

Algorithm Configuration

Challenges, Methods and Perspectives

Marius Lindauer André Biedenkapp



What can you expect?

- Content** Overview over well-established and new research directions
- Structure** Combination of research insights and practical recommendations
- Hands-On** Some simple coding examples for algorithm configuration s.t. you can directly start to play with it
- Outlook** to new research directions and open challenges

Algorithms have Parameters!

Evolutionary Algorithms population size, selection strategies, mutation operators, recombination operators, step size, probabilities, ...

Algorithms have Parameters!

Evolutionary Algorithms population size, selection strategies, mutation operators, recombination operators, step size, probabilities,
...

Satisfiability Solving Search strategies (local vs. tree-based), variable selection heuristic, restart probability, deletion heuristics,
...

Algorithms have Parameters!

Evolutionary Algorithms population size, selection strategies, mutation operators, recombination operators, step size, probabilities,
...

Satisfiability Solving Search strategies (local vs. tree-based), variable selection heuristic, restart probability, deletion heuristics,
...

Machine Learning model class (linear, non-linear, tree, ...), model complexity, regularization, preprocessing, data augmentation,
...

Algorithms have Parameters!

Evolutionary Algorithms population size, selection strategies, mutation operators, recombination operators, step size, probabilities,
...

Satisfiability Solving Search strategies (local vs. tree-based), variable selection heuristic, restart probability, deletion heuristics,
...

Machine Learning model class (linear, non-linear, tree, ...), model complexity, regularization, preprocessing, data augmentation,
...

Deep Learning network depth, network width, network architecture, activation function, optimizer, learning rate, momentum,
...

Algorithms have Parameters!

Evolutionary Algorithms population size, selection strategies, mutation operators, recombination operators, step size, probabilities,
...

Satisfiability Solving Search strategies (local vs. tree-based), variable selection heuristic, restart probability, deletion heuristics,
...

Machine Learning model class (linear, non-linear, tree, ...), model complexity, regularization, preprocessing, data augmentation,
...

Deep Learning network depth, network width, network architecture, activation function, optimizer, learning rate, momentum,
...

Remark: The ML folks call an algorithm's parameters hyperparameters.

But wait! (I)

🤔 But *my* algorithm has no parameters!

But wait! (I)



But *my* algorithm has no parameters!

- ▶ Correct: There are some algorithms that do not have parameters
- ▶ However: Algorithm parameters are simply not exposed to the user in some cases.
- ▶ Programming by Optimization could be a paradigm to address these hidden parameters [Hoos 2012]

But wait! (II)



Can't I simply use its default settings?

But wait! (II)



Can't I simply use its default settings?

Domain	Default vs. Optimized
Answer Set Solving [Gebser et al. 2011]	up to $14\times$ speedup
AI Planning [Vallati et al. 2013]	up to $40\times$ speedup
Mixed Integer Programming [Hutter et al. 2010]	up to $52\times$ speedup
Satisfiability Solving [Hutter et al. 2017]	up to $3000\times$ speedup

But wait! (II)



Can't I simply use its default settings?

Domain	Default vs. Optimized
Answer Set Solving [Gebser et al. 2011]	up to $14\times$ speedup
AI Planning [Vallati et al. 2013]	up to $40\times$ speedup
Mixed Integer Programming [Hutter et al. 2010]	up to $52\times$ speedup
Satisfiability Solving [Hutter et al. 2017]	up to $3000\times$ speedup
Minimum Vertex Cover [Wagner et al. 2017]	up to 9% absolute impr.
Machine Learning [Feuer et al. 2015]	up to 35% absolute impr.
Deep Learning [Zimmer et al. 2020]	up to 49% absolute impr.

But wait! (II)



Can't I simply use its default settings?

Domain	Default vs. Optimized
Answer Set Solving [Gebser et al. 2011]	up to $14\times$ speedup
AI Planning [Vallati et al. 2013]	up to $40\times$ speedup
Mixed Integer Programming [Hutter et al. 2010]	up to $52\times$ speedup
Satisfiability Solving [Hutter et al. 2017]	up to $3000\times$ speedup
Minimum Vertex Cover [Wagner et al. 2017]	up to 9% absolute impr.
Machine Learning [Feuer et al. 2015]	up to 35% absolute impr.
Deep Learning [Zimmer et al. 2020]	up to 49% absolute impr.

Remark: Default configurations can nevertheless *sometimes* be very strong.

But wait! (III)



I don't want to bother with parameters,
because manual parameter tuning is ...

- 1 time-consuming
- 2 tedious
- 3 error-prone

But wait! (III)



I don't want to bother with parameters,
because manual parameter tuning is ...

- 1 time-consuming
- 2 tedious
- 3 error-prone



Can't I simply use these AI algorithms everyone is talking about?

But wait! (III)



I don't want to bother with parameters,
because manual parameter tuning is ...

- ❶ time-consuming
- ❷ tedious
- ❸ error-prone



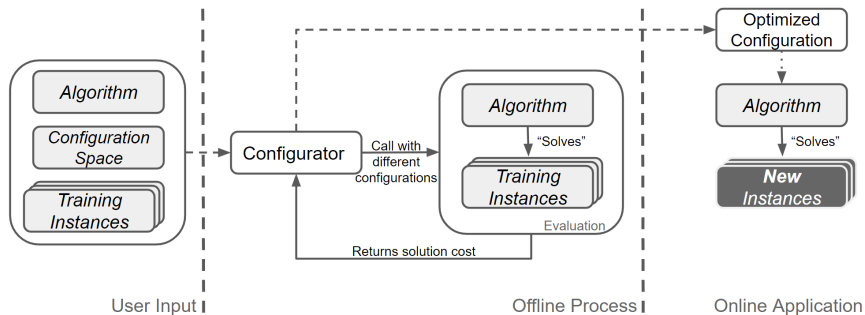
Can't I simply use these AI algorithms everyone is talking about?

↪ Sure! Let's use automatic algorithm configuration!

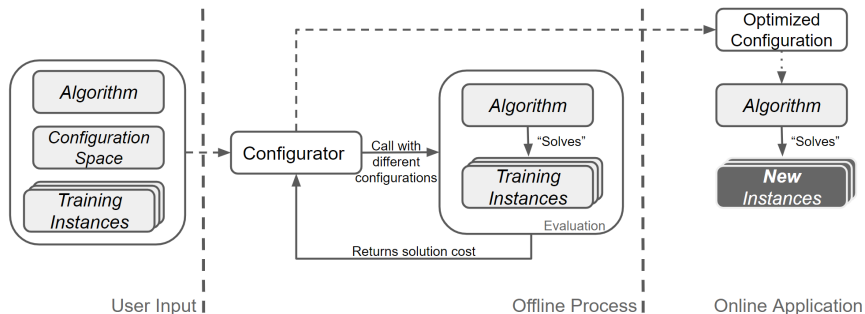
Outline

- 1 What is Algorithm Configuration?
- 2 How can we solve AC?
- 3 What to watch out for?
- 4 Is Solving AC hard?
- 5 Can we learn how to solve AC?
- 6 Top-Down or Bottom-Up Censoring?
- 7 Evolving from Static to Dynamic
- 8 What's next?

Algorithm Configuration Visualized




Algorithm Configuration Visualized



Offline tuning phase: Tune parameters on some training instances

Online application phase: Apply the found configuration to new instances

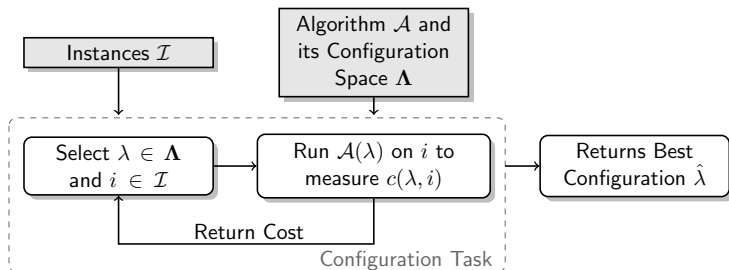
Algorithm Configuration – in More Detail



Definition

Given a parameterized algorithm \mathcal{A} with possible (hyper-)parameter settings Λ ,

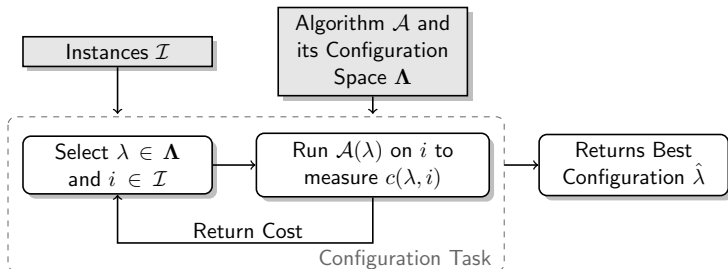
Algorithm Configuration – in More Detail



Definition

Given a parameterized algorithm \mathcal{A} with possible (hyper-)parameter settings Λ , a set of training problem instances \mathcal{I} ,

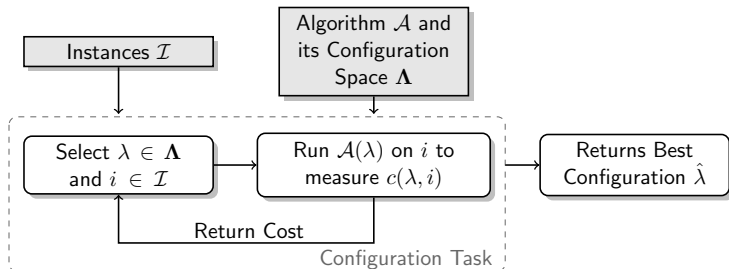
Algorithm Configuration – in More Detail



Definition

Given a parameterized algorithm \mathcal{A} with possible (hyper-)parameter settings Λ , a set of training problem instances \mathcal{I} , and a cost metric $c : \Lambda \times \mathcal{I} \rightarrow \mathbb{R}$,

Algorithm Configuration – in More Detail



Definition

Given a parameterized algorithm \mathcal{A} with possible (hyper-)parameter settings Λ , a set of training problem instances \mathcal{I} , and a cost metric $c : \Lambda \times \mathcal{I} \rightarrow \mathbb{R}$, the algorithm configuration problem is to find a (hyper-)parameter configuration $\lambda^* \in \Lambda$ that minimizes c across the instances in \mathcal{I} .

Challenges of Algorithm Configuration

- Structured high-dimensional parameter space
 - ▶ categorical vs. continuous parameters
 - ▶ conditionals between parameters
 - ▶ between 5 and > 300 parameters
 - ▶ low effective dimensionality

Challenges of Algorithm Configuration

- Structured high-dimensional parameter space
 - ▶ categorical vs. continuous parameters
 - ▶ conditionals between parameters
 - ▶ between 5 and > 300 parameters
 - ▶ low effective dimensionality
- Stochastic optimization
 - ▶ Randomized algorithms: optimization across various seeds
 - ▶ Distribution of benchmark instances (often wide range of hardness)
 - ▶ Subsumes so-called *multi-armed bandit problem*

Challenges of Algorithm Configuration

- Structured high-dimensional parameter space
 - ▶ categorical vs. continuous parameters
 - ▶ conditionals between parameters
 - ▶ between 5 and > 300 parameters
 - ▶ low effective dimensionality
- Stochastic optimization
 - ▶ Randomized algorithms: optimization across various seeds
 - ▶ Distribution of benchmark instances (often wide range of hardness)
 - ▶ Subsumes so-called *multi-armed bandit problem*
- Generalization across instances
 - ▶ apply algorithm configuration to **homogeneous** instance sets
 - ▶ Instance sets can also be **heterogeneous**,
i.e., no single configuration performs well on all instances
 - ~> combination of algorithm configuration and selection

Challenges of Algorithm Configuration

- Structured high-dimensional parameter space

- ▶ categorical vs. continuous parameters
- ▶ conditionals between parameters
- ▶ between 5 and > 300 parameters
- ▶ low effective dimensionality

- Stochastic optimization

- ▶ Randomized algorithms: optimization across various seeds
- ▶ Distribution of benchmark instances (often wide range of hardness)
- ▶ Subsumes so-called *multi-armed bandit problem*

- Generalization across instances

- ▶ apply algorithm configuration to **homogeneous** instance sets
- ▶ Instance sets can also be **heterogeneous**,
i.e., no single configuration performs well on all instances
- ↪ combination of algorithm configuration and selection

↪ Hyperparameter optimization is a subproblem of algorithm configuration
[Eggensperger et al. 2019]

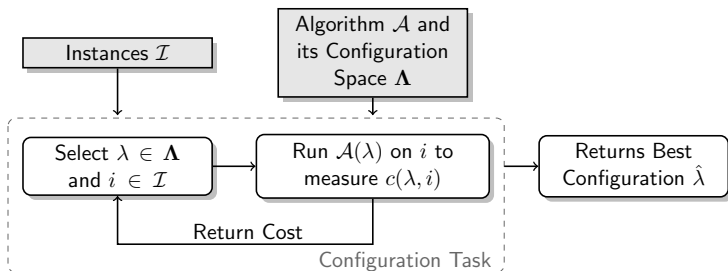
Let's configure some algorithms! (I)

Use this [Google CoLab link](#) to configure some example algorithms.

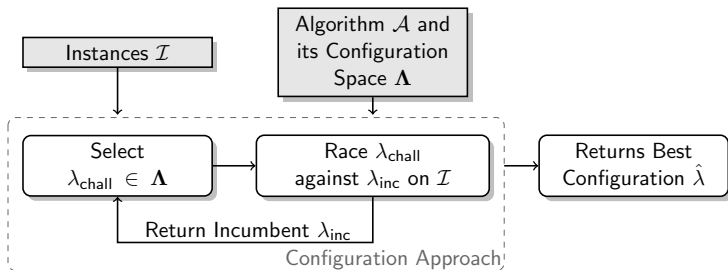
Outline

- 1 What is Algorithm Configuration?
- 2 How can we solve AC?**
- 3 What to watch out for?
- 4 Is Solving AC hard?
- 5 Can we learn how to solve AC?
- 6 Top-Down or Bottom-Up Censoring?
- 7 Evolving from Static to Dynamic
- 8 What's next?

The AC Problem Again




The Two Main Components



Sampling of Challengers: Iterated Local Search

[Hutter et al. 2007]



Sampling of Challengers: Bayesian Optimization

[Hutter et al. 2011]


Require: Search space Λ , cost function c , acquisition function u , maximal number of function evaluations T

Result : Best configuration $\hat{\lambda}$ (according to \mathcal{D} or \hat{c})

```
1 Initialize data  $\mathcal{D}^{(0)}$  with initial observations
2 for  $t = 1$  to  $T$  do
3   Fit predictive model  $\hat{c}^{(t)}$  on  $\mathcal{D}^{(t-1)}$ 
4   Select next query point:  $\lambda_t \in \arg \max_{\lambda \in \Lambda} u(\lambda; \mathcal{D}^{(t-1)}, \hat{c}^{(t)})$ 
5   Query  $c(\lambda)$ ;
6   Update data:  $\mathcal{D}^{(t)} \leftarrow \mathcal{D}^{(t-1)} \cup \{\langle \lambda, c(\lambda) \rangle\}$ 
```

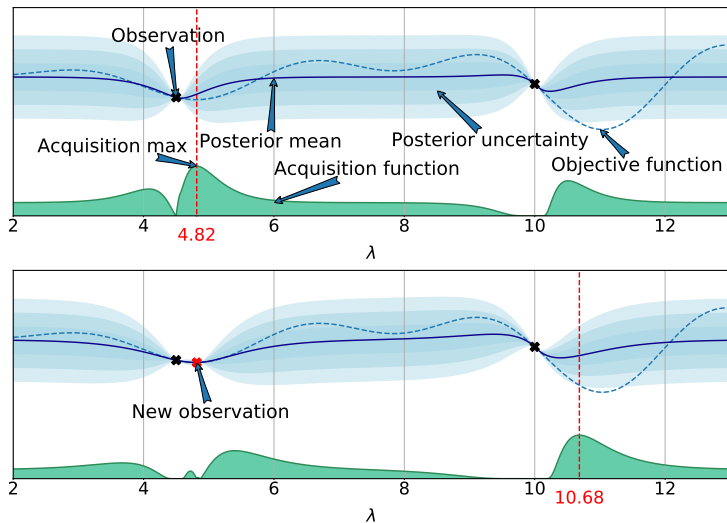
Sampling of Challengers: Bayesian Optimization

[Hutter et al. 2011]



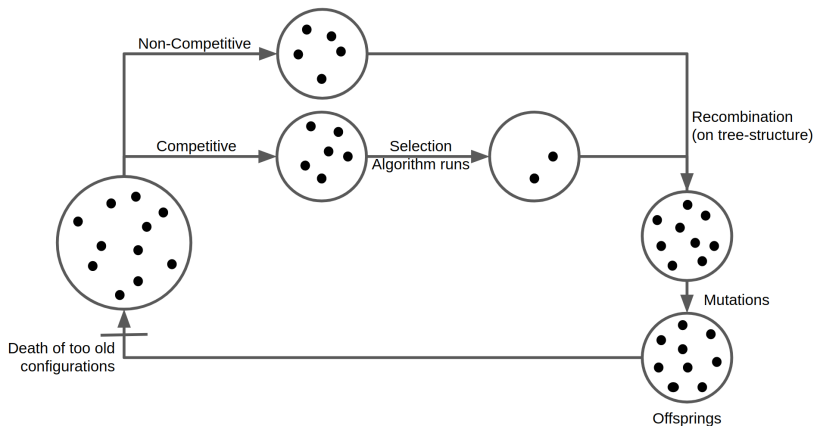
Sampling of Challengers: Bayesian Optimization

[Hutter et al. 2011]




Sampling of Challengers: Genetic Algorithms

[Ansótegui et al. 2009]



Sampling of Challengers: Genetic Algorithms

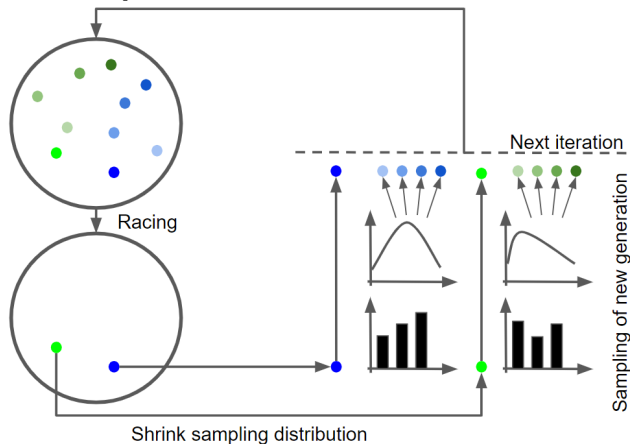
[Ansótegui et al. 2009]



- can be extended by using a surrogate model [Ansótegui et al. 2015]

Sampling of Challengers: Est. of Distributions

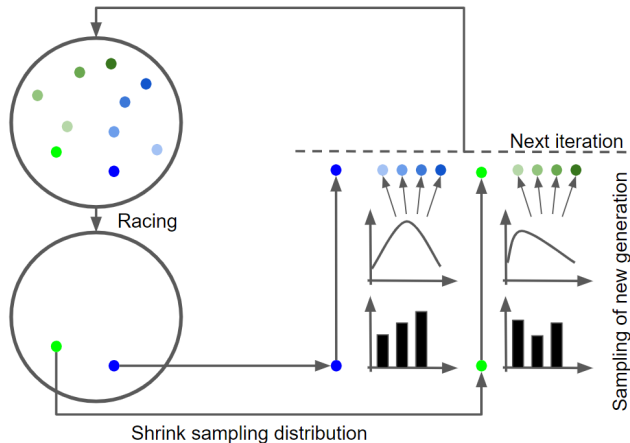
[López-Ibáñez et al. 2016]



- the sampling distribution around well-performing configurations gets more peaked after each iteration

Sampling of Challengers: Est. of Distributions

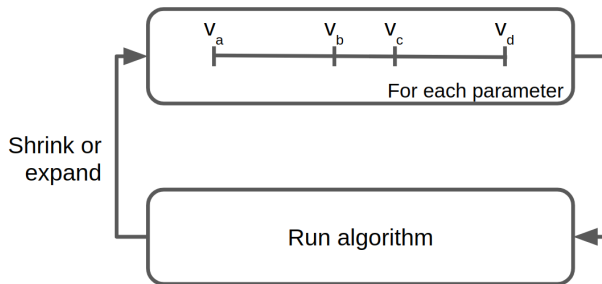
[López-Ibáñez et al. 2016]



- the sampling distribution around well-performing configurations gets more peaked after each iteration
- can be extended by using a surrogate model [Pérez Cáceres et al. 2017]


Sampling of Challengers: Golden Search

[Pushak and Hoos. 2020]



Sampling of Challengers: Golden Search

[Pushak and Hoos. 2020]



Assumes:

- uni-modal & convex search space
- independent parameters

~> more often applicable than one might expect [Pushak and Hoos. 2020]

Let's configure some algorithms! (II)

Use this [Google CoLab link](#) to configure some example algorithms.

Racing of Challengers: Main Idea

Observations

Obs 1 : We cannot afford to evaluate all configurations $\lambda \in \Lambda$ on all instances \mathcal{I}

Racing of Challengers: Main Idea

Observations

Obs I : We cannot afford to evaluate all configurations $\lambda \in \Lambda$ on all instances \mathcal{I}

Obs II : We do not want to waste time on poor λ

Racing of Challengers: Main Idea

Observations

- Obs I : We cannot afford to evaluate all configurations $\lambda \in \Lambda$ on all instances \mathcal{I}
- Obs II : We do not want to waste time on poor λ
- Obs III : We need enough empirical evidence to distinguish between well performing λ s

Racing of Challengers: Main Idea

Observations

- Obs I : We cannot afford to evaluate all configurations $\lambda \in \Lambda$ on all instances \mathcal{I}
- Obs II : We do not want to waste time on poor λ
- Obs III : We need enough empirical evidence to distinguish between well performing λ s

Idea

- Idea I : Discard poorly performing λ s early on
- Idea II : Run promising λ s on many instances

Selection of Challengers: Aggressive Racing

[Hutter et al. 2009]

Race a challenger configuration λ against incumbent configuration $\hat{\lambda}$:

Selection of Challengers: Aggressive Racing

[Hutter et al. 2009]

Race a challenger configuration λ against incumbent configuration $\hat{\lambda}$:

Reject Challenger

Reject λ if

$$\sum_{i \in \mathcal{I}'} c(\lambda, i) > \sum_{i \in \mathcal{I}'} c(\hat{\lambda}, i)$$

where \mathcal{I}' are the instances λ was evaluated on

Selection of Challengers: Aggressive Racing

[Hutter et al. 2009]

Race a challenger configuration λ against incumbent configuration $\hat{\lambda}$:

Reject Challenger

Reject λ if

$$\sum_{i \in \mathcal{I}'} c(\lambda, i) > \sum_{i \in \mathcal{I}'} c(\hat{\lambda}, i)$$

where \mathcal{I}' are the instances λ was evaluated on


Update Incumbent

Update $\hat{\lambda} \rightarrow \lambda$ if

- (i) λ was evaluated on the same instances as $\hat{\lambda}$
(i.e. same evidence level)
- (ii) $\sum_{i \in \mathcal{I}'} c(\lambda, i) \leq \sum_{i \in \mathcal{I}'} c(\hat{\lambda}, i)$

Selection of Challengers: Statistical testing

[López-Ibáñez et al. 2016]




Objective: Decide from a set of configurations $\hat{\Lambda}$ which one(s) is the best

- 1 Evaluate all $\lambda \in \hat{\Lambda}$ on one or more additional instances (i.e., add more empirical evidence)

Selection of Challengers: Statistical testing

[López-Ibáñez et al. 2016]




Objective: Decide from a set of configurations $\hat{\Lambda}$ which one(s) is the best

- 1 Evaluate all $\lambda \in \hat{\Lambda}$ on one or more additional instances (i.e., add more empirical evidence)
- 2 Drop all configurations λ from $\hat{\Lambda}$ that perform statistically worse than the best configuration $\hat{\lambda} \in \hat{\Lambda}$

Selection of Challengers: Statistical testing

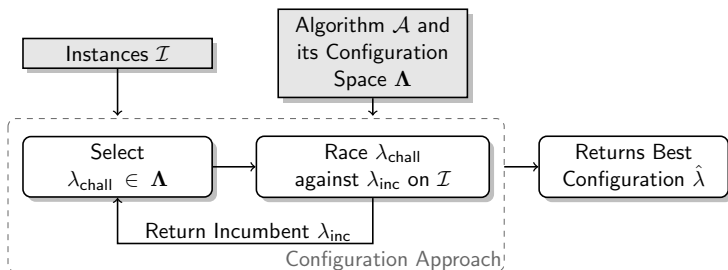
[López-Ibáñez et al. 2016]



Objective: Decide from a set of configurations $\hat{\Lambda}$ which one(s) is the best

- 1 Evaluate all $\lambda \in \hat{\Lambda}$ on one or more additional instances (i.e., add more empirical evidence)
- 2 Drop all configurations λ from $\hat{\Lambda}$ that perform statistically worse than the best configuration $\hat{\lambda} \in \hat{\Lambda}$
- 3 Continue with 1. if budget remains and $|\hat{\Lambda}| > 1$

The Two Main Components



Many systems combining these two components in different ways, incl.
[Hutter et al. 2009, Ansótegui et al. 2009, Hutter et al. 2011, López-Ibáñez et al. 2016, Pushak and Hoos. 2020]

Let's configure some algorithms! (III)

Use this [Google CoLab link](#) to configure some example algorithms.

Outline

- 1 What is Algorithm Configuration?
- 2 How can we solve AC?
- 3 What to watch out for?**
- 4 Is Solving AC hard?
- 5 Can we learn how to solve AC?
- 6 Top-Down or Bottom-Up Censoring?
- 7 Evolving from Static to Dynamic
- 8 What's next?

Pitfalls & Best Practices

The **goals** of automated algorithm configuration include:

Pitfalls & Best Practices

The **goals** of automated algorithm configuration include:

- ➊ reducing the expertise required to use an algorithm

Pitfalls & Best Practices

The **goals** of automated algorithm configuration include:

- ➊ reducing the expertise required to use an algorithm
- ➋ less human-time

Pitfalls & Best Practices

The **goals** of automated algorithm configuration include:

- ➊ reducing the expertise required to use an algorithm
- ➋ less human-time
- ➌ tuning algorithms to the task at hand

Pitfalls & Best Practices

The **goals** of automated algorithm configuration include:

- ➊ reducing the expertise required to use an algorithm
- ➋ less human-time
- ➌ tuning algorithms to the task at hand
- ➍ faster development of algorithms

Pitfalls & Best Practices

The **goals** of automated algorithm configuration include:

- ➊ reducing the expertise required to use an algorithm
- ➋ less human-time
- ➌ tuning algorithms to the task at hand
- ➍ faster development of algorithms
- ➎ facilitating systematic and reproducible research

BUT:

- ➊ algorithm configuration can lead to over-tuning

Pitfalls & Best Practices

The **goals** of automated algorithm configuration include:

- ➊ reducing the expertise required to use an algorithm
- ➋ less human-time
- ➌ tuning algorithms to the task at hand
- ➍ faster development of algorithms
- ➎ facilitating systematic and reproducible research

BUT:

- ➊ algorithm configuration can lead to over-tuning
- ➋ using algorithm configuration requires (at least some) expertise in algorithm configuration

Pitfalls & Best Practices

The **goals** of automated algorithm configuration include:

- ➊ reducing the expertise required to use an algorithm
- ➋ less human-time
- ➌ tuning algorithms to the task at hand
- ➍ faster development of algorithms
- ➎ facilitating systematic and reproducible research

BUT:

- ➊ algorithm configuration can lead to over-tuning
- ➋ using algorithm configuration requires (at least some) expertise in algorithm configuration
- ➌ if done wrong, waste of time and compute resources

Setting up AC

9 Steps to your well-performing algorithm:

- 1 Define your cost metric

Setting up AC

9 Steps to your well-performing algorithm:

- 1 Define your cost metric
- 2 Define instance set

Setting up AC

9 Steps to your well-performing algorithm:

- 1 Define your cost metric
- 2 Define instance set
- 3 Split your instances in training and test

Setting up AC

9 Steps to your well-performing algorithm:

- 1 Define your cost metric
- 2 Define instance set
- 3 Split your instances in training and test
- 4 Define the configuration space

Setting up AC

9 Steps to your well-performing algorithm:

- 1 Define your cost metric
- 2 Define instance set
- 3 Split your instances in training and test
- 4 Define the configuration space
- 5 Choose your preferred configurator

Setting up AC

9 Steps to your well-performing algorithm:

- 1 Define your cost metric
- 2 Define instance set
- 3 Split your instances in training and test
- 4 Define the configuration space
- 5 Choose your preferred configurator
- 6 Implement an interface between your algorithm and the configurator

Setting up AC

9 Steps to your well-performing algorithm:

- 1 Define your cost metric
- 2 Define instance set
- 3 Split your instances in training and test
- 4 Define the configuration space
- 5 Choose your preferred configurator
- 6 Implement an interface between your algorithm and the configurator
- 7 Define resource limitations of your algorithm

Setting up AC

9 Steps to your well-performing algorithm:

- 1 Define your cost metric
- 2 Define instance set
- 3 Split your instances in training and test
- 4 Define the configuration space
- 5 Choose your preferred configurator
- 6 Implement an interface between your algorithm and the configurator
- 7 Define resource limitations of your algorithm
- 8 Run the configurator on your algorithm and the training instances

Setting up AC

9 Steps to your well-performing algorithm:

- 1 Define your cost metric
- 2 Define instance set
- 3 Split your instances in training and test
- 4 Define the configuration space
- 5 Choose your preferred configurator
- 6 Implement an interface between your algorithm and the configurator
- 7 Define resource limitations of your algorithm
- 8 Run the configurator on your algorithm and the training instances
- 9 Validate the eventually returned configuration on your test instances

Pitfall 1: Trust your algorithm

We have encountered algorithms that

- ignored resource limitations
- returned wrong solutions
- even returned negative runtimes

Pitfall 1: Trust your algorithm

We have encountered algorithms that

- ignored resource limitations
- returned wrong solutions
- even returned negative runtimes

Best Practice 1: Never trust your algorithm

Explicitly check and use external software to:

- 1 ensure resource limitations
- 2 terminate your algorithm
- 3 verify returned solutions
- 4 measure cost

Pitfall 2: File System

Algorithm configurators ...

- often produce quite a lot of log files (e.g., for each algorithm run)
- are often used on HPC clusters with a shared file system

Pitfall 2: File System

Algorithm configurators ...

- often produce quite a lot of log files (e.g., for each algorithm run)
- are often used on HPC clusters with a shared file system

⇒ It's surprisingly easy to substantially slow down a shared file system

Pitfall 2: File System

Algorithm configurators ...

- often produce quite a lot of log files (e.g., for each algorithm run)
- are often used on HPC clusters with a shared file system

⇒ It's surprisingly easy to substantially slow down a shared file system

Best Practice 2: Don't use the Shared File System

To lower the burden the file system on a HPC cluster:

- design well which files are required and which are not
- use a local (SSD) disc

Pitfall 3: Over-tuning

It's easy to over-tune to different aspects, including:

- training instances
- random seeds
- machine type

Pitfall 3: Over-tuning

It's easy to over-tune to different aspects, including:

- training instances
- random seeds
- machine type

It can be overcome, e.g., by

- using larger instance sets
- tuning on the target hardware

~> Still this is very hard in practice

Pitfall 3: Over-tuning

It's easy to over-tune to different aspects, including:

- training instances
- random seeds
- machine type

It can be overcome, e.g., by

- using larger instance sets
- tuning on the target hardware

↪ Still this is very hard in practice

Best Practice 3: Check for Over-Tuning

Check for over-tuning by validating your final configuration on

- many random seeds
- a large set of unused test instances
- different hardware

More Pitfalls and Best Practices

... can be found in [Eggersperger et al. 2019]

Outline

- 1 What is Algorithm Configuration?
- 2 How can we solve AC?
- 3 What to watch out for?
- 4 Is Solving AC hard?**
- 5 Can we learn how to solve AC?
- 6 Top-Down or Bottom-Up Censoring?
- 7 Evolving from Static to Dynamic
- 8 What's next?

Worst Case

In the worst case, AC problems are very hard:

- ① very hard instances with long evaluation time
- ② interactions between all parameters
- ③ multi-modal with many local optima
- ④ heterogeneous instance sets


Worst Case

In the worst case, AC problems are very hard:

- ① very hard instances with long evaluation time
- ② interactions between all parameters
- ③ multi-modal with many local optima
- ④ heterogeneous instance sets


⇒ However, does this actually apply in practice?

Low Effective Dimensionality




- Algorithms have often between tenths and hundreds of parameters
- Often only less than 10 matter (\rightsquigarrow low effective dimensionality)
[Bergstra and Bengio. 2012, Hutter et al. 2013, Hutter et al. 2014, Fawcett and Hoos. 2015, Biedenkapp et al. 2017]

Low Effective Dimensionality



- Algorithms have often between tenths and hundreds of parameters
- Often only less than 10 matter (\rightsquigarrow low effective dimensionality)
[Bergstra and Bengio. 2012, Hutter et al. 2013, Hutter et al. 2014, Fawcett and Hoos. 2015, Biedenkapp et al. 2017]
- Importance of parameter changes depending on instance set
[Bergstra and Bengio. 2012, Biedenkapp et al. 2018]


Uni-Modal? Convex? Smooth? (I)



- Intuition: parameters are either set to too high or too low
- We can analyse it for example by using fANOVA
[Hutter et al. 2014, Biedenkapp et al. 2018]

PPO on cartpole [Lindauer et al. 2019]

Uni-Modal? Convex? Smooth? (II) [Pushak and Hoos. 2018]



Observations:

- On a single instances, the landscape might not be uni-modal, convex or smooth
- On average across many instances, the landscapes are often uni-modal, convex and smooth (at least on solvers for discrete combinatorial problems)

Homogeneity vs. Heterogeneity (I) [Schneider and Hoos. 2012]

Algorithm configurators...

- use racing to not evaluate each configuration on all instances

Homogeneity vs. Heterogeneity (I) [Schneider and Hoos. 2012]

Algorithm configurators...

- use racing to not evaluate each configuration on all instances
- misleading on subsets of instances if they are heterogeneous overall

Homogeneity vs. Heterogeneity (I) [Schneider and Hoos. 2012]

Algorithm configurators...

- use racing to not evaluate each configuration on all instances
- misleading on subsets of instances if they are heterogeneous overall

⇒ returned configurations often perform worse than default configurations in the validation phase

Homogeneity vs. Heterogeneity (II)

[Schneider and Hoos. 2012]

Rule of Thumb


- An instance set is homogeneous if all instances encode the same task (or problem)
- An homogeneous instance can have instances of different hardness

Homogeneity vs. Heterogeneity (II)

[Schneider and Hoos. 2012]



Rule of Thumb

- An instance set is homogeneous if all instances encode the same task (or problem)
- An homogeneous instance can have instances of different hardness



← smooth transition \rightsquigarrow homogeneous

← check board pattern \rightsquigarrow heterogeneous




Homogeneity vs. Heterogeneity (II)

[Schneider and Hoos. 2012]

Rule of Thumb



- An instance set is homogeneous if all instances encode the same task (or problem)
- An homogeneous instance can have instances of different hardness



← smooth transition \rightsquigarrow homogeneous

← check board pattern \rightsquigarrow heterogeneous

\rightsquigarrow Rule of thumb does not always apply!



Per-Instance Algorithm Configuration: Hydra

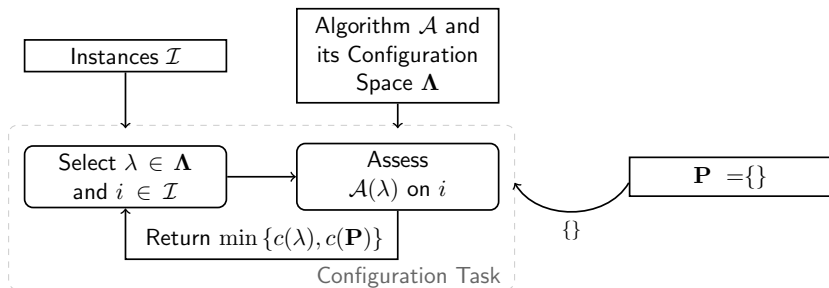
[Xu et al. 2010; Xu et al. 2011]

Idea: Find a configuration for each homogeneous subset of instances

Per-Instance Algorithm Configuration: Hydra

[Xu et al. 2010; Xu et al. 2011]

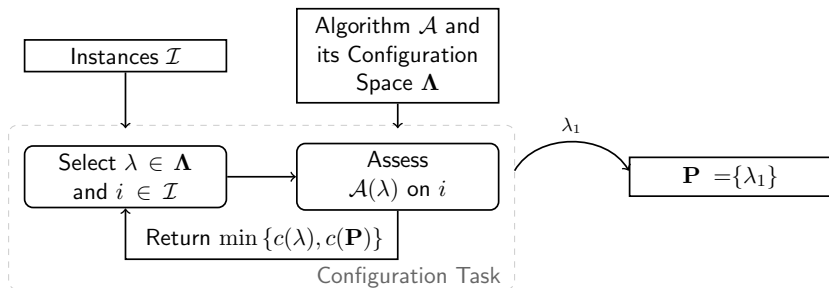
Idea: Find a configuration for each homogeneous subset of instances



Per-Instance Algorithm Configuration: Hydra

[Xu et al. 2010; Xu et al. 2011]

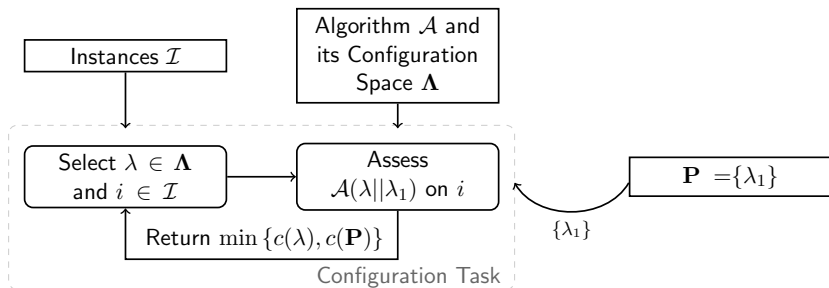
Idea: Find a configuration for each homogeneous subset of instances



Per-Instance Algorithm Configuration: Hydra

[Xu et al. 2010; Xu et al. 2011]

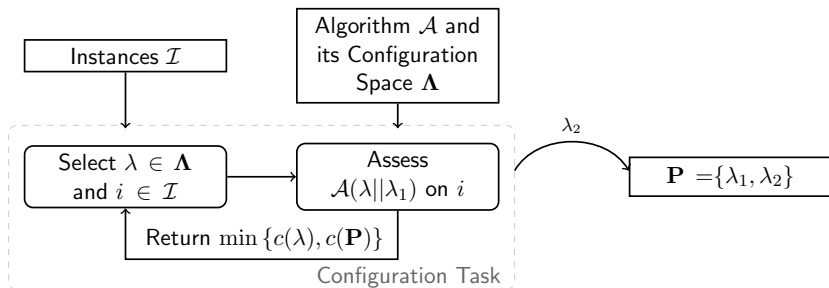
Idea: Find a configuration for each homogeneous subset of instances



Per-Instance Algorithm Configuration: Hydra

[Xu et al. 2010; Xu et al. 2011]

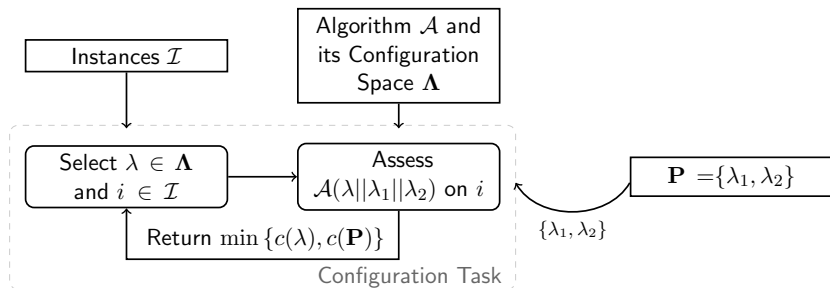
Idea: Find a configuration for each homogeneous subset of instances



Per-Instance Algorithm Configuration: Hydra

[Xu et al. 2010; Xu et al. 2011]

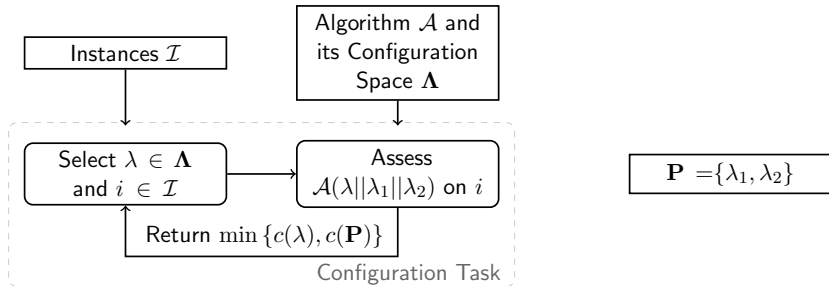
Idea: Find a configuration for each homogeneous subset of instances



Per-Instance Algorithm Configuration: Hydra

[Xu et al. 2010; Xu et al. 2011]

Idea: Find a configuration for each homogeneous subset of instances



Similar ideas:

- ISAC: 1. Cluster instances; 2. configure on each instance cluster [Kadioglu et al. 2010, Ansótegui et al. 2016]
- Combination of iterative configuration and clustering [Liu et al. 2019]

Outline


- 1 What is Algorithm Configuration?
- 2 How can we solve AC?
- 3 What to watch out for?
- 4 Is Solving AC hard?
- 5 Can we learn how to solve AC?**
- 6 Top-Down or Bottom-Up Censoring?
- 7 Evolving from Static to Dynamic
- 8 What's next?

Solving AC more than once?

Assumption: We tune the parameters of the same algorithms again and again on different instance sets.


Solving AC more than once?

Assumption: We tune the parameters of the same algorithms again and again on different instance sets.



Solving AC more than once?

Assumption: We tune the parameters of the same algorithms again and again on different instance sets.




↪ Learn on the first k instance sets,
how to optimize on the $k + 1$ instance set.
(Common idea in the optimization community.)

Idea 1: Try the best ones [Lindauer and Hutter. 2018]

Idea: Take the best configurations of previous runs and try them as initial design on new instances.


Idea 1: Try the best ones [Lindauer and Hutter. 2018]

Idea: Take the best configurations of previous runs and try them as initial design on new instances.



Idea 1: Try the best ones [Lindauer and Hutter. 2018]


Idea: Take the best configurations of previous runs and try them as initial design on new instances.



Problem: If we have too many previous AC runs, the initial design will be too expensive.

Idea 1: Try the best ones [Lindauer and Hutter. 2018]

Idea: Take the best configurations of previous runs and try them as initial design on new instances.



Problem: If we have too many previous AC runs, the initial design will be too expensive.

Idea: Select complementary configurations (Λ') across all previous instance sets ($\bigcup_j \mathcal{I}_j$):

$$\sum_{i \in \bigcup_j \mathcal{I}_j} \min_{\lambda \in \Lambda'} c(\lambda, i)$$

Idea 2: Combine Surrogate Models


[Lindauer and Hutter. 2018]

Idea: Weighted combination of already trained surrogate models.

Idea 2: Combine Surrogate Models

[Lindauer and Hutter. 2018]


Idea: Weighted combination of already trained surrogate models.



Idea 2: Combine Surrogate Models

[Lindauer and Hutter. 2018]

Idea: Weighted combination of already trained surrogate models.



$$\hat{c}(\lambda, i) = w_0 + w_c \cdot \hat{c}_c(\lambda, i) + \sum_j w_j \cdot \hat{c}_j(\lambda, i)$$

↪ Fit linear combination wrt \mathbf{w} on hold-out set of the current obs.

Insights from Warmstarting [Lindauer and Hutter. 2018]

- ① If landscape changed too much, AC should be able to recover.
 - ▶ both methods can do that

Insights from Warmstarting [Lindauer and Hutter. 2018]

- ❶ If landscape changed too much, AC should be able to recover.
 - ▶ both methods can do that
- ❷ Weights of combined surrogate models, prefer previous surrogates first, and later on the surrogate fitted on the current data.

Insights from Warmstarting [Lindauer and Hutter. 2018]

- ① If landscape changed too much, AC should be able to recover.
 - ▶ both methods can do that
- ② Weights of combined surrogate models, prefer previous surrogates first, and later on the surrogate fitted on the current data.
- ③ Combining both ideas lead to the best speedups.

Insights from Warmstarting [Lindauer and Hutter. 2018]


- ❶ If landscape changed too much, AC should be able to recover.
 - ▶ both methods can do that
- ❷ Weights of combined surrogate models, prefer previous surrogates first, and later on the surrogate fitted on the current data.
- ❸ Combining both ideas lead to the best speedups.
- ❹ We sped up SMAC by a factor of 4.3 on average and up to a factor of 165.

Outline

- 1 What is Algorithm Configuration?
- 2 How can we solve AC?
- 3 What to watch out for?
- 4 Is Solving AC hard?
- 5 Can we learn how to solve AC?
- 6 Top-Down or Bottom-Up Censoring?**
- 7 Evolving from Static to Dynamic
- 8 What's next?

Running until the End?


- Assumption: we would like to minimize runtime of algorithms



⇒ overall invested time: 450 sec

Top-Down Capping [Hutter et al. 2009]

- Assumption: we would like to minimize runtime of algorithms




⇒ overall invested time: 275 sec (instead of 450 sec)

Bottom-Up Procrastination

[Kleinberg et al. 2017, Kleinberg et al. 2019]

- Assumption: we would like to minimize runtime of algorithms




- ~> overall invested time: 225 sec (instead of 450 and 275 sec)
- Generalize this idea to evaluate configurations on pairs of random seeds and instances
- ~> performance guarantees of AC methods

Bottom-Up Procrastination

[Kleinberg et al. 2017, Kleinberg et al. 2019]

- Assumption: we would like to minimize runtime of algorithms




- ~> overall invested time: 225 sec (instead of 450 and 275 sec)
- Generalize this idea to evaluate configurations on pairs of random seeds and instances
- ~> performance guarantees of AC methods

Bottom-Up Procrastination

[Kleinberg et al. 2017, Kleinberg et al. 2019]

- Assumption: we would like to minimize runtime of algorithms

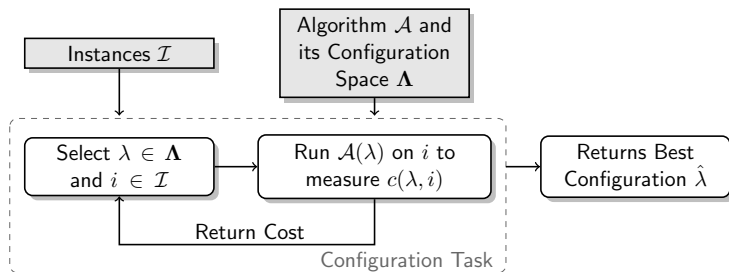


- ~> overall invested time: 225 sec (instead of 450 and 275 sec)
- Generalize this idea to evaluate configurations on pairs of random seeds and instances
- ~> performance guarantees of AC methods


Outline

- 1 What is Algorithm Configuration?
- 2 How can we solve AC?
- 3 What to watch out for?
- 4 Is Solving AC hard?
- 5 Can we learn how to solve AC?
- 6 Top-Down or Bottom-Up Censoring?
- 7 Evolving from Static to Dynamic**
- 8 What's next?

Algorithm Configuration Recap



Algorithm Configuration Recap



- Important observation: Many algorithms are iterative by design
- ↪ fixed configurations are not optimal for each iteration

Algorithms have Dynamic Parameters!

Evolutionary Algorithms population size, selection strategies, mutation operators, recombination operators, step size, probabilities, ...

Algorithms have Dynamic Parameters!

Evolutionary Algorithms **population size**, selection strategies, mutation operators, recombination operators, **step size**, probabilities, ...

Algorithms have Dynamic Parameters!

Evolutionary Algorithms **population size**, selection strategies, mutation operators, recombination operators, **step size**, probabilities,
...

Satisfiability Solving Search strategies (local vs. tree-based), variable selection heuristic, restart probability, deletion heuristics,
...

Algorithms have Dynamic Parameters!

Evolutionary Algorithms **population size**, selection strategies, mutation operators, recombination operators, **step size**, probabilities,
...

Satisfiability Solving Search strategies (local vs. tree-based), **variable selection heuristic**, **restart probability**, **deletion heuristic**,
...

Algorithms have Dynamic Parameters!

Evolutionary Algorithms **population size**, selection strategies, mutation operators, recombination operators, **step size**, probabilities, ...

Satisfiability Solving Search strategies (local vs. tree-based), **variable selection heuristic**, **restart probability**, **deletion heuristic**, ...

Machine Learning model class (linear, non-linear, tree, ...), model complexity, regularization, preprocessing, data augmentation, ...

Algorithms have Dynamic Parameters!

Evolutionary Algorithms **population size**, selection strategies, mutation operators, recombination operators, **step size**, probabilities, ...

Satisfiability Solving Search strategies (local vs. tree-based), **variable selection heuristic**, **restart probability**, **deletion heuristic**, ...

Machine Learning model class (linear, non-linear, tree, ...), model complexity, **regularization**, preprocessing, data augmentation, ...

Algorithms have Dynamic Parameters!

Evolutionary Algorithms population size, selection strategies, mutation operators, recombination operators, step size, probabilities, ...

Satisfiability Solving Search strategies (local vs. tree-based), variable selection heuristic, restart probability, deletion heuristic, ...

Machine Learning model class (linear, non-linear, tree, ...), model complexity, regularization, preprocessing, data augmentation, ...

Deep Learning network depth, network width, network architecture, activation function, optimizer, learning rate, momentum, ...

Algorithms have Dynamic Parameters!

Evolutionary Algorithms **population size**, selection strategies, mutation operators, recombination operators, **step size**, probabilities,
...

Satisfiability Solving Search strategies (local vs. tree-based), **variable selection heuristic**, **restart probability**, **deletion heuristic**,
...

Machine Learning model class (linear, non-linear, tree, ...), model complexity, **regularization**, preprocessing, data augmentation,
...

Deep Learning network depth, network width, network architecture, activation function, optimizer, **learning rate**, **momentum**,
...

Algorithms have Dynamic Parameters!

Evolutionary Algorithms¹ population size, selection strategies, mutation operators, recombination operators, step size, probabilities,

...

Satisfiability Solving Search strategies (local vs. tree-based), variable selection heuristic, restart probability, deletion heuristic,

...

Machine Learning model class (linear, non-linear, tree, ...), model complexity, regularization, preprocessing, data augmentation,

...

Deep Learning network depth, network width, network architecture, activation function, optimizer, learning rate, momentum,

...

¹Gregor Papa: "Dynamic Parameter Changing During the Run to Speed Up Evolutionary Algorithms"

What Can We Do? (I)



Manually design heuristics to adapt the parameter online.

- ▶ Requires substantial expert knowledge
- ~> specialized heuristics for specific domains
- ▶ Very time-consuming

What Can We Do? (I)



Manually design heuristics to adapt the parameter online.

- ▶ Requires substantial expert knowledge
- ~> specialized heuristics for specific domains
- ▶ Very time-consuming



Use Algorithm Configuration/Selection to choose a heuristic.

- ▶ Limited approach
- ~> does not make use of information during the algorithms execution

What Can We Do? (II)


- 🤔 Learn to configure from scratch
 - ▶ Requires access to the algorithms internal statistics
 - ▶ Data-driven approach

What Can We Do? (II)

- 🤔 Learn to configure from scratch
 - ▶ Requires access to the algorithms internal statistics
 - ▶ Data-driven approach
- 🤔 Warmstart from expert knowledge
 - ▶ Potentially more sample efficient
 - ▶ Which expert should we learn from?

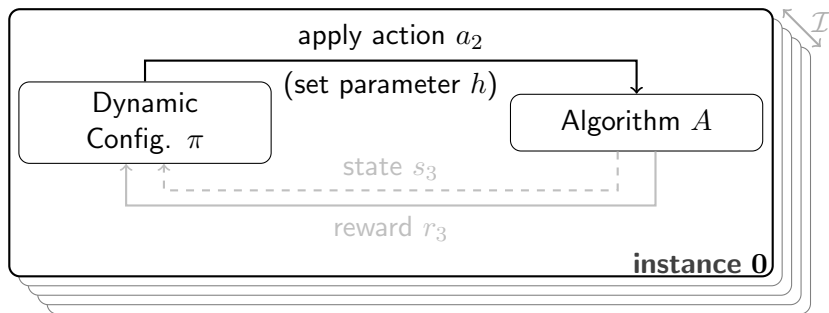
Dynamic Algorithm Configuration

[Biedenkapp et al. 2020]



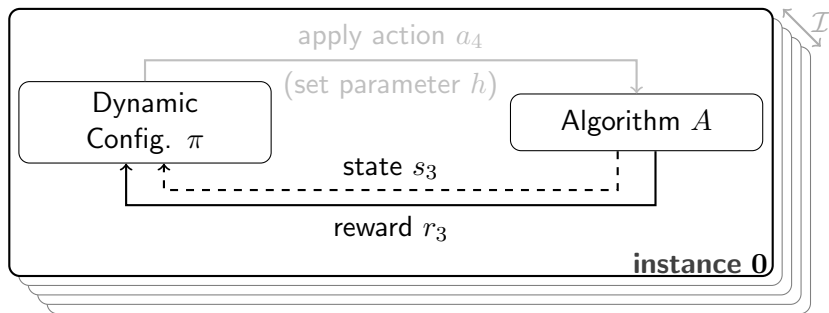
Dynamic Algorithm Configuration

[Biedenkapp et al. 2020]




Dynamic Algorithm Configuration

[Biedenkapp et al. 2020]



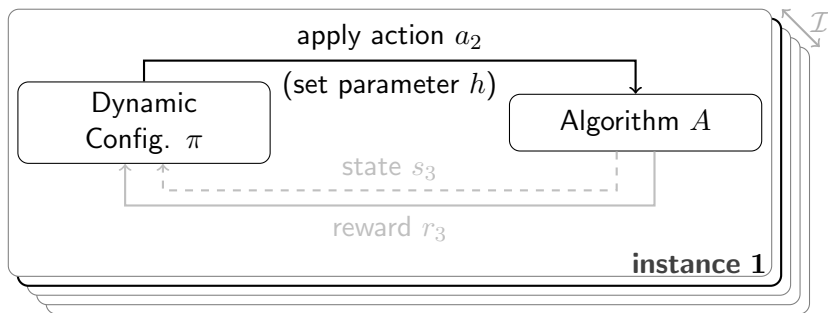
Dynamic Algorithm Configuration

[Biedenkapp et al. 2020]




Dynamic Algorithm Configuration

[Biedenkapp et al. 2020]




Dynamic Algorithm Configuration

[Biedenkapp et al. 2020]



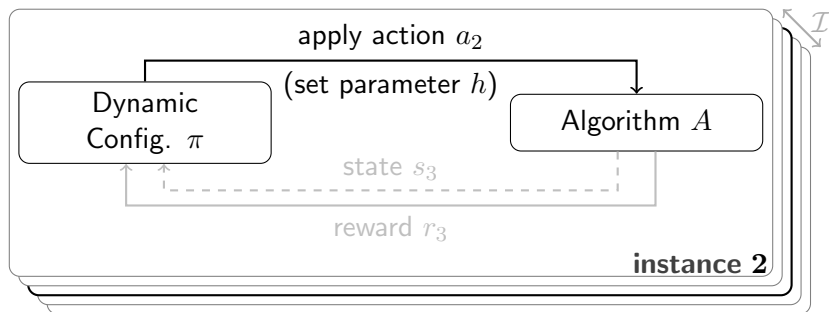
Dynamic Algorithm Configuration

[Biedenkapp et al. 2020]



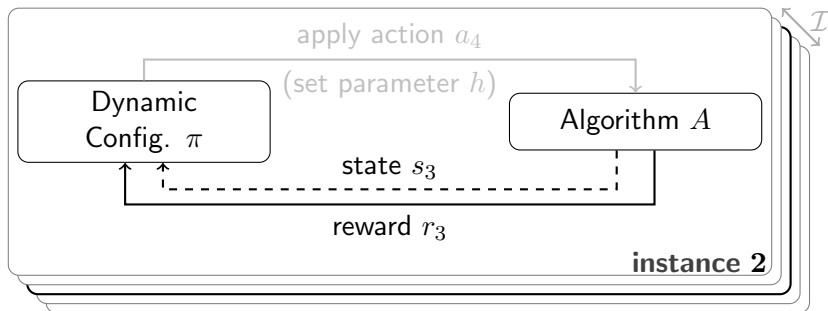
Dynamic Algorithm Configuration

[Biedenkapp et al. 2020]



Dynamic Algorithm Configuration

[Biedenkapp et al. 2020]



⇒ How can DAC look inside the algorithm? I.e. what is a state?

Looking Inside Algorithms: Desiderata

- ① Cheap to compute
- ② Quantify the progress of the algorithm
- ③ Available at each decision point

Looking Inside Algorithms: Desiderata

- ① Cheap to compute
- ② Quantify the progress of the algorithm
- ③ Available at each decision point

⇒ Make use of internal statistics

Looking Inside Algorithms

- EAs
 - ▶ population fitness [Sharma et al. 2019, Shala et al. 2020]
 - ▶ stdev population fitness [Sharma et al. 2019]
 - ▶ cumulative evolution path length [Shala et al. 2020]
 - ▶ ...
- AI Planning [Speck et al. 2020]
 - ▶ average heuristic value
 - ▶ minimal heuristic value
 - ▶ #possible next planning states
 - ▶ ...
- NN Optimization [Daniel et al. 2016]
 - ▶ predictive change in function value
 - ▶ disagreement of function values
 - ▶ uncertainty
 - ▶ ...

Looking Inside Algorithms

- EAs \rightsquigarrow better generalization
 - ▶ population fitness [Sharma et al. 2019, Shala et al. 2020]
 - ▶ stdev population fitness [Sharma et al. 2019]
 - ▶ cumulative evolution path length [Shala et al. 2020]
 - ▶ ...
- AI Planning [Speck et al. 2020] \rightsquigarrow more solved instances
 - ▶ average heuristic value
 - ▶ minimal heuristic value
 - ▶ #possible next planning states
 - ▶ ...
- NN Optimization [Daniel et al. 2016] \rightsquigarrow robust learning rates
 - ▶ predictive change in function value
 - ▶ disagreement of function values
 - ▶ uncertainty
 - ▶ ...

Looking Inside Algorithms

- EAs \rightsquigarrow better generalization
 - ▶ population fitness [Sharma et al. 2019, Shala et al. 2020²]
 - ▶ stdev population fitness [Sharma et al. 2019]
 - ▶ cumulative evolution path length [Shala et al. 2020²]
 - ▶ ...
- AI Planning [Speck et al. 2020] \rightsquigarrow more solved instances
 - ▶ average heuristic value
 - ▶ minimal heuristic value
 - ▶ #possible next planning states
 - ▶ ...
- NN Optimization [Daniel et al. 2016] \rightsquigarrow robust learning rates
 - ▶ predictive change in function value
 - ▶ disagreement of function values
 - ▶ uncertainty
 - ▶ ...

²presented at PPSN in poster session 2

Outline

- 1 What is Algorithm Configuration?
- 2 How can we solve AC?
- 3 What to watch out for?
- 4 Is Solving AC hard?
- 5 Can we learn how to solve AC?
- 6 Top-Down or Bottom-Up Censoring?
- 7 Evolving from Static to Dynamic
- 8 What's next?**

Next I: Neural Networks?



"everyone" uses neural networks these days.

Models in algorithm configuration are still random forests?

Next I: Neural Networks?




"everyone" uses neural networks these days.

Models in algorithm configuration are still random forests?

Towards DNNs for AC (I): Prediction of Runtime Distributions

[Eggersperger et al. 2018]



Next II: Neural Networks?



How to handle censored data (from capping of runs) in DNNs?

Next II: Neural Networks?

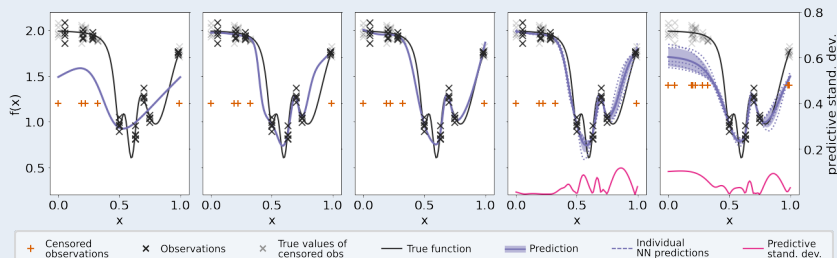


How to handle censored data (from capping of runs) in DNNs?

Towards DNNs for AC (II): Handling of Censored Data

[Eggensperger et al. 2020]


$$-\log \mathcal{L} \left((\hat{\mu}_i, \hat{\sigma}_i^2)_{i=1}^n \mid \mathcal{D} \right) = - \sum_{i=1}^n \log \left(\phi(Z_i)^{1-I_i} (1 - \Phi(Z_i))^{I_i} \right)$$



Next II: Multi-Objective Optimization?




🤔 Can we consider more than one objective?



Next II: Multi-Objective Optimization?



Can we consider more than one objective?



Challenges:

- 1 How to sample promising (non-dominated) configurations?
- 2 How to efficiently determine that a configuration is really on the Pareto front across all instances?

[Blot et al. 2016]

Next III: What are interesting AC benchmarks?



I developed a new configurator,
but on which benchmarks should I benchmark it?

<https://bitbucket.org/mlindauer/aclib2/>

Next III: What are interesting AC benchmarks?



I developed a new configurator,
but on which benchmarks should I benchmark it?

AClib [Hutter et al. 2014]

- 84 AC scenarios
- 18 target algorithms and their configuration spaces
- 45 instance sets
- 6 AI domains
- 6 configurators

<https://bitbucket.org/mlindauer/aclib2/>

Next III: What are interesting AC benchmarks?



I developed a new configurator,
but on which benchmarks should I benchmark it?

AClib [Hutter et al. 2014]

- 84 AC scenarios
- 18 target algorithms and their configuration spaces
- 45 instance sets
- 6 AI domains
- 6 configurators

- ~> To make development cheaper, you can benchmark on surrogate benchmarks [Eggensperger et al. 2018]
- ~> Open which of these are interesting?
 - ▶ not too easy, not too hard, diverse set, different characteristics, representative, ...

<https://bitbucket.org/mlindauer/aclib2/>

Take Home Message

- ❶ Algorithm Configuration improves the performance of your algorithms

Take Home Message

- ① Algorithm Configuration improves the performance of your algorithms
- ② There are different state-of-the-art AC approaches

Take Home Message

- ① Algorithm Configuration improves the performance of your algorithms
- ② There are different state-of-the-art AC approaches
- ③ Using algorithm configuration is nevertheless fairly easy these days

Take Home Message

- ① Algorithm Configuration improves the performance of your algorithms
- ② There are different state-of-the-art AC approaches
- ③ Using algorithm configuration is nevertheless fairly easy these days
- ④ Many open questions for the research community to be answered

Thank you!

