Future State Projection as Planning

Shirin Sohrabi and **Octavian Udrea** and **Anton V. Riabov**
IBM T.J. Watson Research Center
1101 Kitchawan Rd, Yorktown Heights, NY 10598, USA
{ssohrab, udrea, riabov}@us.ibm.com

Abstract

We extend the plan-recognition-as-planning technique to both explain the past and project the future. The Planning Projector helps users understand future possibilities in order to make better decisions.

1 Introduction and Motivation

Being prepared for the future is of great importance to many, if not all, human endeavors. In this work, we develop the Planning Projector system prototype, which applies planning to project alternative future state trajectories, where one state is created from another by an application of an action, in accordance with a planning domain definition that captures expert knowledge. Additionally, instead of a fully specified initial state, our system receives a sequence of observations from the past and, possibly, the present. These observations, each expressed as a set of predicates, must then be mapped to inferred states as well, which again must be connected by actions, and hence the system must infer (recognize) plans for the past and the present, as well as project the future.

Consider the following energy domain example, where the objective is to project the price of oil and volume of oil produced 15 years into the future. Note, our objective is not to find a precise estimate of the price of oil, but rather to project the possible ranges, as well as provide state trajectories that lead to those ranges. The Planning Projector relies on domain knowledge that can either be provided by domain experts, or encoded by non-experts after reviewing various sources of available knowledge, such as research papers, textbooks or Wikipedia. The domain knowledge in this example would describe possible actions affecting oil price directly or indirectly, for example by affecting supply levels. For instance, the decision of the leaders of OPEC (Organization of the Petroleum Exporting Countries) to meet is an action that is likely to affect both the price and the supply of oil, depending on the outcome of such a meeting. The decision to limit production will decrease the supply and increase the price, or the decision to increase supply can lead to lower prices. The observations associated with these actions, confirming or denying them, can be derived from news reports. Similarly, several other events or actions can be modeled: the discovery of a new oil field, drilling activity in known fields, hurricanes or other natural disasters affecting oil production, and changes in currency rates.

There are multiple challenges that the Planning Projector must address: (1) the observations can be unreliable (i.e., noisy, missing, inconsistent) and the domain model may be incomplete, (2) no specific set of goals may be given, (3) the projected future state space can become extremely large, even with short time horizons. To address the latter challenge, our implementation generates many scenarios efficiently and uses special techniques to produce summary reports.

2 Planning Projector

We model the future state projection problem as plan recognition by adding a sufficient number of future observations to the end of the observation trace, as many as needed to fill the time horizon of the projection. These future observations are defined to match any allowed future state. Following an extended version of Ramírez and Geffner 2010 approach, we translate the resulting plan recognition problem in to a planning problem [Sohrabi, Riabov, and Udrea, 2016].

We introduce two extensions to plan recognition as planning. First, we do not require observations to be exclusively over actions, allowing any fact or fluent as an observation of a state. Furthermore, we do not require the observations to be complete and consistent, nor the domain model description to be complete. To address the missing or noisy observations we use the theory described in [Sohrabi, Udrea, and Riabov, 2013] and use a planner capable of finding a set of plans; we have experimented with diverse planning [Nguyen et al., 2012] as well as top-k planning [Riabov et al. 2014]. As a result, the Planning Projector is a domain-independent system that can not only generate multiple explanations of the past observations, but can also extend these explanations past the last observation, generating possible future scenarios.

3 User Interface

Figure 1 shows the main input screen of the Planning Projector. At the top, the observation sequence is entered. The observations can be entered one by one by selecting from a drop down menu containing known fluents. In this example, the observations are of ongoing exploration in a field in North America, and of bidding on an oil field called Terra Nova.

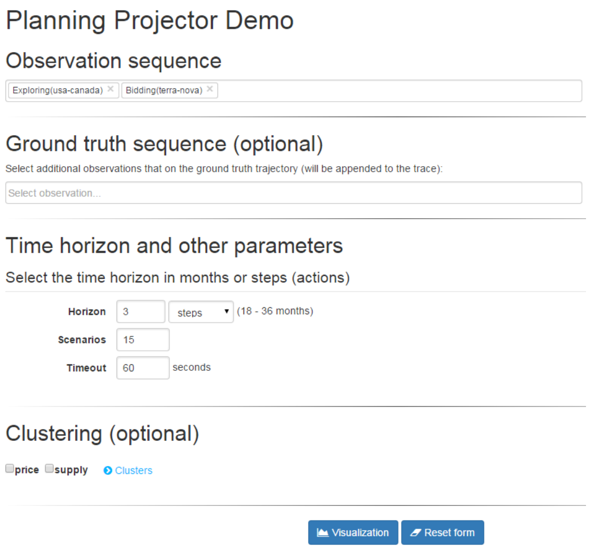

Planning Projector Demo

Observation sequence

Exploring(usa-canada) × Bidding(terra-nova) ×

Ground truth sequence (optional)

Select additional observations that on the ground truth trajectory (will be appended to the trace):

Select observation...

Time horizon and other parameters

Select the time horizon in months or steps (actions)

Horizon 3 steps (18 - 36 months)

Scenarios 15

Timeout 60 seconds

Clustering (optional)

☐ **price** ☐ **supply** Clusters

Figure 1: Planning Projector input screen

The ground truth is an optional sequence of future observations that is highlighted in the output if found, and is used for testing the domain models, helping to ensure the expected scenario occurs among the generated future states. The time horizon can be specified as the number of steps or the number of months into the future. The number of scenarios can also be specified, as well as the planning timeout in seconds. Optionally the user can express interest in grouping the results in clusters around their values of interest. These numbers can be entered by clicking on “Clusters” at the bottom of the page.

Figure 2 shows the results page which is generated by formulating a plan recognition problem, translating it to a planning problem, and running a planner to compute a set of plans, as described in the previous section. The x-axis shows the time horizon, where 0 indicates the current time, the positive numbers indicate steps into the future, and the negative numbers indicate the past. Each circle corresponds to a state, and each rectangle represents an action, where longer rectangles represent actions of longer duration. Hence, the set of plans are shown by their state-action sequences. The small blue triangle indicates the action that explains an observation in the past. The user can move the mouse over the states and actions in order to get a more detailed description, including all facts that are true in the state, and action name, duration, and how the action influences the price and supply of oil. The actions are color-coded for a quick visual analysis, where red indicates a decrease in the price of oil and green indicates an increase. The intensity of the colors indicates the degree of the change in price.

The last states, represented by circles at step 24, can be of particular interest to the user. Theses states can be clustered optionally around given values of price or volume of oil. We use a simple clustering technique to cluster the states based on a configurable constant threshold. There are two kinds of clusters, as shown in blue and in orange. The blue clusters are those that are centered around the given values provided by the user, while the orange clusters are discovered by the system without user guidance. Hence the states in the blue clusters, and the state sequences leading to them, may confirm the prior information that the user may have, while the orange clusters may be unexpected and investigating those can provide valuable information to the users of the system.

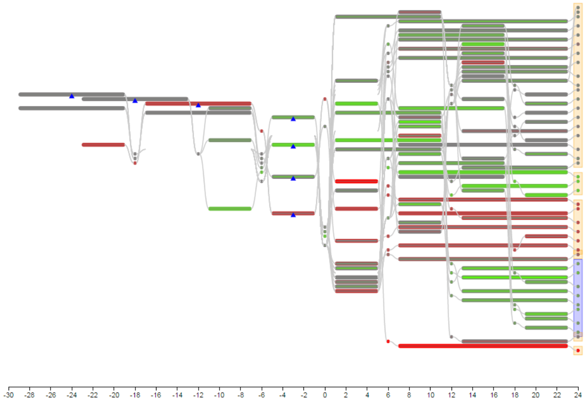

Figure 2: Planning Projector results page

Note, the user may interact with the system through the command line and/or view and manipulate the generated plans which can be exported in JSON object format.

4 Summary

We present the Planning Projector system developed by extending the prior work on plan-recognition-as-planning. The system is able to infer the past and project the future. The web interface presents the generated future states, and further clusters them to facilitate analysis and interpretation. In the future, we plan to extend the formalism to better capture the interactions between concurrent actions and simplify the modeling of independent agents acting within one domain.

5 Acknowledgements

This material is based upon work supported in whole or in part with funding from the Laboratory for Analytic Sciences (LAS). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the LAS and/or any agency or entity of the United States Government.

References

[Nguyen et al., 2012] Nguyen, T. A.; Do, M. B.; Gerevini, A.; Serina, I.; Srivastava, B.; and Kambhampati, S. 2012. Generating diverse plans to handle unknown and partially known user preferences. *AIJ* 190:1–31.

[Ramírez and Geffner, 2010] Ramírez, M., and Geffner, H. 2010. Probabilistic plan recognition using off-the-shelf classical planners. In *AAAI*.

[Riabov, Sohrabi, and Udrea, 2014] Riabov, A.; Sohrabi, S.; and Udrea, O. 2014. New algorithms for the top-k planning problem. In *SPARK*, 10–16.

[Sohrabi, Riabov, and Udrea, 2016] Sohrabi, S.; Riabov, A.; and Udrea, O. 2016. Plan recognition as planning revisited. In *IJCAI*, To appear.

[Sohrabi, Udrea, and Riabov, 2013] Sohrabi, S.; Udrea, O.; and Riabov, A. 2013. Hypothesis exploration for malware detection using planning. In *AAAI*, 883–889.