Abstract: Empirical Model Learning, and the integration of machine learning models in CP

Tias Guns, Michele Lombardi, Michela Milano, Pascal Van Hentenryck

An overarching goal in Artificial Intelligence (AI) is to develop general reasoning techniques. Constraint Programming (CP) is one such technique, and it excels at solving discrete satisfaction and optimization problems involving complex constraints such as those appearing in routing, scheduling, and planning.

However, as the problems move to model complex human and physical behavior, as well as large-scale systems and infrastructures, it also becomes increasingly harder for experts to design CP models with the appropriate accuracy and efficiency.

Empirical Model Learning [1, 2] has been proposed as a way to overcome this. The key idea is to learn a function linking decision variables with observable quantities of interest. The observations may come from the real system we are modeling or from a simulator. Machine learning techniques are then used to extract an empirical model that accurately generalizes the observations. This empirical model is then embedded in the combinatorial optimization specification; and because the input to the model involves decision variables, it has to be evaluated during search. CP is especially well-suited for EML since it easily supports the integration of heterogeneous techniques through global constraints.

This is a promising research direction to enhance the capability of CP to model real problems, however, research into this integration has only just been started. We have identified the following four main challenges:

Effective propagation of ML models. Many machine learning models exist, from decision trees to neural networks and Bayesian networks. Their applicability for EML depends on the ability to effectively embed the *inference* over the ML model in a CP solver. Traditional issues such as incrementality and the tradeoff between the strength of propagation and its performance take another dimension in this context.

Trade-off between model accuracy and optimization efficiency. An accurate ML model (e.g. a deep leraning network) may lead to weaker propagation, and therefore lower quality solutions. On the other hand, an ML model with a strong propagation algorithm (e.g. a decision tree) may be less accurate from a learning point of view. This cost-benefit relation has barely been studied, yet should be core to choosing which ML models to use.

Representative observations. ML learns from observations, but a lack of observations in a part of the feature space can lead to false predictions. A good representative sample can avoid learning spurious correlations that are far from causal. Moreover, an open question is whether to sample the whole observation space or just the feasible region. Selecting a good set of samples has been studied in the field of Design of Experiments [3]. Some functions may be so complex that one needs to interleave sampling/learning with optimization, which is studied for unconstrained problems in surrogate optimization [4]. An interesting challenge would be to adapt these techniques to discrete, constrained optimization problems.

Uncertainty information. A number of ML models can provide some confidence estimation on their predictions. To increase robustness, we may wish to use this information in the optimization to optimize expection or various risk measures.

Answering these challenges will require a good understanding of both the machine learning and optimization methods, and is of interest to both research domains in AI. Advances on this topic will allow one to tackle novel combinatorial problems that are out of scope of today's constraint-programming technology. This will increase the reach of AI problems for which CP brings significant benefits.

References

[1] Andrea Bartolini, Michele Lombardi, Michela Milano, and Luca Benini. Neuron constraints to model complex real-world problems. In *Principles and Practice of Constraint Programming, CP*, pages 115–129, 2011.

[2] Michele Lombardi, Michela Milano, and Andrea Bartolini. Empirical decision model learning. *Artificial Intelligence*, pages –, 2016.

[3] D. C. Montgomery. *Design And Analysis Of Experiments*. Wiley New York, 1997.

[4] N V. Queipo, R. T. Haftka, W. Shyy, T. Goel, R. Vaidyanathan, and P. K. Tucker. Surrogate-based analysis and optimization. *Progress In Aerospace Sciences*, 41(1):1–28, 2005.