

Feasibility Criteria for Investigating Potential Application Areas of AI Planning

T. L. McCluskey

School of Computing and Engineering
The University of Huddersfield, Huddersfield HD1 3DH, UK
lee@hud.ac.uk

Abstract

In this paper we address the problem of deciding whether it is feasible to apply AI planning technology (involving currently available planning engines) to an application area. We develop some criteria based on motivation, technological infrastructure and knowledge engineering aspects of an application, and we go on to apply these criteria to two application areas. The criteria both help to evaluate the overall feasibility, and in cases where development continues, help us to focus on the parts of the application which are likely to be most troublesome.

Introduction

In recent AI conferences (ICAPS, ECAI) there have been a number of workshops devoted to AI planning applications, and ICAPS itself gives an award to 'best application' paper. While many applications tend to be in AI-rich environments such as Space Technology, there is a growing body of applications from a wider range of areas. A notable example is the SIADEx (Fdez-Olivares et al. 2006) project, developing tools for helping people to manage forest fire fighting resources. Several other notable applications were described in the recent ICAPS 'Moving Planning and Scheduling to the Real World' workshop (Myers et al. 2007). However, we still appear to be very far away from the point where automated planning technology can be franchised to the software engineering community.

Our work is motivated by investigations into the use of AI planning in large-scale *control* applications. Automated assessment and prediction via monitoring and modelling is quite well developed in these kinds of applications, but there is a need to develop software support that enables active decision support or even autonomous control eg in water/flood control (Rob 2007), or road transport network control (Various 1999). However, how feasible is the use of AI planning tools within such an application area? How could we evaluate an application in terms of whether it can benefit from AI

planning technology, and how can we determine what areas of the application would cause the most problems? In this paper we explore the characteristics of an application area that make the application of AI planning feasible. To motivate the discussion, we use two particular applications from the Transport and Water Management service industries respectively. These are wide ranging, complex, involve many stake holders and organisations, and have allied research and development areas.

This endeavour has much in common with the general area of business process change through the introduction of new technology, and in particular the introduction of knowledge-based AI technology. The potential problem areas in the application of automated planning are in some cases similar to challenges already well known when implementing KBS systems. These include the 'knowledge bottleneck' - the difficulty of knowledge elicitation and formulation, the availability of experts and expertise, and the verification, validation, and maintenance of knowledge bases. The subject of this paper can be taken in the context of the well known reasons for failure of early KBS, to do with their brittleness and stand-alone nature. However, applications of automated planning can also take advantage of the more recent developments that alleviate the 'knowledge bottleneck': the development of shared ontologies and globally accessible knowledge, and the development of standard, tool support environments for the engineering of knowledge. For a discussion of the similarities and distinguishing features between knowledge engineering for AI planning and KBS, the reader is referred to section 7 of PLANET's Roadmap (Biundo et al. 2003).

In this paper we address the problem of evaluating the feasibility of applying AI planning technology, by devising a set of evaluation criteria based on motivation, technological infrastructure and knowledge engineering aspects of an application. To both illustrate and evaluate the usefulness of these criteria we use them to investigate the feasibility of applying automated planning technology to the applications. For each feature we rank

it as low, medium or high, indicating its contribution towards an overall feasibility factor. We conclude with a short discussion of the use of the criteria.

Feasibility Evaluation Criteria

We assume that an 'application area' has been identified, and there is a prima facie case for the use of automated planning within it. In the case of the two applications considered below: (i) road network management: planning can be used for drawing up plans to ease congestion or alleviate the effects of incidents (ii) flood prevention and management: planning may be applied to form plans for evacuations. Given this context, we postulate a number of key questions that need to be considered in evaluating the feasibility and effectiveness of utilising automated planning. We group them into 3 areas:

1. Motivation Factors

Motivation factors include the fundamental and underlying reasons for the introduction of planning technology. If a current system delivers an optimum solution, or a subset of stakeholders are satisfied with the operation of the system using current technology, then the motivation may be too weak. An example of low motivation is where there may be pressure to introduce advanced technology for its own sake, rather to satisfy a perceived need.

Overall the questions that should be asked include: are there compelling reasons for the introduction of technology: is it likely to deliver a step change in quality of service being offered eg increased reliability of plans, correctness of plans, real-time speed-up in the generation, and execution or distribution of plans? Will AI planning enable a significantly more cost-effective solution to some perceived problem?

2. Technological Context and Human Factors

Feasibility with respect to the application's context increases if there is already a high level of technological development within the service or industry. In human controlled systems there are several well defined phases in control: understanding what is happening in the system, evaluating that understanding (is there a problem?) and generating an effective plan to help alleviate the problem. The introduction of planning technology is more likely to be feasible if IT is already heavily used in the collecting, processing and interpretation of data, and in providing support for the current decision processes. If the data collected has an uncertain interpretation, or is incomplete, then the feasibility factor is lessened.

Given the nature of the technological change, it is help-

ful if there already exists experimental platforms to support the introduction of new technology. Typically this would comprise of historical data and a simulation system which can be used to investigate the effectiveness of techniques off-line.

Feasibility also depends on human factors: if key stakeholders are unwilling to accept the kind of autonomy delivered by automated planning then it will not be feasible. An example is where the current problem owners contract out the planning task to a third party. The third party is not necessarily going to be a willing partner in the venture if the new technology threatens that contract; the third party, however, may hold knowledge that is necessary for success. In summary, the key questions are:

Existing technological infrastructure:

Within the computer systems that are currently being used in the potential application area, are sophisticated systems used extensively for management information, and/or for decision support? What is the level of technological take-up in the area? Are the current systems stand-alone or fully interoperable?

Data availability and quality:

Is there a ready supply of data to supply state information on the observed system? Is the data in high level (information extracted) form, or is it in a very low level (eg numerical) form? Is the data trustworthy or does it contain a significant amount of uncertainty? Can data be extracted from a standard data interface? Is there historical data and/or a simulation environment that can be used to test new technology off line?

Human Factors:

Are the problem owners (the current providers of solutions) and other stakeholders open and supportive of innovation to help improve their methods and systems?

3. Knowledge Engineering Factors

There is a well known characterisation of AI planning technology that it requires the pre-engineering of a specific database of actions, heuristics etc. The task of engineering knowledge into such a particular form is itself made feasible by the presence of a number of factors, such as: existing high level formalisations of the domain, existing high level formalisations of plans, or the existence of similar planning domain models. These factors are very relevant in knowledge engineering for KBS in general, as it is well known that if all the expertise lies solely in the brains of experts, then the amount of effort involved in knowledge elicitation and knowledge formulation can be very high indeed. Applications where there are existing encodings of actions and plans are thus very attractive. Hence, if knowledge of the cur-

rent planning process in the application area is not in written form, or there are no examples of precisely encoded plans, then the feasibility is low, or at least the amount of resource needed to create a domain model and planning heuristics may be prohibitive. The key questions to consider are:

Closeness to previous applications:

Is the application area close or analogous to a previously defined planning benchmark, or a current fielded system? Can parts of domain models or previously engineered constraints be re-used?

Procedure formalisation in the problem area:

Are there existing encoded, formalised or written-down procedures or plans? Is the current system managed by experts using their own experience, or do they have recourse to manuals and training aids? Are there readily available examples of the kinds of plans that are needed to be generated?

Appropriateness for AI planning solution:

How far does the construction of a plan fall into the classic definition of generating orderings of instantiated action schema to achieve goals or decompose tasks? Does plan generation involve a great deal of uncertainty, mixed discrete and continuous variables, or large amounts of human skill?

Application Area: Road Network Management

Description

Road network management (RNM) relies on complex, integrated systems to meet increasing requirements upon the road network specified within policy documents from central and local government. The responsibility for managing the road infrastructure in the UK rests with the Highways Agency (HA) for the motorway and trunk road network and the Local Authority (LA) for the urban network. Short term traffic events, such as road works, accident, adverse weather conditions, occurring on either the motorway or urban network can have devastating affects on one another. Currently human operators respond to this kind of problem using their expert knowledge, but their effectiveness is limited as they have to interpret complex information fed to them, decide on which of an array of actions to take, and deal with the interface between urban and motorway traffic control. Within the UK there is a duty on LAs to manage their traffic networks efficiently and reduce traffic pollution. Clearly there is a need to develop systems that will support the road network operators objectives when they try to tackle congestion or other

problems, such as excessive fuel emissions, in an increasingly complex environment.

Evaluation using Criteria

Motivation Factors: high

Within RNM, there is a well defined split between understanding what is happening in the system, and generating an effective plan to help alleviate the problem. In the former case, there are many real time data feeds from which knowledge about the system can be extracted, including loop detectors, ANPR (automatic number plate recognition), and CCTV. In the latter case, the traffic manager can manage a situation by initiating a range of actions; this includes the setting of traffic light timings, variable message signs (VMS), variable speed limits (VSL), ramp metering and radio broadcasting. In real time traffic control of large road networks it has been demonstrated that necessary processing and decision making is beyond the capabilities of human operators alone, and as the demand for road usage increases, this difficulty in managing traffic effectively becomes more acute. Additionally, the cost of congestion is increasing over time and in the UK alone is expected to rise to £30 billion by 2010. Improvement to the efficiency of traffic control and management also can be linked to the reduction of emissions of air pollutants produced by road traffic.

An application for AI planning could be to generate traffic and transport system plans and courses of actions in real-time to enable more effective control of incidents and events. A similar application might be to help with crisis management across the LA and HA controlled networks by generating plans which take account of LA and HA priorities and interactions. Hence there is a clear aim and motivation for the introduction of this kind of technology: to increase the quality of plans (which involve lights, VSLs, VSMs, etc), taking into account an increasing amount of information flow, which will benefit the quality of life through reduced congestion and pollutant emissions.

Technological Context and Human Factors

Existing technological infrastructure: high

There has been a good record of adoption of computer systems in road network traffic management, and currently there are emerging common service platforms which will be beneficial to products and services delivered by technology providers. High level data platforms such as the HA's *Travel Information Highway* allow sophisticated software packages to both monitor and disseminate traffic information both to other services and to the general public (eg in the uk we have www.trafficengland.co.uk). The development of

self-adapting computer systems such as SCOOT¹ has been one of the most important single developments. SCOOT systems are used worldwide to control the timings and offsets of groups of traffic lights connected by a local road network. They adapt to different traffic levels, automatically adjusting light timings at related junctions in reaction to sudden or gradual changes in traffic flows.

Data availability and quality: medium to high

In the UK, the UTMC² is a relational database convention for data collected and distributed in the course of traffic management. UTMC provides a high level, standard platform for traffic applications to use and inter-operate. Local Authorities use systems such as SCOOT and UTMC to make effective and efficient use of technology in managing the local road network. However, there are some shortfalls with current systems, and rising traffic levels will only exacerbate the situation. The most serious problem is that, although motorways and city centres have traffic flow monitored, traffic flow outside of these areas is largely unknown, and there is still a high degree of uncertainty of the status of some networks.

Regarding evaluation platforms for testing new technology off-line, there is a long standing history of transport research using such methods, with large amounts of data available for testing and simulation.

Human Factors: high

Expertise in management and operation of the network appears to be thin, and there is a realisation within the service that this, and the growing complexity of the problem, will require more technological investment. There are a range of high-tech service providers in the sector who are experienced in technological innovation. All stakeholders appear ready to embrace further technological innovation (especially given the past success of SCOOT).

Knowledge Engineering Factors

Closeness to previous applications: medium

Within the area, there have been attempts to incorporate some kinds of specific automated reasoning systems into the control of motorway incidents eg in the MOLA system (Still, P.B and Harbord, B.J. 1998). No such attempt has been made in local authority-controlled roads in the UK.

Regarding similar domains, the Pipesworld domain from IPC-4 shares some characteristics with road transport: the basic domain consists of an arcs and nodes

network, with some arcs (roads) bi and some uni-directional. Also, the 'transporter' (pipe or road) does not move - objects move along them. Despite there being ways to abstract the complexity of road networks (eg by bundling traffic into distinct quanta) the complexity of the road network may well cause a problem of scale to current planning engines.

Procedure formalisation in the problem area: medium

Plans do exist on paper, but are not plentiful. Decisions and plans are made by experts on the basis of collated information of the road network. Current procedure formulation is at the level of SQL constructs.

Appropriateness for AI planning solution: medium

Parameterised actions can be formed to model the actions mentioned above, although the effects of such actions may be difficult to encode in propositional form. Propositional descriptions of road network status and goal criteria are not generally used in current systems.

Application Area: Flood prevention and management

Flood prevention and management (FPM) involves, as in RNM, local and national authorities, service industries, and research institutes. This is due to its perceived importance: throughout many parts of the world the prevention, early warning, crisis and post-crisis management of water inundation is an important factor in human well-being. We have identified two areas which incorporate two potential applications of AI planning: for long term planning of infrastructure to prevent or lessen the risk of flooding, and for real-time planning to support flood event management. The former area considers such criteria as climatic change and population change, and may involve flood defence design or even river design. The latter area falls under the heading of crisis management, and may incorporate evacuation management. Here (as in RNM) there is the need to understand what the status of the event is - this is essential to support the active management of any identified problems.

Below we concentrate on evaluating the feasibility of AI planning to support flood event management, and use the information from deliverables of the current EU project 'FLOODsite' to support it³ FLOODsite aims to develop tools to help in evacuation management, particularly meta-tools and frameworks for the building of specific decision support systems (DSS).

¹<http://www.scoot-utc.com/>

²<http://www.utmc.gov.uk/>

³<http://www.floodsite.net/>

Evaluation using Criteria

Motivation Factors: high

The need for plan generation support in the real-time scenario is directly supported by FLOODsite research: 'Given the large variety of possible scenarios generating flash floods, the pre-flood generation of all the corresponding emergency plans is out of reach' (FLOODsite workplan, page 23). This implies that the motivation is similar to incident management in the RNM application - to be able to produce sound plans in real time in response to a crisis in which there are a number and mix of information streams.

Context and Human Factors

Existing technological infrastructure: medium - high

There are many decision support systems that have been created to help in flood event management in the UK, France and Netherlands alone (Rob 2007). These DSS are typically GIS-based simulation systems with user-friendly interfaces. They can inform on flood distributions, identify population, transport and properties at risk; evaluate the likely effectiveness of flood defences etc. Communication between 'actors' is very important in flood event management (as in other incident/crisis management) and hence systems are aimed at information dissemination among emergency services and connected organisations.

The number of decision support systems suggests a high level of technological infrastructure. However, real-time use of technology within the sector appears to be targeted at disseminating information about the unfolding crisis to the range of emergency services that are called upon to assist. No systems seem to exist that perform support for flood event management in general, or evacuation planning for flood events in particular, by generating plans in response to a specific disaster. Indeed, in the area of flood event management, we could find no evidence that there exists systems that can validate or simulate pre-existing evacuation plans; that is systems that input water distribution models, and simulate the execution of disaster plans in real time, and evaluate them.

Data availability and quality: medium - high

Data from meteorological predictions, data concerning population densities, population characteristics, physical assets (safe building etc) and evacuation routes is readily available. On the other hand, while obtaining data for simulation is possible, it is currently not possible (according to FLOODsite) to generate up to date models of water levels, velocities etc in real time, due to the amount of computational time required. Hence any simulation systems would need to use precomputed

models.

Human Factors: medium - high

Research and innovation in this area is accepted as an essential ongoing activities by stakeholders in the field, hence there would be no threats to feasibility. Many of the potential users, however, would not be IT literate and hence any AI software would need to be embedded within user-friendly interfaces.

Knowledge Engineering Factors

Closeness to previous applications: medium

This area is clearly related to the more general area of crisis prevention and management. There has been a great deal of work on decision support for crisis or disaster management, ranging back more than 20 - 30 years, although only a fraction of this work has attempted to automate generation of plans. An exception is the ongoing work aimed at disaster management for eruptions of the Popocatepetl volcano in Mexico, where the techniques used are based on answer-set programming (Cortes, Solnon, and Martnez 2004). This work is aimed at integrating a planning function with existing GIS systems. The language used for representation incorporates some measures of uncertainty, and the system has the potential for generating simple emergency evacuation plans. However, the application appears as yet not implemented.

Evacuation planning is an activity that has already been used with the Planning community - it is used as an example within the recent textbook (Ghallab, Nau, and Traverso 2004). SIADEX (Fdez-Olivares et al. 2006) is a system that is currently undergoing tests in real fire fighting situations. It produces plans, monitors execution, and interacts with human experts to support management in forest fire fighting. The insights resulting from the SIADEX implementation would certainly contribute to the success of a flood event management application.

Procedure formalisation in the problem area: low - medium

In general, plans and procedures in the area are not formalised and if they exist are stated in natural language. However, there are some DSS that expect emergency response plans as an input, and evaluate them by calculating the effect. This implies the existence of some plan formulations.

Appropriateness for AI planning solution: medium

Many of the inputs required in a planning domain model have been formalised in past decision support systems: actions and methods representing resources to be used for evacuation, and objects such as carriers and

routes (eg road networks). The planning state would likely consist of flood levels, safe evacuation zones, spatial distribution of inhabitants, types of inhabitants (eg able-bodied or not). While this appears appropriate for AI planning technology (and the closeness of this domain to general disaster management is evident) the continuous nature of the domain, in particular flood distribution, may be difficult to represent with current domain model languages.

Discussion

For the RNM application, strengths lie in the technological infrastructure, the motivation for the work, and (in the UK at least) the availability of interoperable services due in part to the standard technology platform (UTMC). The existence of software in the industry with AI characteristics (SCOOT and MOLA) is an important factor. The main problems seem to be the lack of high level information about road network status (which equates to the 'world state' in planning), and the lack of precisely defined plan databases. For FPM, again, motivation and technological infrastructure and innovation is generally high. The main areas of concern are within the knowledge engineering aspects, particularly plan reasoning and representation aspects. In both areas then, it would seem that the applications are feasible, but more work is required to quantify the resources required to complete the knowledge engineering task.

Another application area we have investigated resulted in a remarkably different result in the feasibility criteria, leading to us not pursuing the application of AI planning. This is an historical example (in that it may not still hold today given the changes in technology) from the area of Air Traffic Control. It is based on our early work in formalisation of ATC separation criteria (McCluskey et al. 1995). The application area is to produce a planning aid for helping conflict resolution of aircraft during en-route control over North Atlantic airspace. Thus the planner would need to take existing route plans, adjust them to clear any airspace conflict that had been detected by a conflict probe, and output the new plans to an air traffic control officer. In this domain the level of current technology was high, data on aircraft positions and plans was very good, and offline evaluation was possible. Additionally, the knowledge engineering aspects were good: there were rule books, formalised plans, propositional state descriptions, and much of the context knowledge had been formalised through our previous work on aircraft separation criteria. The criteria that scored low were motivation and human factors: investigation showed that although automated aids were desired by some state holders, the en route air traffic control officers were quite happy with their current method, which was capable of delivering the plans without the need for extra technology.

Conclusions

In this paper we introduced a set of criteria for evaluating the feasibility of introducing planning technology into an application area. We applied these criteria to two application areas which (as yet) have not seen AI planning applications. Although with the applications considered the introduction of AI planning was thought to be feasible (with some reservations), the exercise appears to illustrate to us the difficulty in finding suitable application areas: an application must score well on all three aspects: motivation, technological infrastructure and knowledge engineering aspects.

Some aspects of the criteria were based on those that would be used when assessing an application for the introduction of a KBS, as the same problems of knowledge elicitation and availability of expertise may be evident. In control applications, however, a further important factor seems to be the level of technological progression within the industry. In order to integrate planning technology into a currently human controlled system, there should already exist high levels of technological use and expertise in the industry, such as examples of past success with AI technology. A parallel can be drawn with the field of Autonomic Computing (Kephart and Chess 2003), which is to do with the manufacture of computer systems which take care of themselves in that they can self-configure, self maintain, self-heal etc. The protagonists of AC portray the deployment of autonomic qualities as the culmination of a technological progression along which the progress of an application area can be tracked. Hence, for an application area to adopt a new system incorporating autonomic features, the current technology must already be far advanced (eg the current system may have software components with intelligent characteristics). This seems the case with AI planning technology also: the application area in general must be technologically sophisticated enough to support knowledge engineering of the required dynamic and heuristic knowledge to make plan generation feasible.

Acknowledgements

I would like to thank personnel from the following organisations for their input: telent, AVLIC Ltd, AeGis, TSEU, TRL, NASA, Wallingford Software, and Kent and Hampshire County Councils.

References

- Biundo, S.; Aylett, R.; Beetz, M.; Borrajo, D.; Cesta, A.; Grant, T.; McCluskey, T.; Milani, A.; and Verfaillie, G. 2003. PLANET technological roadmap on AI planning and scheduling. Electronically available at <http://www.planet-noe.org/service/Resources/Roadmap/Roadmap2.pdf>.

- Cortes, C. Z.; Solnon, C.; and Martnez, D. S. 2004. Planning operation: An extension of a geographical information system. Proceedings of the 1st Intl. LANMR04 Workshop, Antiguo Colegio de San Ildefonso, Mexico City, D.F., Mexico.
- Fdez-Olivares, J.; Castillo, L.; Garcia-Perez, O.; and Reins, F. P. 2006. Bringing users and planning technology together: experiences in SIADEX. In *Proceedings of the Sixteenth International Conference on Automated Planning and Scheduling (ICAPS 2006)*, 11 – 20.
- Ghallab, M.; Nau, D.; and Traverso, P. 2004. *Automated Planning: Theory and Practice*. Morgan Kaufmann ISBN 1-55860-856-7.
- Kephart, J. O., and Chess, D. M. 2003. The vision of autonomic computing. *Computer* 36(1):41–50.
- McCluskey, T. L.; Porteous, J. M.; Naik, Y.; Taylor, C. N.; and Jones, S. 1995. A Requirements Capture Method and its use in an Air Traffic Control Application. *Software-Practice and Experience* 25:47–41.
- Myers, K.; Frank, J.; McCluskey, L.; and Yorke-Smith, N. 2007. Workshop on Moving Planning and Scheduling Systems into the Real World. <http://icaps07.icaps-conference.org/>.
- Rob, M. 2007. Review of flood event management decision support systems. Technical Report FLOODsite deliverable T19-07-01, Delft Hydraulics.
- Still, P.B and Harbord, B.J. 1998. Strategic management of traffic in Kent using MOLA. In *9th International Conference on Road Transport Information and Control*.
- Various. 1999. UMTC 04: Research Review and Requirements Report. Birmingham City Council and Leicester City Council and University of Leeds, <http://www.utmc.gov.uk/research/>.