Towards Machine-Generated Algorithms

Lars Kotthoff
University of Wyoming
larsko@uwyo.edu

Much progress has been made in recent years leveraging machine learning to improve solving hard combinatorial problems. While machine learning is used in almost all parts of the solving process, it only improves the human-designed solver and does not replace it. Designing and implementing state-of-the-art solvers for combinatorial problems is still beyond the reach of current machine learning approaches.

In contrast, instances of hard combinatorial problems are relatively well understood, thanks to decades of effort to analyze them. Features extracted from such problem instances provide the information machine learning needs to assist the solving process. The solvers and solving algorithms themselves are treated as black boxes, without any knowledge of their inner workings.

I challenge the community to do better than that. I will present recent advances in automatically extracting features from solvers for hard combinatorial problems [1]. Such features further enable machine learning approaches to speed up the solving process, but also allow to analyze algorithms in much the same way as problem instances have been analyzed for decades. This opens up new research directions, enabling a better understanding of why some algorithms perform well while others do not, and how to change and synthesize algorithms with better performance solving combinatorial problems than existing ones.

Can we extract features from algorithms that allow to characterize their performance, and learn what makes an algorithm fast? Can machine learning identify characteristics of the joint algorithm-instance space that help to solve challenging problems faster? Can we generate algorithms given their feature values, much like problem instances are generated today? Answering such under-explored questions will eventually allow automatically generate algorithms tailored to the problems they are to solve.

The time for this presentation could range from a short talk (5 minutes) to a conference talk (20-25 minutes).

References

[1] Damir Pulatov, Marie Anastacio, Lars Kotthoff, and Holger Hoos. Opening the Black Box: Automated Software Analysis for Algorithm Selection. In *First Conference on Automated Machine Learning (Main Track)*, 2022.