

IBM Scenario Planning Advisor: Plan Recognition as AI Planning in Practice

Shirin Sohrabi, Michael Katz, Otkie Hassanzadeh, Octavian Udrea, Mark D. Feblowitz

IBM Research, USA
ssohrab@us.ibm.com

Abstract

We present the IBM Research Scenario Planning Advisor (SPA), a decision support system that allows users to generate diverse alternate scenarios of the future and enhance their ability to imagine the different possible outcomes, including unlikely but potentially impactful futures. The system includes tooling for experts to intuitively encode their domain knowledge, and uses AI Planning to reason about this knowledge and the current state of the world, including news and social media, when generating scenarios.

1 Introduction and Motivation

Scenario planning is a widely accepted technique by which organizations develop their long-term plans [Schoemaker, 1995]. Scenario planning for risk management puts an added emphasis on identifying the extreme yet possible risks and opportunities that are not usually considered in daily operations. Scenario planning involves analyzing the relationship between forces - such as social, technical, economic, environmental, and political trends - in order to explain the current situation and to provide insights about the future. A major benefit to scenario planning is that it helps businesses or policy-makers to learn about possible alternative futures and to anticipate them. We use AI planning, informed by expert domain knowledge, because some scenarios have never yet occurred and thus cannot be projected by probabilistic means. And we generate many different scenarios, exploring a variety of possible futures; because we want to be prepared for both expected and surprising futures [Peterson et al., 2003].

Risk management is a set of principles that focus on the outcome for risk-taking [Stulz, 1996]. A variety of (manual) methods and standards for risk management have been developed [Avanesov, 2009]. Our approach in addressing scenario planning for risk management is different from previous work in that we reason about emerging risks based on observations from the news and social media trends, and produce scenarios that both describe the current situation and project the future possible effects of these observations. Furthermore, each scenario we produce highlights the potential *leading indicators*, the set of facts that are likely to lead to a scenario, the *scenario and emerging risk*, the combined set of consequences

or effects in that scenario, in addition to the *business implications*, a subset of potential effects of that scenario that the users care about and are relevant to the enterprise. For example, prior to the Brexit referendum in 2016, an international company operating in the United Kingdom could consider alternative future scenarios for changes in trade and employment treaties assuming the majority voted to leave the European Union, identifying the implications for the company's finances and its ability to hire, enabling the company to act immediately to minimize the negative implications.

The main functions of Scenario Planning Advisor (SPA) are: 1) discovering active risk drivers by aggregating relevant news from the Web and social media, and generating lists of candidate observations corresponding to the detected risk drivers¹; 2) generating multiple alternative scenarios based on the user-selected observations, the domain knowledge about the driver relations, and the potential likelihoods and impact of their cascading effects². The main idea of the approach in SPA is to characterize the scenario planning problem for enterprise risk management as a plan recognition problem [Sohrabi et al., 2018] and use AI planning to address the plan recognition problem [Ramírez and Geffner, 2009; Ramírez and Geffner, 2010; Sohrabi et al., 2016a; 2017a]. To this end, we transform the domain knowledge as captured by Mind Maps, a graphical representation that captures concepts and their relations, into an AI planning task. Furthermore, we generate multiple high-quality plans [Riabov et al., 2014; Sohrabi et al., 2016b; Katz et al., 2018] and cluster them into scenarios. These scenarios are then used to start a risk conversation between the analysts and decision makers.

2 News Aggregation & Observations Selection

SPA continuously monitors multiple real-world sources (e.g., news channels from RSS and Atom feeds, social media posts) in multiple languages in order to provide continuous awareness of an evolving world, filtered and aggregated based on users' interest profiles. To this end, several text analytics are implemented to find the information relevant for a particular domain in the vast amount of information available to crawl. To further refine and filter the information, SPA uses

¹Demo video is here: <https://youtu.be/6BRXwSkRC04>

²Demo video is here: <https://youtu.be/N0AsRbdL7WM>




Figure 1: Part of the Mind Map for the currency depreciation.

the structured semantic knowledge available in Wikidata using Wikidata Query Service to find important people, organization, and relevant sources [Sohrabi *et al.*, 2017b]. Analysts then review the generated results and select the observations that are the most relevant and important for them. Note that SPA can deal with unreliable observations (i.e., noisy, inconsistent, missing observations) as it exploits previous work on plan recognition as planning [Sohrabi *et al.*, 2016a].

3 Scenario Generation via AI Planning

3.1 Domain Knowledge

The Mind Maps and the customization components encode knowledge about risk drivers and business implications elicited from the domain experts and the local country experts correspondingly. While the reasoning engine in SPA supports a rich representation of risk drivers as actions in Planning Domain Description Language (PDDL) [McDermott, 1998], the knowledge representation used by domain experts is drastically simplified, to prevent conflicts and reduce overheads in knowledge elicitation and maintenance. The domain experts use Mind Maps created in FreeMind (freemind.sourceforge.net/wiki/), a graphical tool that encodes concepts and relations, to capture directed graphs of risk drivers and business implications, with edges having hidden semantics of pairwise cause and effect. An example of a Mind Map for the currency depreciation force is shown in Figure 1. The customizations are elicited using generated questionnaires that request country-specific likelihood and impact for selected cause and effect pairs. Due to customization, the same observations may generate different scenarios in different countries. The domain knowledge encoded in the Mind Maps, together with the answers from the questionnaire are automatically translated into a planning task. To do so we define an action for each cause, as well as each edge in the Mind Map. For details see [Sohrabi *et al.*, 2018].

3.2 Plan and Scenarios Computation

To compute a set of scenarios (see Figure 2 for an example of a scenario) we perform the following steps: (i) follow previous work on plan-recognition-as-planning to compile away the observations and ensure that at least one goal is satisfied, (ii) compute a set of high-quality plans on the transformed planning problem, and (iii) cluster the resulting plans into scenarios so that similar plans are grouped together. The scenarios are clusters of high-quality plans that include a trajectory of cause-effect transitions from the Mind Maps, explaining the largest possible subset of observations, and such that each plan ends with a business implication. To compute a large set of high-quality plans on the transformed planning task, we use top- k planning techniques [Sohrabi *et al.*, 2016b;

Leading indicators	Scenario and Emerging Risks	Business Implications
Competitive FX rate benefits net exporter country	<ul style="list-style-type: none"> • Currency depreciation against US dollar 	<ul style="list-style-type: none"> • IBM workforce capital available at better rates
High inflation	<ul style="list-style-type: none"> • Increase in taxes 	<ul style="list-style-type: none"> • Rising cost to attract and retain talent in IBM
Currency depreciation against US dollar	<ul style="list-style-type: none"> • Increase in government spending (education) 	<ul style="list-style-type: none"> • Increased client use of newer technology creates opportunity for IBM offerings
Reduced tax incentives	<ul style="list-style-type: none"> • Decrease in price of commodity > Weakening economic environment 	<ul style="list-style-type: none"> • Increased opportunity for IBM to invest and expand
Increased motivation for businesses to outsource work	<ul style="list-style-type: none"> • Reduced tax incentives 	

Figure 2: Sample generated scenario

Katz *et al.*, 2018]. Top- k planning is defined in as the problem of finding k set of plans that have the highest quality. We use the K^* algorithm [Pommerening and Helmert, 2012] in the SPA system. To cluster the plans, we apply a hierarchical clustering algorithm on the resulting plans [Defays, 1977]. To compare plans with each other, we consider the union of the set of states traversed by that plan. Hence, rather than presenting all plans, we group similar plans and only present 3-6 clusters of plans or scenarios to the end user.

4 User Interface

We present the scenarios as both text summaries and graphically. We perform several tasks to prepare the scenarios for presentation. First, we separate the predicates in each cluster (scenario) into business implications and regular predicates. Second, we compute a set of leading indicators. These are predicates that appear early on the plans that are part of one scenario and are also discriminating; these are useful to monitor in order to determine early on whether a scenario is likely to occur. Third, we compute a summary of all plans that are part of the scenario and present this as a graph to the user. This graph serves as an explanatory tool for the predicates that are presented in each scenario and shows how the different Mind Maps are connected with each other.

5 Summary

We present the SPA system developed by applying plan-recognition-as-planning techniques for scenario planning for risk management. To this end, we addressed knowledge engineering challenge of encoding the domain knowledge, and designed a system that takes as input news and social media posts aligned with user interests, interacts with the business user to obtain observations, and generates scenarios by clustering the generated plans. We believe AI planning provides a very natural formulation for the efficient exploration of possible outcomes required for scenario planning.

Acknowledgements

We thank Claudio Martinez, Deepak Turaga, Nagui Halim, and Edward Shay for their guidance and support. 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 authors and do not necessarily reflect the views of the LAS and/or any agency or entity of the United States Government.

References

- [Avanesov, 2009] Evgeny Avanesov. Risk management in ISO 9000 series standards. In *International Conference on Risk Assessment and Management*, volume 24, page 25, 2009.
- [Defays, 1977] D. Defays. An efficient algorithm for a complete link method. *Computer Journal*, 20(4):364–366, 1977.
- [Katz *et al.*, 2018] Michael Katz, Shirin Sohrabi, Octavian Udrea, and Dominik Winterer. A novel iterative approach to Top-k planning. In *Proceedings of the 28th International Conference on Automated Planning and Scheduling*, 2018.
- [McDermott, 1998] Drew V. McDermott. PDDL — The Planning Domain Definition Language. Technical Report TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control, 1998.
- [Peterson *et al.*, 2003] Garry D Peterson, Graeme S Cumming, and Stephen R Carpenter. Scenario planning: a tool for conservation in an uncertain world. *Conservation biology*, 17(2):358–366, 2003.
- [Pommerening and Helmert, 2012] Florian Pommerening and Malte Helmert. Optimal planning for delete-free tasks with incremental LM-Cut. In *Proceedings of the 22nd International Conference on Automated Planning and Scheduling (ICAPS)*, 2012.
- [Ramírez and Geffner, 2009] Miquel Ramírez and Hector Geffner. Plan recognition as planning. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1778–1783, 2009.
- [Ramírez and Geffner, 2010] Miquel Ramírez and Hector Geffner. Probabilistic plan recognition using off-the-shelf classical planners. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 1121–1126, 2010.
- [Riabov *et al.*, 2014] Anton Riabov, Shirin Sohrabi, and Octavian Udrea. New algorithms for the top-k planning problem. In *Proceedings of the Scheduling and Planning Applications workshop (SPARK) at the 24th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 10–16, 2014.
- [Schoemaker, 1995] Paul JH Schoemaker. Scenario planning: a tool for strategic thinking. *Sloan management review*, 36(2):25, 1995.
- [Sohrabi *et al.*, 2016a] Shirin Sohrabi, Anton Riabov, and Octavian Udrea. Plan recognition as planning revisited. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3258–3264, 2016.
- [Sohrabi *et al.*, 2016b] Shirin Sohrabi, Anton Riabov, Octavian Udrea, and Otkie Hassanzadeh. Finding diverse high-quality plans for hypothesis generation. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*, pages 1581–1582, 2016.
- [Sohrabi *et al.*, 2017a] Shirin Sohrabi, Anton Riabov, and Octavian Udrea. State projection via AI planning. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI-17)*, pages 4611–4617, 2017.
- [Sohrabi *et al.*, 2017b] Shirin Sohrabi, Anton Riabov, Octavian Udrea, and Fang Yuan. Using lightweight semantic models to assist risk management in a large enterprise. In *Proceedings of the 16th International Semantic Web Conference - Industry Track (ISWC-17)*, 2017.
- [Sohrabi *et al.*, 2018] Shirin Sohrabi, Anton V. Riabov, Michael Katz, and Octavian Udrea. An AI planning solution to scenario generation for enterprise risk management. In *Proceedings of the 32nd National Conference on Artificial Intelligence (AAAI)*, pages 160–167, 2018.
- [Stulz, 1996] René M Stulz. Rethinking risk management. *Journal of applied corporate finance*, 9(3):8–25, 1996.