Adaptive Modeling for Risk-Aware Decision Making

Sandhya Saisubramanian
College of Information and Computer Sciences
University of Massachusetts
Amherst, MA 01003, USA
saisubramanian@cs.umass.edu

Introduction

Decision making under uncertainty is a core capability of an autonomous agent. In model-based reasoning, agents reason about the effects of their own actions and the events in the environment based on the model parameters. A cornerstone for long-term autonomy with safety guarantees is risk-aware decision making. The importance of accounting for risks in AI systems is attracting growing interest (Kulić and Croft 2005; Zilberstein 2015; Amodei et al. 2016). A *risk-aware* model fully accounts for a known set of risks in the environment, with respect to the problem under consideration. The process of decision making using such a model is risk-aware decision making. Formulating risk-aware models is critical for robust reasoning under uncertainty, since the impact of using less accurate models may be catastrophic in extreme cases due to overly optimistic view of problems. The risk awareness in a model can be increased by improving the model fidelity and thereby, the solution quality.

In practice, agents often operate based on models that do not fully capture the underlying dynamics of the decision process, leading to sub-optimal solutions that may be risky. There are three key reasons for this behavior: (i) the notion and representation of risks is problem-specific, which makes it challenging to design a model that is risk-aware across multiple problems; (ii) improving the model fidelity leads to significant increase in computational complexity of decision making; and (iii) complete knowledge of the system is unavailable ahead of the planning process which the agent discovers during plan execution. Existing techniques that minimize risks in planning under uncertainty, using Markov Decision Processes (MDPs), assume the models are risk-aware and optimize a risk-sensitive criterion or maximize goal reachability in the presence of unavoidable risky states (Kolobov, Mausam, and Weld 2012; Trevizan, Teichteil-Königsbuch, and Thiébaux 2017).

This thesis specifically focuses on investigating the question of how and when to improve risk awareness and model fidelity, taking into account the complexity of planning under uncertainty. I explore risk-aware decision making under two broad settings: planning with full information and planning with incomplete information. In the first setting, the agent encounters risks primarily due to sub-optimal action selection that arises when planning using a simplified model, which trades the solution quality for computational gains. The second setting deals with the impact of missing information and the associated risks. I propose adaptive modeling that allows an agent to balance the trade-off between model simplicity and risk awareness, for different notions of risks, while remaining computationally tractable. Outside of AI, models with varying fidelity have been employed for optimizing the trade-off between multiple parameters in computational fluid dynamics, combustion modeling in turbulent flows, and video coding (Pellegrini et al. 2016; Wu et al. 2015; Lu and Zhang 2011).

The overall scope of my thesis is to design decision theoretic mechanisms for risk-aware decision making by considering the impact of the risks and computational complexity of accounting for them, focusing on the following objectives:

- Identify and characterize different forms of risks, in the context of planning under uncertainty;
- Develop metrics for risk assessment, which quantifies the impact of the risks, and risk awareness, which is measured with respect to the environment;
- Develop a systematic approach for adjusting model fidelity that would guide the automated construction of good risk-aware models, which can balance the trade-off between model simplicity and risk awareness;
- Design robust and scalable solution techniques for solving this class of problems, along with identifying conditions under which bounded-optimality can be achieved; and
- Perform a comprehensive evaluation and comparison of the approach with the existing techniques.

Current Progress

My recent works (Saisubramanian, Zilberstein, and Shenoy 2018; Saisubramanian and Zilberstein 2018) formalize risk-aware decision making for problems modeled as a Stochastic Shortest Path (SSP) problem. Given the computational complexity of solving large SSPs optimally, there has been considerable interest in developing efficient approximations, such as reduced models, that trade solution quality for computational gains. Reduced models simplify the problem by partially or completely ignoring uncertainty; I consider reduced models in which the number of outcomes per action is reduced relative to the original model. If the ignored outcomes are "risky" — significantly affecting the cost of reaching a goal — then the reduced model is too optimistic and not risk-aware. Existing reduced model techniques have focused on the reduction of planning time (Yoon, Fern, and Givan 2007; Keller and Eyerich 2011), but they do not address explicitly the associated risks. While the risk awareness can be improved by always using the full model for planning, it defeats the purpose of using reduced models. The key question addressed in this work is how to formulate risk-aware reduced models that balance this trade-off.

Intuitively, the trade-off between model simplicity and risk awareness can be optimized by identifying when to use a simple model and when to use a more informed model. I have designed a planning paradigm that employs a *portfolio of reduced models with cost adjustments*, that improves risk awareness when using a reduced model by selectively improving its fidelity in certain states and a means to account for the ignored details by adjusting the actions costs in the reduced model. This adaptive outcome selection improves the risk awareness by providing the flexibility to create reduced models with different levels of detail using a portfolio, without compromising the runtime gains of using a reduced model. The cost adjustments also act as a heuristic for choosing outcome selection principles in a portfolio. My results on various domains have demonstrated the potential benefits of this framework in formulating risk-aware reduced models, without significantly affecting the runtime gains of using a reduced model.

Currently, I am extending this framework to handle the negative side effects of planning with reduced models. The reduction in reachable states, as a result of using reduced models, leads to replanning to the goal states when the agent encounters states without an action during plan execution. However, this may not be always acceptable and the agent is expected to be better prepared to handle risky situations and avoid replanning in risky states. For example, when using reduced models for autonomous navigation, it is unsafe to replan in the middle of an intersection if the agent misses a turn or an exit. Therefore, replanning in states that are unsafe for deliberation is considered as a *negative side effect* of planning with reduced models, since the agent will never replan when using the complete model for planning optimally. An agent should thus only replan in states which the user has approved as safe for replanning. The challenge is to plan with reduced models, but somehow account for the possibility of encountering risky situations. Preliminary experimental results indicate that portfolio of reduced models achieves significant reduction in negative side effects, without compromising the run time gains and expected cost.

Research Plan

In the coming years, I intend to address several open questions related to risk-aware decision making. First, I plan to investigate and design adaptive modeling approaches that work well in settings with incomplete information and across multiple forms of risks. This would potentially involve abstractions and querying an oracle to estimate the missing information. Second, I intend to explore other notions of risks. I have recently started to work on minimizing negative side effects, which I will continue to explore and aim to devise mechanisms to automatically classify side effects from negative side effects. This would help broaden the scope of my framework, while simultaneously opening up new challenges. This extension will also include more comprehensive empirical analysis of risk awareness of a model. Finally, I plan to employ machine learning techniques such as regression and decision stump to automatically identify the features that represent risks, which would help in risk assessment and automated generation risk-aware models.

References

Amodei, D.; Olah, C.; Steinhardt, J.; Christiano, P.; Schulman, J.; and Mané, D. 2016. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*.

Keller, T., and Eyerich, P. 2011. A polynomial all outcome determinization for probabilistic planning. In *Proc. of the 21st International Conference on Automated Planning and Scheduling*.

Kolobov, A.; Mausam; and Weld, D. 2012. A theory of goal-oriented MDPs with dead ends. In *Conference on Uncertainty in Artificial Intelligence*, 438–447.

Kulić, D., and Croft, E. A. 2005. Safe planning for human-robot interaction. *Journal of Field Robotics* 22:383–396.

Lu, S.-P., and Zhang, S.-H. 2011. Saliency-based fidelity adaptation preprocessing for video coding. *Journal of Computer Science and Technology* 26(1):195–202.

Pellegrini, R.; Iemma, U.; Leotardi, C.; Campana, E. F.; and Diez, M. 2016. Multi-fidelity adaptive global metamodel of expensive computer simulations. In *Proc.of the IEEE Congress on Evolutionary Computation (CEC)*.

Saisubramanian, S., and Zilberstein, S. 2018. Safe reduced models for probabilistic planning. In *ICML/IJCAI/AAMAS Workshop on Planning and Learning (PAL)*.

Saisubramanian, S.; Zilberstein, S.; and Shenoy, P. 2018. Planning using a portfolio of reduced models. In *Proc. of the 17th International Conference on Autonomous Agents and MultiAgent Systems*.

Trevizan, F.; Teichteil-Königsbuch, F.; and Thiébaux, S. 2017. Efficient solutions for stochastic shortest path problems with dead ends. In *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence*.

Wu, H.; See, Y. C.; Wang, Q.; and Ihme, M. 2015. A fidelity adaptive modeling framework for combustion systems based on model trust-region. In *Proc. of the 53rd AIAA Aerospace Sciences Meeting*.

Yoon, S.; Fern, A.; and Givan, R. 2007. FF-Replan: A baseline for probabilistic planning. In *Proc. of the 17th International Conference on Automated Planning and Scheduling*.

Zilberstein, S. 2015. Building strong semi-autonomous systems. In *Proc. of the 29th AAAI Conference on Artificial Intelligence*.