Algorithmic configuration by learning and optimization

Claudia D'Ambrosio[1], Antonio Frangioni[2],
Gabriele Iommazzo [1,2], Leo Liberti[1]

[1]DASCIM Team, CNRS LIX, Ecole Polytechnique, Palaiseau, France

[2]Operations Research Group, Dip. di Informatica, Università di Pisa, Italy

Dec 5, 2019. Pisa
THESES

The Algorithm Configuration Problem [Rice, 1976]

Running a target algorithm depends on:

- **features** of the input (*"problem space"*);
- **controls** by the user, i.e. the algorithmic parameters (*"algorithm space"*).

Algorithm Configuration problem (ACP): given a problem input $\boldsymbol{f}$ and a target algorithm $\boldsymbol{\mathcal{A}}$, find the controls $\boldsymbol{c}$ optimizing the performance $\boldsymbol{p}_{\boldsymbol{\mathcal{A}}}$ of the target algorithm.

Default algorithm configuration

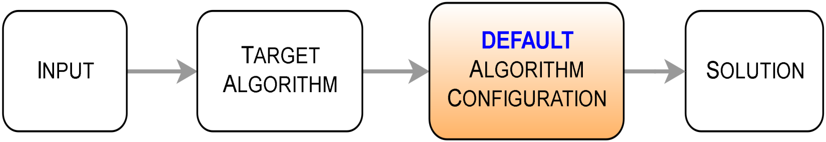

Figure: No tuning: let the algorithm roam free

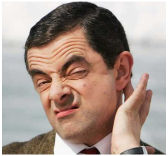

Default algorithm configuration

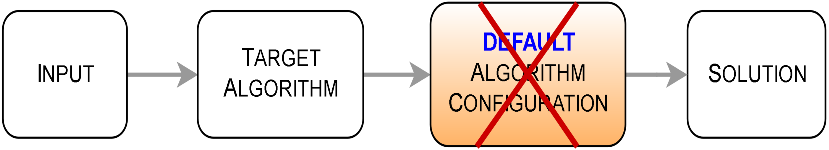

Figure: No tuning: let the algorithm roam free

WARNING: Default configurations often suboptimal!

Manual algorithm configuration

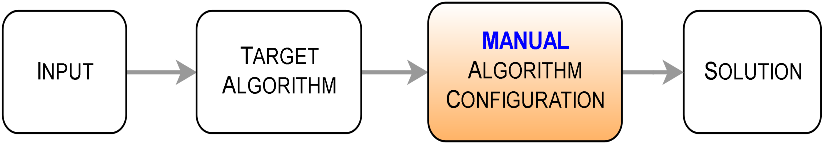

Figure: Manual tuning: steer the algorithm

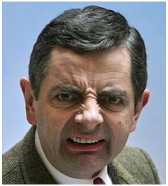

Manual algorithm configuration

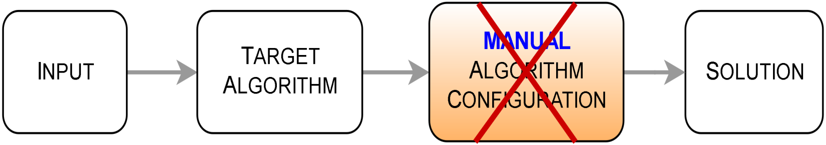

Figure: Manual tuning: steer the algorithm

WARNING: Manual tuning is time-consuming + trial & error process

Automatic algorithm configuration [Hutter et al., 2009]

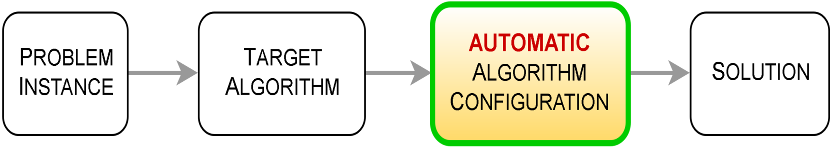

Figure: Automatic tuning: take educated guesses

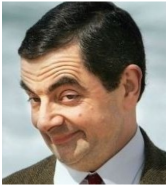

Our approach to the ACP: "learn then optimize"

1. Performance Map Learning Phase (PMLP)

Supervised Machine Learning (ML) predictor [Shalev-Shwartz and Ben-David, 2014] learns coefficients θ^* for prediction model $\bar{p}_{\mathcal{A}}(f, c, \theta)$ of $p_{\mathcal{A}}(f, c)$.

2. Configuration Space Search Problem (CSSP)

Optimization problem encoding properties of PMLP in closed form (by mathematical programming). For given f and θ CSSP(f, θ) finds algorithmic configuration with minimum estimated performance:

$$\mathrm{CSSP}(f, \theta) \equiv \min_{c \in \mathcal{C}_{\mathcal{A}}} \quad \bar{p}_{\mathcal{A}}(f, c, \theta).$$

The challenges

You may be enticed by:

- in CSSP: **ML $\neg$black-box** $\implies$ CSSP solved with known optimization techniques
- **learning-based optimization**: **learn** components of an optimization problem **then** find clever way to **optimize over huge combinatorial sets**
- **learning under hard constraints**: **have learning** methodology **enforce** dependency and compatibility **constraints** on the controls

Possible research applications...

. . . encompass but are not limited to:

- **medicine and bio-informatics**: find best treatment for a specific patient
- **optimization solvers**: find best solver parameters for specific problem instance
- **machine learning**: hyperparameter tuning
- . . .

Possible future work

Algorithmic ideas to develop:

- **test** SVR kernels and other **learning methodologies** (trees, ad-hoc neural networks, ...) for PMLP
- explore ways to **embed downstream optimization problem into the learning** phase
- find **approximations of the CSSP** for **faster** solutions
- ...

Hutter, F., Hoos, H. H., Leyton-Brown, K., and Stützle, T. (2009).
ParamILS: An automatic algorithm configuration framework.
J. Artif. Int. Res., 36(1):267–306.

Rice, J. R. (1976).
The algorithm selection problem.
Advances in Computers, 15:65–118.

Shalev-Shwartz, S. and Ben-David, S. (2014).
Understanding Machine Learning: From Theory to Algorithms.
Cambridge University Press.