful tactic in situations where you are underconstrained and stuck. AI has made enormous advances in constraint management using search and other methods over the last 5 years - so much so that some of its proponents argue that we do not need to bother with domain expertise or being knowledge-based about many of the problems we are addressing. It's this latter overenthusiasm for one approach which I seek to counter. Even very powerful search can be made more useful if put into a sensible knowledge-based context. This is, of course, more relevant when humans are involved in the decisions as then a more naturalistic style of mutually progressing towards a solution become a key to successful use of the technology.