Profitable Directions for AI-Planning Research: A Personal View

Peter Jarvis

Artificial Intelligence Center
SRI International
333 Ravenswood Ave, Menlo Park, California 94025, USA.
Jarvis@ai.sri.com

Abstract

I reflect on the progress that the AI Planning field has made over the past 30 years and define the directions where I believe that we should focus our future efforts if we are to continue as a successful and vibrant scientific field. More concretely, I argue that Fikes and Nilsson's original framing of the planning problem is too far removed from the requirements of real-world applications. We must accept that it will not be feasible to obtain complete and consistent domain theories in the near future and that users will want to influence the plans they receive in dimensions other than goal state and shortest path. In light of this, I argue that we must focus on computer-aided planning instead of computer-replaced planning. I outline the technical challenges this course offers and the exciting work already emerging.

Keywords: Computer Aided Planning, and STRIPS Assumption.

Introduction

> "There are three principal means of acquiring knowledge available to us: observation of nature, reflection, and experimentation."
>
> **Denis Diderot** (1713–84)

All fields must periodically reflect upon their achievements, reexamine their goals, and direct future efforts accordingly. Here, I outline my personal reflection on the progress made in the AI planning field over the past 30 years and identify where I believe we should focus our future effort. I highlight some promising directions that have emerged before setting out some fundamental changes in research program structure and performance evaluation that we must bring about if these new directions are to mature and enable the field to move to a new level of accomplishment.

My Background

It is important for me to set out my experience in the field before proceeding. This is not an attempt to proclaim that I am a great oracle with far-reaching insights; rather, it is important that you understand the path I have taken so that you may judge appropriately my experience and bias.

During my Ph.D. program I focused on applying planning techniques to civil engineering projects (Jarvis & Winstanley 1996; Bloomfield et al. 1999) before joining the O-Plan team in Edinburgh. At Edinburgh, I worked on a range of applications, including process management for the chemical industry (Jarvis et al. 2000), army small unit tactical planning (Tate et al. 2000), and strategic military planning (Dyke et al. 2000). Since moving to SRI International, I have worked on air campaign (Myers et al. 2001) and Special Forces problems (Myers, Jarvis & Lee 2001; 2002). I have also worked on incorporating concepts from fuzzy logic into graph and SAT-based planners (Jarvis, Miguel & Shen 2000; Miguel, Jarvis & Shen 2000). To summarize, over the last 8 years I have experienced four large DARPA programs, several commercial consultancy projects, and a handful of Ph.D. programs. My goal here is to add what I have learned from this path to the debate on the directions that we should take in the future.

What We Have Achieved

AI researchers have worked primarily with Fikes and Nilsson's (1971) original framing of the planning problem for just over 30 years now. Quoting directly from their paper:

> *...the problem space for STRIPS* [or your planner] *is defined by three entities:*
>
> *(1) An Initial world mode, which is a set of wffs describing the present state of the world*
>
> *(2) A set of operators, including a description of their effects and their precondition wff schemata*
>
> *(3) A goal condition stated as a wff*
>
> *The problem is solved when STRIPS* [or your planner] *produces a world model that satisfies the goal wff.*

The publication mass of the planning field is dominated by two thrusts:

A Minimize the time a planner takes to find the shortest operator sequence necessary to produce a world model that satisfies the goal wff.

B Without damage to progress on thrust A, remove the (explicit) simplifications made by Fikes and Nilsson: atomic time, deterministic action effects, omniscience, and the planning agent being the sole cause of change.

We have made extraordinary progress over the past 30 years. Today's planners can solve problems orders of magnitude more complex than those of a few years ago (read Weld 1994, then Weld 1999 for an excellent perspective of how the field changed in those 5 years). We also have the first fully fielded and well-documented applications (Muscettola et al. 1998; [a subset of those in] Knoblock ed. 1996).

While we have a right to be proud of our accomplishments, we must continue to move forward with new and exciting innovations if the field is to remain alive and funded. I now consider the question of what is missing from our portfolio.

What We Have Not Achieved

A scientific field must produce concepts that can be taken on by engineers to produce artifacts of value to society. While it should not be a field's only focus, this link to application is important, as it provides many fascinating intellectual challenges that help keep scientists grounded and funding agencies interested.

As I noted above, our field has recently produced well-documented applications. This landmark accomplishment is rightly being celebrated in the literature and at conferences. My concern, however, is that people planning military operations still use paper or Microsoft PowerPoint™ while civil engineers use Artimis™ type tools. These tools offer benefits in facilitating the documentation and communication of plans, but they provide no assistance in the decision-making, specifically deciding what actions must be included in the plan or detecting and resolving complex action interactions.

The question I ask is, will the current focus of our field produce applications that scale to support the type of large-scale planning tasks I mention above? It is clear that we have solutions offering much impact at the physical device level. Taking a military example, we might soon be able to comfortably control a tank in achieving goals such as "stay alive" and "engage the enemy. " In the next section, I argue that the field is not tackling with sufficient emphasis the requirements of large-scale planning problems. Returning to the military example, I do not see the field's current direction leading to a technology capable of supporting a general's staff officers in designing the high-level strategy of a battle.

Where AI Planning Must Focus Its Efforts

I group my opinions on this topic into two sections. In the first, I consider promising work that seeks to address the implicit STRIPS assumptions and encourage more work in these areas. I then turn my attention to the more pragmatic issues of changes in research program structure and research evaluation necessary to encourage a broader range of work under the planning banner.

Relaxing the Implicit STRIPS Assumptions

While Fikes and Nilsson's explicit simplifying assumptions (often referred to as the STRIPS assumptions) have been the focus of much effort, little research has been devoted to the implicit assumptions in their framing of the planning problem. My thesis is that the Fikes and Nilsson's casting is too restrictive to be of value outside of device-level application domains. Here I work through the implicit STRIPS assumptions and promising work in each area. Only by bringing more effort to bear in these areas will the AI planning community reach the broad range of applications that could benefit from tool support.

Complete Operator Sets

Fikes and Nilsson's assumption that we can build an operator set that completely covers a domain is proving most difficult to realize in practice. In applications with many degrees of freedom, it is impractical to expect full coverage.

Consider the military problem of evacuating U.S. citizens from a hostile country. The number of factors that must be taken into account is simply enormous. People charged with planning operations of this type need to bring around 20 years of domain experience to bear on the problem. The AI community has long been trying to encode knowledge in these quantities, but with limited payoff (Leant & Guha 1990). Why does the AI planning community believe it can do better than its colleagues?

Consider planning's sister activity, design. Much design is now computer aided, as researchers have sought to complement rather than replace human designers. Encoding the knowledge needed to design products is several orders of magnitude more complex than encoding that required in assisting a human. Why does the AI Planning community insist on focusing solely on automated planning when computer-aided planning could provide tools of social value in a much closer time scale? Should we balance our portfolio to provide near as well as long-term payback?

Exciting work is emerging in the computer-aided planning direction. Dyer's SOFTools (GDATS 2002) provides one extreme on the continuum of computer-aided planning tools that is already in operational use with U.S. Special Forces units around the world. SOFTools provides a simple temporal planning interface with domain-specific icons. This speeds planning from the users' perspective, as it is more focused upon the types of diagrams with which they represent plans while also providing a more structured representation for researchers to exploit than alternative documentation aids such as Microsoft PowerPoint (their previous tool of choice).

SRI's CODA system (Myers, Jarvis & Lee 2001; 2002) integrates with SOFTools to provide a higher level of computer support. With CODA, users can describe aspects of a plan where changes are likely to affect them adversely. CODA automatically generates alerts if another user changes such an area. This helps human planners coordinate when distributed in both time and space.

The mixed initiative paradigm is well suited to computer aided planning. Ferguson, Allen, and Miller's TRIPS system (1996) hooks a person and a planner together to solve travel problems, where the computer's role is to maintain constraints and inform the users when they are likely to be violated.

This direction offers significant challenges. For example, it necessitates that a user be allowed to add structures to a plan. With such editing allowed, it is no longer possible to check automatically that a plan is correct (all required causal links in place and unthreatened, for example) as the user may have neglected to input some important precondition or effect. What level of consistency checking is appropriate and possible in a plan authoring context is an interesting research question.

Goal Specification

Applications demand the specification of more than just the goals a plan is to achieve. Users may want to specify the strategies to use in solving the problem (avoid using F-14s for combat air patrols or stay in first class accommodation on the business legs of a trip) or even some of the actions that must be included in the solution (fly United between SFO and LAX). Myers has explored both types of user guidance in the form of Advisable Planners (1996) and Plan Sketch Completion (1997).

While Myers' work is an important step toward providing comprehensive mechanisms for specifying user objectives, the work assumes a complete operator base that as we argue above, is not likely to be available in practice. Some significant and interesting research challenges are still to be addressed in this area.

Plan Evaluation

We have assumed that the user of a planning system is looking for a single plan and that this should be the shortest plan. In practice, people plan for many reasons. In military planning (when, as the old adage suggests, no plan survives first contact with the enemy) planning is often used to ensure that the course of action committed to is readily adaptable to a changing situation. Planning in this context is an exploratory task where multiple plans are produced under different assumptions or advice directives (use F14s for Combat Air Patrols (CAPS), don't use F14s for CAPS) and compared.

Myers and Lee (1999) have considered this problem of generating multiple plans automatically. Again, this work assumes that complete operator sets are available. Swartout and Gil (1996) in their INSPECT system provide support for evaluating a course of action against user-defined criteria.

To move forward we must consider in more detail the need to explore multiple plans, perhaps even combining parts to form a new option. Open questions remain at many levels. How can we efficiently reason with large-scale plans that contain contingency branches? What are the salient features that users use to choose between courses of action? How can we present those features to users so that they can rapidly compare plans?

Flexible Preconditions

In mainstream planning, all an operator's preconditions must be satisfied for it to be applicable. This framework does not support the flexibility in constraint satisfaction necessary in many applications.

Consider the military problem of infiltrating a small team by swimming from a submarine to a beach. The standard operating procedure for this task contains many constraints. The submarine must remain concealed under the ocean surface (minimum operating depth and maximum illumination from the moon), the swimmer must avoid hypothermia (function of the distance to swim and sea temperature), and the infiltration must be completed within the time frame demanded by the overall mission of which it is a component.

It is rare that one can find the right combination of ocean temperatures, tide, and lunar illumination within the timescale of an operation of this type. Typically, something has to be compromised. For example, asking the swimmers to cover a greater distance keeps the submarine concealed while reducing the effectiveness of the divers when they reach their target because of the additional fatigue.

While there has been much work on looking at the probability of mission success (send two teams of divers rather than one) there has been little considering the effect of constraint violation on plan quality. Miguel, Jarvis, and Shen (2001) consider this problem. The approach is preliminary and asks more questions that it answers. In particular, we take a fuzzy-logic based approach to reasoning about the damage a violated constraint inflicts on a solution. Is this the right way to represent the importance of constraints?

Necessary Environmental Changes

Here, I examine the broader environmental issues that must be considered if we, as a field, are to produce more readily applicable technology.

Multidisciplinary Research

The problem with AI Planning research is that AI Planning researchers have undertaken it! We are focused on algorithms, as that is what interests us most. What we are not generally interested in is modeling how people go about solving planning problems and identifying the niches for tool support. We have rather assumed that in solving Fikes and Nilsson's categorization of the problem we will produce the tool support that people require.

There are two barriers to carrying out the multidisciplinary research that I am suggesting. First, we need to build relationships with people with expertise in human factors, systems analysis, and cognitive psychology as these are the types of people who can help us understand the "as is" situation with human-level planning. We can then work on generating the "to be" process and finally the tool support needed for it. Second, we have not sought such multidisciplinary funding. This might be a function of the compartmentalization of funding agencies or just the comfort of working with colleagues with similar expertise.

Planner Evaluation and Planning Competitions

One of the attractions of Fikes and Nilsson's characterization is the ease with which progress can be measured. If Planner A solves problem 1 faster than Planner B, then we can conclude that Planner A's performance is superior. This ease of comparison has driven the community towards its current focus on planning competition where the group walking away with the most prestige from a conference is likely to be the one that provides the fastest system in the competition track.

Computer aided planning systems are going to be more difficult to evaluate. However, we must find appropriate metrics or will be difficult to determine if progress has been made. Again, we reach a difficult research question. How do we evaluate applied research? Who will pay for the evaluation effort given that it might need access to many users in controlled conditions?

Good Application Papers

A justifiable criticism raised at much of the previous applied research is that it is not well documented. Authors have not always clearly laid out the computational procedures and domain models that they have used. This has made it difficult to determine exactly what an "application" is doing and the compromises that have been made in its design.

While omitting this detail is understandable given funding constraints, it is necessary if the applied community is to maintain the respect of the more formal community. We should be careful to take the time to make our applications assessable by our more formally focused colleagues.

End the Divide on Search Control Knowledge

The planning field has been divided into the mutually exclusive "formal" and "applied" camps for too long. The applied camp has centered on Hierarchical Task Network (HTN) (Sacerdoti 1974; Tate 1977) techniques that encode knowledge about the actions available in a domain together with knowledge about how to go about solving problems in that domain. The "formal" camp has resisted the encoding of this search control knowledge as its members correctly argue that it leads to less flexible solutions.

There are two points here. First, applied work has been discounted because it has almost always has made use of search control knowledge. However, the problems posed by real applications exist independently of this design decision. Indeed, I have not had to mention this design decision until now. Second, HTN approaches couple search control knowledge tightly with operators. Huang, Selman and Kautz (1999) show that search control knowledge can be loosely coupled with the operator base, allowing it to be swapped in and out more easily.

We should proceed with the understanding that search control knowledge should be avoided. However, when we have to use it we must ensure that it is declarative so that it may be replaced as search speeds increase or the control knowledge becomes outdated.

Conclusion

As a field, we must move beyond Fikes and Nilsson's characterization of the planning problem and center our efforts upon a computer-aided rather than computer-replaced planning process.

There is no shortage of challenging research questions to answer on this path. My closing question is, how do we motivate a field to move in a new direction given that it will cause significant discomfort in the short term?

References

Bloomfield, D., Faraj, I., Jarvis, P., and Anumba, C., 1999, Managing and Exploiting Knowledge Assets in the Construction Industry, *In Proceedings of the 8th International Conference on Durability of Building Materials and Components, Vancouver, Canada.*

Dyke, D., Salt, M., Jarvis, P., and Desimone, R., 2000, Experimental Results from Integrating Planning Systems and Simulation Models. *In Proceedings of the 2000 Command and Control Research Technology Symposium, Vienna, VA.*

Ferguson, G., Allen, J., and Miller, B., TRAINS-95: Towards a Mixed-Initiative Planning Assistant. in *Proceedings of the Third Conference on Artificial Intelligence Planning Systems (AIPS-96), Edinburgh, UK, 29-31 May 1996, pp. 70-77.*

Fikes, R., and Nilsson, N., 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence* 5(2). North Holland Publishing Company.

GDATS, 2002. SOFTools V2.0 User Guide. General Dynamics Corporation, http://www.gdats.com.

Huang, Y., Selman, B., and Kautz, H., 1999. Control Knowledge in Planning: Benefits and Tradeoffs. *Proc. AAAI-99*, Orlando, FL.

Jarvis, P., Miguel, I, and Shen, Q., 2000, Flexible Blackbox *In Proceedings of the Workshop on Representational Issues for Real-World Planning Systems held within AAAI-00, Austin, TX.*

Jarvis, P., Moore, J., Stader, J., Macintosh, A., and Chung, P., 2000, Harnessing AI Technologies to Meet the Requirements of Adaptive Workflow Systems, In J. Filipe (ed) *Enterprise Information Systems,* Kluwer Academic Publishers, pp173-180.

Jarvis, P., and Winstanley, G., 1996, Dynamically Assessed and Reasoned Task (DART) Networks. *In Proceedings of Expert Systems 1996, the 16th Annual Technical Conference of the British Computer Society Specialist Group on Expert Systems*, Cambridge, UK, December 1996, pp92-105. ISBN- 1-899621-15-6.

Knoblock, C., (editor), 1996, AI planning systems in the real world. *IEEE Expert,* December, p 4 – 12.

Leant, D., and Guha, R., 1990, *Building Large Knowledge Based Systems*. Addison Wesley

Miguel, I., Jarvis, P., and Shen, Q., 2001, Efficient Flexible Planning via Dynamic Flexible Constraint Satisfaction. *Engineering Applications of Artificial Intelligence*, 14, pp301-327.

Muscettola, N., Nayak, P., P, Pell, P, and William, B, 1998., Remote Agent: to boldly go where no ai system has gone before. *Artificial Intelligence* 103(1-2) 5-48

Myers, K., 1997., Abductive Completion of Plan Sketches, In Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-97).

Myers, K., 1996., Strategic Advice for Hierarchical Planners., In Principles of Knowledge Principles of Knowledge Representation and Reasoning: Proceedings of the Fifth International Conference (KR '96), Morgan Kaufmann Publishers, San Francisco, CA.

Myers, K., Jarvis, P., and Lee, T., 2002, CODA: Coordination of Distributed Human Planners. *In Proceedings of the 6th International Conference on Artificial Intelligence Planning and Scheduling Systems,* France.

Myers, K., Jarvis, P., and Lee, T., 2001, Active Coordination of Distributed Human Planners. *In Proceedings of the 6th European Conference on Planning (ECP-01)*, Toledo, Spain.

Myers, K., Smith, S., Hildum, D., Jarvis, P., and de Lacaze, R., 2001, Integrating Planning and Scheduling through Adaptation of Resource Intensity Estimates. *In Proceedings of the 6th European Conference on Planning (ECP-01)*, Toledo, Spain.

Myers, K., and Lee, T., 1999, Generating Qualitatively Different Plans through Metatheoretic Biases. in *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI-99)*, AAAI Press, Menlo Park, CA, 1999.

Tate, A., Levine, J., Jarvis, P., and Dalton, J., 2000, Using AI Planning Technology for Army Small Unit Operations. *In Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling Systems, Colorado, USA, April 2000.*

Tate, A., 1977, Generating Project Networks, IJCAI, pp 888-893.

Sacerdoti., E., 1974, Planning in a Hierarchy of Abstraction Spaces. *Artificial Intelligence,* 5 pp115-135.

Swartout, W., and Gil, Y., 1996, EXPECT: A User-Centered Environment for the Development and Adaptation of Knowledge-Based Planning Aids. In *Advanced Planning Technology: Technological Achievements of the ARPA/Rome Laboratory Planning Initiative*,. Menlo Park, Calif.: AAAI Press, 1996.

Weld, D., 1999, Recent Advances in AI Planning. AI Magazine. 20(2), pages 93-123.

Weld, D., 1994 An Introduction to Least Commitment Planning. AI Magazine, 15(4), pages 27-61.