Pathways to Impact

Automated planning and scheduling (APS) is of great importance for autonomous intelligent systems. APS enables such systems to exhibit goal-directed behaviour that goes beyond immediate goals and is considered to be an important aspect of general intelligence. This is why planning is an active research area that has been investigated in Artificial Intelligence (AI) almost from the beginning of the field. Planning can be implemented by purpose-built algorithms in specific domains where the cost-benefit ratio justifies their development (e.g. for scheduling problems). However, for autonomous systems in a complex environment a general and domain-independent planning capability is necessary to achieve goal-directed behaviour. Thus, the ability to plan is an important component of rational thought and action in an autonomous system.

Impact to Knowledge

Domain-independent planning algorithms have made enormous progress over the past 15 years. Most planning systems require as input a description of the initial state of the world, a goal specification (or a task in the case HTN planning), and a domain model. While the former two are problem-specific and offer little potential for reuse, the domain model is intended to be common to many problems. Thus tools that support domain modelling support the creation of reusable knowledge.

The domain model essentially describes the actions, objects, resources and events that are possible in a given abstraction of the world. In many applications this constitutes a bottleneck, as the development of a good domain model requires a planning expert for a formalization at an adequate level of abstraction to allow for efficient planning, and it requires a subject matter expert with detailed knowledge of the domain to be modelled. And even with such experts at hand, the resulting models are often incorrect, difficult to maintain, and elaboration intolerant. The tools and methods developed in the proposed work aim to address this problem, *enabling the development of formal domain models more rapidly, on a larger scale, and with fewer conceptual flaws*. This reusable knowledge will facilitate the development of autonomous intelligent systems.

We plan to disseminate the research results initially via workshop papers, then more developed theoretical results and software demonstrations will be submitted to leading conferences such as ICAPS and AAAI. Longer papers will be submitted to journals such as the Knowledge Engineering Review and $JAIR$. A project website will be established and maintained to provide a further route for dissemination of publications and research prototype software where appropriate. The results of the project will impact significantly on the Knowledge Engineering (KE) area of APS since the aim of the research project is known to be a major challenge in this area. The KE community is particularly concerned with looking to exploit APS research, since it involves aspects of integration and context that are necessary to apply APS technology. Hence through the KE for APS community, the project will impact on those looking to deploy their research ideas and embed planning within real applications.

The proposed work will impact on the wider planning community: more domain models will become available, which will provide new challenges for planning algorithms. Developments in planning algorithms are currently driven by the international planning competition which has an emphasis on transportation domains. As a result of this project domain experts from various fields will be able to develop formal domain models, and this should make a much wider spectrum of domains available to the competition. This in turn will enable planning algorithm developers to extend their algorithms and heuristics to broaden their applicability. It will lead to a better understanding of which techniques work in which type of domain, and informed by AIS-DDL (WP3 in the workplan), to new extensions to the underlying representation languages used.

The research will strengthen UK expertise in this area: the proposers have strong links with current UK APS leaders, primarily with the University of Strathclyde, where Prof Maria Fox and Prof Derek Long are carrying out research into automated reasoning in mixed discrete/continuous systems especially applied to policy learning.

This proposal will advance the state of knowledge in four research areas of the Research Interests table:

Model Building and Learning (8.0). The proposal concerns the acquisition, learning, validation, maintenance and adaptation of "reference models" (here called domain models) and therefore directly addresses the research problems in this area.

Planning and Decision Making (4.0). The main role of the domain model (referred to above) is to enable the system to construct plans to achieve desired goals, as the domain model is an input to automated planning.

Situational Awareness and Information Abstraction (3.0). To be able to adapt and change the domain model requires information inferred from sensor data. Furthermore, the learned domain model will be useful for assessing an agent's capabilities.

Verification and Validation of Autonomous Systems (7.0). The proposed project will contribute to this in so far as it will address the verification and validation of the domain model.

Impact to the Economy

The tools developed in this project will increase adoption of automated planning in autonomous systems as they make them accessible to domain experts with little knowledge of AI planning techniques or formal knowledge representation. This will lead to a competitive advantage as planing technology can be included in many new application areas.

This project will create a community of domain experts who are not expert in AI planning technology to develop systems that include formal domain models and planning tools. The knowledge engineering tool (WP4 part c) will help these users create initial domain model knowledge, simulation and domain analysis tools will help them to understand the resulting model better, and expert users will be able to critique the acquired domain model.

A major pathway to economic impact is via the AIS Programme Call collaborators. The research is seen by the proposers as fundamental to all the collaborators' scenarios as described in the Call. More specifically, the proposed work will have the following impact:

Scenario: Multi-Vehicle Cooperative Autonomy. This project will enable planning in a collaborative setup, where communication and constraint satisfaction of capabilities are essential, and can be captured within a domain model. Model analysis and verification techniques will increase the robustness of the system.

Scenario: Long Range Mars Rover. Autonomous vehicles that cannot be remotely controlled must be able to plan. This project will enable such systems to learn and adjust their domain model to cope with unknown and unforseen environments. As a result Mars rovers using this kind of technology will be able to reason about their own capabilities and explore new ways of acting in the environment, enabling them to extend their capabilities and create novel plans.

Scenario: Investigate and Repair Defective Infrastructure. Assuming that there is a wide range of infrastructures and different autonomous systems to diagnose and repair these infrastructures, it will be necessary to develop different domain models, which can be achieved with the results of this project. This will enhance the flexibility of the whole system of agents by increasing its elaboration tolerance and speed with which it adapts to new circumstances.

Scenario: Teleoperation to Teleautonomy. This group of scenarios is aimed at agents operating in difficult environments that behave in unusual ways, which requires agents to adapt their behaviour. Again, the learning and model analysis techniques of this project are features that will enable autonomous, long term goal-directed behaviour.

Impact to Society

While the tools and methods of the project are not aimed at any particular application area (except for autonomous systems in general), we intend to exploit links with the ISCRAM (Information Systems for Crisis Response and Management) and the ITS (Intelligent Transport Systems) communities. Dr Wickler is a member of the board of directors of the ISCRAM association. The aim is to adopt more planning technology for crisis response domains. While crisis response and management has adopted a planning approach in the military (a lot of planning technology was developed in this context), many civil agencies are only just getting to the point where planning technology could be helpful to them. This project and the resulting tools will lower the threshold for entry into this technology for this community and we will promote results of the project at the annual ISCRAM conference. With respect to ITS, Prof McCluskey is leading a COST Action Network aiming to apply ideas of autonomic computing to road network support systems. The proposed research will impact on this network in the form of producing advanced techniques that can be applied to increase the autonomy in transport network. For example, in the area of road traffic control, automated planning and scheduling may be used to create plans for traffic lights, diversions, variable speed limits, hard shoulder operation, variable message signs etc.

While the planned outputs on verification and validation, including the mixed-initiative tool (WP4) will have the greatest impact on application communities, the fully automatic methods (WP5) will have impact on domains where human intervention is limited or impossible (e.g. Mars rover, underwater vehicles). Of course, it will still be possible to use the tools at design time in a mixed-initiative mode. In such domains the automated learning can extend the capabilities of autonomous vehicles beyond what has been anticipated by the designers, who may not be able to test in the real environment. This is difficult to evaluate, but one can easily imagine a scenario where a Mars rover is stuck; in this case it could switch to automated domain learning mode to explore new options.

Many areas of industry use software tools to underlie their monitoring, modelling, and management functions. Automated assessment and prediction via monitoring and modelling is well developed in such application areas; however there is a pressing need to develop software to support or automate the management process. The work in this project will lead to underpinning theory, knowledge acquisition tools and application prototypes demonstrating the potential for automated management in these kinds of control applications.

Impact to People

This project will have impact on the people currently working as developers of autonomous systems. More people will have the skills to use planning technology in their application areas, as the results of this project will greatly increase the accessability of this technology.

Furthermore, autonomous systems that can learn and adapt the underlying domain model will be more flexible and therefore usable in contexts where previously humans were required, e.g. in highly dangerous situations that are common in disaster response. This will result in increased quality of life for emergency responders and more security for society as a whole living in a safer environment.