

Stochastic Constraint Programming as Reinforcement Learning

Steven D. Prestwich¹ S. Armagan Tarim² Roberto Rossi³

¹Insight Centre for Data Analytics, University College Cork, Ireland

²Department of Management, Cankaya University, Ankara, Turkey

³University of Edinburgh Business School, Edinburgh, UK

s.prestwich@ucc.ie, at@cankaya.edu.tr, roberto.rossi@ed.ac.uk

Stochastic Constraint Programming (SCP) [1] is an extension of Constraint Programming (CP) used for modelling and solving problems involving constraints and uncertainty. SCP inherits excellent modelling abilities and filtering algorithms from CP.

Reinforcement Learning (RL) [2] extends Dynamic Programming to large stochastic problems, but is problem-specific and has no generic solvers.

We propose [3] a hybrid combining the scalability of RL with the modelling and constraint filtering methods of CP. We implemented a simple RL algorithm in a CP solver, and obtained a novel algorithm for solving SCP problems. We showed that this RL/CP hybrid can find high-quality solutions to hard problems.

We believe that exploiting Machine Learning methods is a good direction for SCP research, to make it a practical tool for real-world problems. This work should also be of interest from an RL perspective. Firstly, implementing RL algorithms in a CP solver enables the user to perform rapid prototyping of RL methods for new problems. For example, simply by specifying a different filtering algorithm for a global constraint we obtain a new RL solver. Secondly, we now have an RL solver for an interesting class of problem (SCP problems). There are no general-purpose RL solvers available because, like Dynamic Programming, RL is a problem-specific approach. Thirdly, allowing the use of constraint filtering methods in RL potentially boosts its ability to solve tightly-constrained problems.

In future work we shall show that our approach extends to multistage SCP problems using different state aggregation techniques — we have preliminary results on an inventory control problem.

References

- [1] T. Walsh. Stochastic Constraint Programming. 15th European Conference on Artificial Intelligence, 2002, pp. 111–115.
- [2] R. S. Sutton, A. G. Barto. Reinforcement Learning: An Introduction. MIT Press, Cambridge, MA, 1998.
- [3] Prestwich, S. D., Tarim, S. A., Rossi, R., Stochastic Constraint Programming as Reinforcement Learning, arXiv:1704.07183