Per-Class Algorithm Selection for Black-Box Optimisation

Koen van der Blom[1] Carola Doerr[1]

[1]Sorbonne Université, CNRS, LIP6, France, koen.vdblom@lip6.fr, carola.doerr@lip6.fr

Keywords: Algorithm selection, black-box optimisation, feature-free algorithm selection, per-class algorithm selection.

Algorithm selection works on the premise that no single algorithm is the best on each individual instance of a problem. With this in mind, the aim is to choose the most suitable algorithm for a problem, out of a set of available algorithms. Most algorithm selection approaches use precomputed features to predict the most suitable algorithm. However, this both requires experts to design high-quality features, and incurs a cost to compute these features for each problem instance before being able to make a prediction.

Feature-free algorithm selection approaches use models like artificial neural networks to predict the most suitable algorithm, with the raw problem instance description as input. Although this avoids the need for expert designed features and their computation, a specific description for each problem instance is still needed.

In black-box optimisation neither features, nor a raw problem instance description may be available a priori. Information about the specific problem instance may only be obtained by evaluating possible solutions to it, which may be computationally expensive.

Recognising that we cannot make predictions without any cost on a per-instance basis, we take a step back and propose an algorithm selector that makes predictions on a per-class basis. Our proposed method predicts the most suitable algorithm for a class of continuous black-box optimisation problems, using only the basic properties of this class. Specifically, we take into account only the number of decision variables, and the available evaluation budget. This requires neither instance-specific information, nor any expensive computation when applying the selector.