Instance-specific algorithm configuration

Yuri Malitsky¹

Published online: 12 September 2015
© Springer Science+Business Media New York 2015

Abstract

When developing a new heuristic or complete algorithm, we frequently face the problem of choice. There may be multiple heuristics that we can employ, different types of inference mechanisms, various settings for the learning rate, or a multitude of neighborhoods to choose from. Furthermore, the way in which the choices we make affect one another is not readily known. The task of making these choices is referred to as algorithm configuration. At the same time, even for instances from the same problem domain, very different algorithmic strategies may prove effective. The task of selecting the most appropriate solver for the provided problem instance is known as algorithm selection. While it is possible to approach both these tasks manually, the potential explosion of the search space necessitates that need for automated methodologies to provide consistently beneficial results. This thesis addresses this exact problem, showing how to effectively perform both algorithm selection and configuration simultaneously, and leading to approaches that have dominated international competitions.

School: Brown University

Supervisors:

Meinolf Sellmann

Graduated: Wednesday, May 16, 2012

Link to full text: http://www.a4cp.org/sites/default/files/yuri_malitsky_-_instance-specific_algorithm_configuration.pdf

Notes:

Runner-Up for the 2014 ACP Doctoral Thesis Award.

✉ Yuri Malitsky
yuri.malitsky@gmail.com

¹ IBM Thomas J. Watson Research Center, Yorktown, NY, USA