

Project Proposal: Model-Based Optimization of Strategy Schedules

Filip Bártek Martin Suda

Czech Technical University in Prague

September 7, 2021



Task definition

Input		Example
Parameterized target algorithm	A	Vampire (an ATP)
Parameter configuration space	Θ	Vampire strategy space
Instance set	Π	FOL problems from TPTP
Cost metric	c	PAR10 ¹

Output	Task	Tool
Configuration	Algorithm configuration	SMAC
Portfolio	Portfolio optimization	Hydra
Schedule	Schedule optimization	HOS-ML

¹Penalized average runtime with a penalty factor of 10

Sequential Model-based Algorithm Configuration (SMAC)²

Input: Target algorithm A ; parameter configuration space Θ ;
instance set Π ; cost metric c

Output: Configuration $\theta_{inc} \in \Theta$

- 1: $[\mathbf{R}, \theta_{inc}] \leftarrow \text{Initialize}(\Theta, \Pi)$
- 2: $\triangleright R = \{([\theta_1, \pi_1], o_1), \dots, ([\theta_n, \pi_n], o_n)\} \subseteq [\Theta \times \Pi] \times \mathbb{R}$
- 3: **repeat**
- 4: $\mathcal{M} \leftarrow \text{FitModel}(\mathbf{R})$
- 5: $\Theta_{new} \leftarrow \text{SelectConfigurations}(\mathcal{M}, \theta_{inc}, \Theta, \mathbf{R})$
- 6: $[\mathbf{R}, \theta_{inc}] \leftarrow \text{Intensify}(A, \Theta_{new}, \theta_{inc}, \mathbf{R}, \Pi, c)$
- 7: **until** total time budget for configuration exhausted
- 8: **return** θ_{inc}

²Hutter et al. [2011]

Empirical performance model (EPM) \mathcal{M}

Input: Configuration θ

Output: Predictive statistics of $c(\theta, \Pi)$

- Mean μ_θ
- Variance σ_θ^2

Architecture: Random forest (10 regression trees)

Prediction:

- 1 Each tree predicts $c(\theta, \pi)$ for configuration θ and instance π (represented by a feature vector)
- 2 Aggregation across instances: mean
- 3 Aggregation across trees: mean and variance

Batch evaluation for multiple configurations and instances is cheap.

Candidate configuration selection (*SelectConfigurations*)

Positive improvement: $I(\boldsymbol{\theta}) = \max\{0, c(\boldsymbol{\theta}_{inc}, \Pi) - c(\boldsymbol{\theta}, \Pi)\}$

Maximize the *expected positive improvement* (EI) over $\boldsymbol{\theta}_{inc}$:

- 10 000 random configurations
- Multi-start local search
 - Initial population: 10 configurations with the highest EI in \mathbf{R}
 - Randomized one-exchange neighborhood

Interleave:

- Configurations with high EI
- Random configurations

Intensification

Input: Target algorithm A ; sequence of candidate configurations Θ_{new} ; incumbent configuration θ_{inc} ; set of target algorithm runs R ; instance set Π ; cost metric c

Output: Set of runs R ; new incumbent configuration θ_{inc}

```

1: for  $\theta_{new} \in \Theta_{new}$  do
2:   Evaluate  $\theta_{inc}$  on a random instance from  $\Pi \setminus \Pi[\theta_{inc}]$ 
3:   loop
4:     Evaluate  $\theta_{new}$  on some subset of  $\Pi[\theta_{inc}] \setminus \Pi[\theta_{new}]$ 
5:     if  $c(\theta_{new}, \Pi[\theta_{new}]) > c(\theta_{inc}, \Pi[\theta_{new}])$  then break
6:     end if
7:     if  $\Pi[\theta_{new}] = \Pi[\theta_{inc}]$  then  $\theta_{inc} \leftarrow \theta_{new}$ ; break
8:     end if
9:   end loop
10: end for
11: return  $[R, \theta_{inc}]$ 

```

Hydra³ with SMAC

Input: Target algorithm A ; parameter configuration space Θ ;
instance set Π ; cost metric c ; number of iterations K

Output: Portfolio P

- 1: $c_1 \leftarrow c$
- 2: $\theta_1 \leftarrow \text{SMAC}(A, \Theta, \Pi, c_1)$
- 3: $P_1 \leftarrow \{\theta_1\}$
- 4: **for** $k \leftarrow 2 \dots K$ **do**
- 5: Define c_k : $c_k(\theta, \pi) = \min(c(\theta, \pi), \min_{\hat{\theta} \in P_{k-1}} c(\hat{\theta}, \pi))$
- 6: $\theta_k \leftarrow \text{SMAC}(A, \Theta, \Pi, c_k)$
- 7: $P_k \leftarrow P_{k-1} \cup \{\theta_k\}$
- 8: **end for**
- 9: **return** P_K

³Xu et al. [2010]

Issues with SMAC3

- Bugs
- If the instance set is large, the iterated local search for configuration takes a lot of time

Initial experiment

- Instance set: 1000 FOL problems from TPTP 7.5.0
- Instance features: 32 syntactic features
- Target algorithm: Vampire
 - Parameter configuration space: 113 parameters
- Training computation budget: 4 CPUs, 8 hours

Results

Solver	Configurations	Time limit [s]	Timeouts ⁴
Hydra	1	8	518
Hydra	2	4	503
Hydra	4	2	479
Hydra	8	1	465
Vampire ⁵	1	8	422

⁴Number of timeouts on 1000 training problems

⁵vampire --mode casc

Future

- Optimize a configuration schedule
 - Interleave optimization of configurations and schedule
 - Modify Hydra: Construct schedule iteratively
 - Modify SMAC: Replace the incumbent configuration with an incumbent schedule
- Analyze the EPM, namely the parameter importance
- Better instance features
 - Runtime statistics of probing runs
- More problem domains (SMT, HOL)

Thank you for your attention!

Filip Bártek

filip.bartek@cvut.cz

Acknowledgement: This work is supported by the Czech Science Foundation project no. 20-06390Y (JUNIOR grant), and the Grant Agency of the Czech Technical University in Prague, grant no. SGS20/215/OHK3/3T/37.

Terminology

ATP	SMAC	Symbol
prover	target algorithm	A
strategy	parameter configuration	θ
strategy space	parameter configuration space (PCS)	Θ
problem	instance	π
problem set	instance set	Π
	runtime budget	B
	portfolio	P
schedule		f
	EPM	\mathcal{M}

Tasks

- Per-instance algorithm selection (AS)
- Algorithm configuration (AC)
 - One configuration for all instances
 - Per-instance
- Portfolio
- Scheduling

Parameter configuration space example

Vampire call example

```
vampire --age_weight_ratio 1:1 --naming 8
```

PCS

```
age_weight_ratio:log_ratio real [-10.0, 3.0] [0.0]  
  
naming:special categorical {regular, 0} [regular]  
naming integer [2, 64] [8] log  
naming | naming:special = regular
```

Experiment setup

Sequential Model-based Algorithm Configuration (SMAC) command line options

```
smac  
--acq_opt_challengers 1000  
--sls_n_steps_plateau_walk 2  
--mode Hydra  
--n_optimizers 4
```

Scenario

```
run_obj = runtime  
overall_obj = par10  
deterministic = true  
initial_incumbent = DEFAULT
```


SMAC summary

- Empirical performance model (EPM): Random forest
 - Input: Configuration
 - Output: Aggregate cost (for example PAR10)
- Optimizes the strategy for multiple instances
- Incumbent and candidate are always compared using only the instances on which they have both been run
- Parameter types: categorical, integer, real
- Acquisition function: EI
- Adaptive capping: a candidate is capped at 1.2 times the runtime of the incumbent

Related optimization tools

- ParamILS (Hutter et al. [2009]): algorithm configuration, model-free, only categorical and ordinal parameters
- BliStr (Urban [2015]): combines instance clustering with ParamILS
- HOS-ML (Holden and Korovin [2021]): combines instance clustering with SMAC, proposes schedules
- MaLeS (Kühlwein and Urban [2015]): dynamically schedules strategies with good predicted performance
- CPHydra (Bridge et al. [2012]): given an instance, produces a schedule of solvers

References I

Derek Bridge, Eoin O'Mahony, and Barry O'Sullivan. Case-based reasoning for autonomous constraint solving. In *Autonomous Search*, volume 9783642214349, pages 73–95. Springer-Verlag Berlin Heidelberg, oct 2012. ISBN 9783642214349. doi: 10.1007/978-3-642-21434-9_4. URL https://link.springer.com/chapter/10.1007/978-3-642-21434-9_4.

Edvard K. Holden and Konstantin Korovin. Heterogeneous Heuristic Optimisation and Scheduling for First-Order Theorem Proving. In *CICM 2021*, pages 107–123, Cham, jul 2021. Springer, Cham. doi: 10.1007/978-3-030-81097-9_8. URL https://link.springer.com/chapter/10.1007/978-3-030-81097-9_8.

References II

- F. Hutter, H. H. Hoos, K. Leyton-Brown, and T. Stuetzle.
ParamLLS: An automatic algorithm configuration framework.
Journal of Artificial Intelligence Research, 36:267–306, oct 2009.
ISSN 1076-9757. doi: 10.1613/JAIR.2861. URL
<https://jair.org/index.php/jair/article/view/10628>.
- Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown.
Sequential model-based optimization for general algorithm
configuration. In *Lecture Notes in Computer Science (including
subseries Lecture Notes in Artificial Intelligence and Lecture
Notes in Bioinformatics)*, volume 6683 LNCS, pages 507–523.
Springer, Berlin, Heidelberg, 2011. ISBN 9783642255656. doi:
10.1007/978-3-642-25566-3_40. URL [https://link.
springer.com/chapter/10.1007/978-3-642-25566-3_40](https://link.springer.com/chapter/10.1007/978-3-642-25566-3_40).

References III

- Laura Kovács and Andrei Voronkov. First-order theorem proving and Vampire. In Natasha Sharygina and Helmut Veith, editors, *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, volume 8044 of *Lecture Notes in Computer Science*, pages 1–35. Springer, 2013. ISBN 978-3-642-39798-1. doi: 10.1007/978-3-642-39799-8_1.
- Daniel Kühlwein and Josef Urban. MaLeS: A framework for automatic tuning of automated theorem provers. *Journal of Automated Reasoning*, 55(2):91–116, August 2015. ISSN 15730670. doi: 10.1007/s10817-015-9329-1. URL <https://github.com/dkuehlwein/males>.
- M. Lindauer, H. Hoos, F. Hutter, and T. Schaub. AutoFolio: An automatically configured algorithm selector. *Journal of Artificial Intelligence Research*, 53:745–778, 2015.

References IV

- Marius Lindauer, Katharina Eggensperger, Matthias Feurer, Stefan Falkner, André Biedenkapp, and Frank Hutter. SMAC v3: Algorithm configuration in Python. <https://github.com/automl/SMAC3>, 2017.
- Geoff Sutcliffe. The TPTP problem library and associated infrastructure. *Journal of Automated Reasoning*, 59(4), December 2017. ISSN 0168-7433. doi: 10.1007/s10817-017-9407-7.
- Josef Urban. BliStr: The blind strategymaker. In *EPiC Series in Computer Science*, volume 36, pages 312–303. EasyChair, January 2015. doi: 10.29007/8n7m. URL <http://www.tptp.org/CASC/J6/Design.html#CompetitionDivisions>.

References V

Lin Xu, Holger H. Hoos, and Kevin Leyton-Brown. Hydra: Automatically configuring algorithms for portfolio-based selection. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI'10*, page 210–216. AAAI Press, 2010.