

Using Lightweight Semantic Models to Assist Risk Management in a Large Enterprise

Shirin Sohrabi, Anton Riabov, Octavian Udrea, and Fang Yuan

IBM T.J. Watson Research Center

1 Introduction

In this paper we summarize our experience and the initial results from implementing and operating IBM Scenario Planning Advisor (SPA), a decision support system that uses lightweight semantic models to assist finance organizations in identifying and managing emerging risk, a category of risk associated with the changes in the global or local economies, politics, technology, society, and others.


SPA is designed to support the business process called “scenario planning” [3] that consists of preparing several future scenarios, followed by identifying the implications for the business, and finally choosing the mitigation actions to be taken. For example, prior to the Brexit referendum in 2016, an international company operating in the UK could consider alternative future scenarios for changes in trade and employment treaties assuming the majority voted to leave the EU, identifying the implications for the company’s finances and its ability to hire, enabling the company to act immediately to minimize the negative impacts.

The main functions of 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 drivers; 2) generating multiple alternative future scenarios highlighting their business implications and leading indicators, based on user-selected observations, and using domain knowledge about driver relations, cascading effects, and implications.

A key design decision in SPA is that the system does not compute probabilities for the generated scenarios. Instead, we recommend that the domain experts assign probabilities to the final 3-5 scenarios when necessary. In our experience that approach delivers value to the users much faster and requires significantly less work from the experts, compared to creating prediction models for over 200 often non-stationary risk drivers (e.g., oil prices or election results), and their interactions, as would be required to derive scenario probabilities automatically.

We have identified the following major challenges in developing SPA: 1) capturing observations from news and social media; 2) capturing the domain knowledge about the risk drivers quickly and efficiently, while preventing conflicts; 3) reasoning with incomplete and biased input to include sufficiently complete and minimally biased sets of risks and opportunities in generated scenarios.

In the rest of the paper, we describe the components of the SPA system, explaining their role in addressing the challenges, and highlighting the use of semantic technologies throughout the SPA system. We then present some details of the deployment that indicate our initial success in overcoming the challenges.



2 Semantic Components in SPA

The above figure shows the interactions between SPA components and its users.

Capturing Observations. The News Aggregator component aggregates news from RSS and Atom feeds and social media posts, e.g., Twitter, in multiple languages, by monitoring user-configured keywords for each candidate observation, for each country. To further refine and filter the information, News Aggregator uses the structured semantic knowledge available in Wikidata. Country relevance is determined based on the mentions of the local people and organizations found using Wikidata Query Service [7]. News Aggregator also uses Wikidata for source discovery. The end-users (local experts) then choose relevant observations for scenario planning.

Capturing Domain Knowledge. The Mind Maps and The Customizations components store 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) [2], 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 [1] to capture directed graphs of risk drivers and business implications, with edges having hidden semantics of pairwise cause and effect. The Customizations are elicited using generated questionnaires that request country-specific likelihood and impact for selected cause and effect pairs. Due to Customizations, the same observations may generate different scenarios in different countries.

Reasoning With Incomplete Knowledge. The AI Planning component applies plan recognition and top-k planning techniques to reason with incomplete knowledge and generate scenarios [4–6]. 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 (rather than achieving a predefined goal), and such that each plan ends with a business implication. The scenarios, presented as generated text summaries and graphically, are then reviewed and refined by scenario planning teams.

3 SPA Deployment

The deployed system has over 30 active teams of users, 700 scenarios generated, and is processing over 50,000 social media messages per hour. There are 230 risk drivers and business implications and 382 edges between them in the lightweight Mind Map representation of the domain knowledge used for scenario generation.

4 Acknowledgements

We thank Fang Yuan and Finn McCoole at IBM for providing the domain expertise. We thank Nagui Halim and Edward Shay for their guidance and support. We also thank our LAS collaborators. 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

1. FreeMind: <http://freemind.sourceforge.net>
2. McDermott, D.V.: PDDL — The Planning Domain Definition Language. Tech. Rep. TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control (1998)
3. Peterson, G.D., Cumming, G.S., Carpenter, S.R.: Scenario planning: a tool for conservation in an uncertain world. *Conservation biology* 17(2), 358–366 (2003)
4. Sohrabi, S., Riabov, A., Udrea, O.: Plan recognition as planning revisited. In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)* (2016)
5. Sohrabi, S., Riabov, A., Udrea, O.: State projection via AI planning. In: *Proceedings of National Conference on Artificial Intelligence (AAAI)* (2017)
6. Sohrabi, S., Riabov, A., Udrea, O., Hassanzadeh, O.: Finding diverse high-quality plans for hypothesis generation. In: *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)* (2016)
7. Wikidata: <https://query.wikidata.org/>