Incorporating Decision-Maker's Preferences into the Automatic Configuration of Bi-objective Optimisation Algorithms

Manuel López-Ibáñez

manuel.lopez-ibanez@manchester.ac.uk
http://lopez-ibanez.eu

The Joint Lectures on Evolutionary Algorithms (JoLEA)

15 February 2023

The University of Manchester
Alliance Manchester Business School

Who am I?

- Senior Lecturer (Assoc. Prof.) at AMBS (UoM), 2015–
- Senior Distinguished Researcher, University of Málaga, 2020–2022
- Editor-in-Chief of GECCO 2019
- Committee Chair of the ACM SIGEVO Dissertation Award
- (co-)Editor-in-Chief of
  ACM Transactions on Evolutionary Learning and Optimization (TELO)

My research is about . . .

- Benchmarking and Empirical Analysis of Optimization Algorithms
  - ☞ Reproducibility in Evolutionary Computation [López-Ibáñez, Branke & Paquete, 2021]
- Multi-objective optimization ☞ EAF package
- Interactive optimization (human-in-the-loop)
  - ☞ Machine Decision Makers [López-Ibáñez & Knowles, 2015]
- Automatic configuration, selection and design of algorithms ☞ irace
- Expensive optimization . . .
- Applications, applications, applications!
  - School bus routing for SEND students
  - Optimization in steel manufacturing
  - Supply chain design for Personalised Medicine
  - Bayesian Optimisation with dynamic constraints
  - Patient scheduling

My research is about . . .

- Benchmarking and Empirical Analysis of Optimization Algorithms
  - ☞ Reproducibility in Evolutionary Computation [López-Ibáñez, Branke & Paquete, 2021]
- *Multi-objective optimization* ☞ EAF package
- *Interactive optimization (human-in-the-loop)*
  - ☞ Machine Decision Makers [López-Ibáñez & Knowles, 2015]
- *Automatic configuration, selection and design of algorithms* ☞ irace
- Expensive optimization . . .
- Applications, applications, applications!
  - School bus routing for SEND students
  - Optimization in steel manufacturing
  - Supply chain design for Personalised Medicine
  - Bayesian Optimisation with dynamic constraints
  - Patient scheduling

Today's topic

Juan E. Diaz Leiva

Universidad San Francisco

Manuel López-Ibáñez

MANCHESTER 1824
The University of Manchester
Alliance Manchester Business School

Juan Esteban Diaz and Manuel López-Ibáñez,
Incorporating decision-maker's preferences into the automatic configuration of bi-objective optimisation algorithms,
European Journal of Operational Research, 289:3, 16 March 2021, pages 1209–1222.
https://doi.org/10.1016/j.ejor.2020.07.059
☆ EJOR Editors' Choice Article, January 2021

Agenda

(1) Automatic algorithm configuration

(2) Bi-objective optimization

(3) The empirical attainment function (EAF)

(4) Incorporating decision-maker's preferences into the automatic configuration of bi-objective optimisation algorithms

Design choices and parameters everywhere

Modern high-performance optimizers involve a large number of design choices and (hyper)-parameter settings

- Exact solvers
  - Design choices: alternative models, pre-processing, variable selection, value selection, branching rules . . .
    $+$ numerical parameters
  - IBM CPLEX: 63 parameters that control the optimization

Design choices and parameters everywhere

Modern high-performance optimizers involve a large number of design choices and (hyper)-parameter settings

- Exact solvers
  - Design choices: alternative models, pre-processing, variable selection, value selection, branching rules . . .

    $+$ numerical parameters
  - IBM CPLEX: 63 parameters that control the optimization

Design choices and parameters everywhere

Modern high-performance optimizers involve a large number of design choices and (hyper)-parameter settings

- Exact solvers
  - Design choices: alternative models, pre-processing, variable selection, value selection, branching rules ...
    + numerical parameters
  - IBM CPLEX: 63 parameters that control the optimization
- (Meta)-heuristic solvers
  - Design choices: solution representation, operators, neighborhoods, pre-processing, strategies, ... + numerical parameters
  - Many are *hidden*

Design choices and parameters everywhere

Modern high-performance ~~optimizers~~ software involve a large number of design choices and (hyper)-parameter settings

Domain	Software	Parameters	
ML	WEKA	768	[Kotthoff et al., 2016]
	Auto-sklearn	110	[Feurer et al., 2015]
AI Planning	LGP,	45 – 66	[Vallati et al., 2011]
	Fast-downward		[Fawcett et al., 2011]
Code optimization	GCC	172 flags	[Pérez Cáceres et al., 2017]
		+ 195 numerical	

Design choices and parameters everywhere

Modern high-performance ~~optimizers~~ software involve a large number of design choices and (hyper)-parameter settings

Design choices and parameters everywhere

- *categorical* parameters

  recombination $\in$ { uniform, one-point, two-point }
  localsearch $\in$ { tabu search, SA, ILS }

- *ordinal* parameters

  neighborhoods $\in$ { small, medium, large }

- *numerical* parameters

  population sizes, acceptance temperature,
  hidden constants, ...

- Parameters may be *conditional* to specific values of other parameters:

  `temperature` only enabled if `localsearch == "SA"`

Design choices and parameters everywhere

- *categorical* parameters

  recombination $\in$ { uniform, one-point, two-point }
  localsearch $\in$ { tabu search, SA, ILS }

- *ordinal* parameters

  neighborhoods $\in$ { small, medium, large }

- *numerical* parameters

  population sizes, acceptance temperature,
  hidden constants, ...

- Parameters may be *conditional* to specific values of other parameters:

  `temperature` only enabled if `localsearch == "SA"`

Configuring an algorithm involves setting categorical, ordinal and numerical parameters

The algorithm configuration problem

1. Find the best algorithm configuration given a set of *training problem instances*
2. Repeatedly use this algorithm configuration to solve *unseen problem instances*

The algorithm configuration problem

1. Find the best algorithm configuration
   given a set of *training problem instances*

2. Repeatedly use this algorithm configuration to solve
   unseen problem instances

A problem with many names:

offline parameter *tuning*,
automatic algorithm configuration,
hyper-parameter optimization,
hyper-heuristics, genetic programming,
meta-optimisation, programming by optimisation [Hoos, 2012], . . .

Automatic Algorithm Configuration: How?

A mixed-integer stochastic black-box optimization problem

- Mixed decision variables:

  categorical, ordinal and numerical (real-valued and integer)

- Stochasticity from algorithm and problem instances
- Black-box: evaluation requires running the algorithm

Automatic Algorithm Configuration: How?

A mixed-integer stochastic black-box optimization problem

- Mixed decision variables:

  categorical, ordinal and numerical (real-valued and integer)

- Stochasticity from algorithm and problem instances
- Black-box: evaluation requires running the algorithm

Methods for Automatic Algorithm Configuration

- **SPOT** [Bartz-Beielstein, Lasarczyk & Preuss, 2005]
- **ParamILS** [Hutter, Hoos & Stützle, 2007]
- **SMAC** [Hutter, Hoos & Leyton-Brown, 2011]
- **IRACE** [López-Ibáñez, Dubois-Lacoste, Stützle & Birattari, 2011]
  [López-Ibáñez, Dubois-Lacoste, Pérez Cáceres, Stützle & Birattari, 2016]

Automatic Algorithm Configuration

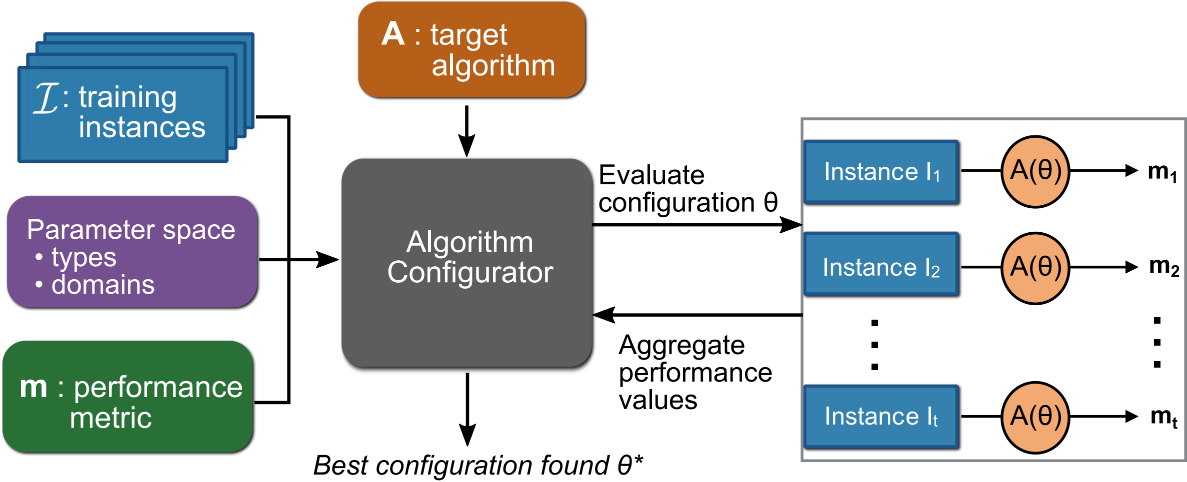

Bi-objective Optimisation

Two objective functions: $\vec{f} = (f_1(x), f_2(x))$

How to minimize both of them?

Bi-objective Optimisation

Two objective functions: $\vec{f} = (f_1(x), f_2(x))$

How to minimize both of them?

- *A priori* knowledge / preferences known
  - ☞ Aggregate into a scalar value:

$$f_s(x) = \lambda \cdot f_1(x) + (1 - \lambda) \cdot f_2(x)$$

Bi-objective Optimisation

Two objective functions: $\vec{f} = (f_1(x), f_2(x))$

How to minimize both of them?

- *A priori* knowledge / preferences known
  - ☞ Aggregate into a scalar value:

$$f_s(x) = \lambda \cdot f_1(x) + (1 - \lambda) \cdot f_2(x)$$

- No a priori knowledge $\Rightarrow$ *Pareto-optimality*

Bi-objective Optimisation

Two objective functions: Minimize $\vec{f} = (f_1(x), f_2(x))$

No a priori knowledge $\Rightarrow$ *Pareto-optimality*

Bi-objective Optimisation

Two objective functions: Minimize $\vec{f} = (f_1(x), f_2(x))$

No a priori knowledge $\Rightarrow$ *Pareto-optimality*

a dominates b

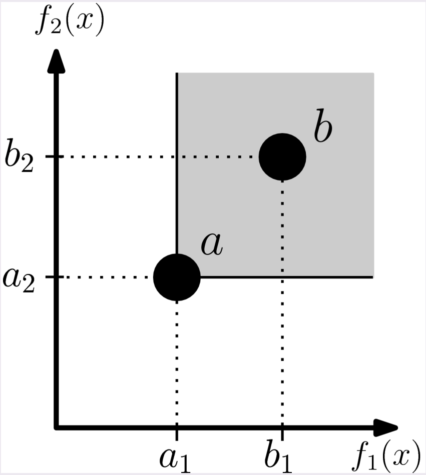

Bi-objective Optimisation

Two objective functions: Minimize $\vec{f} = (f_1(x), f_2(x))$

No a priori knowledge $\Rightarrow$ *Pareto-optimality*

a dominates b

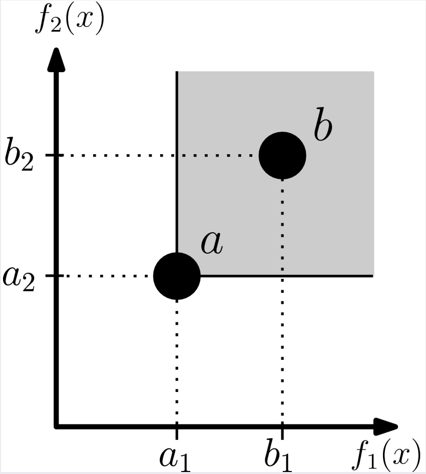

a and b nondominated

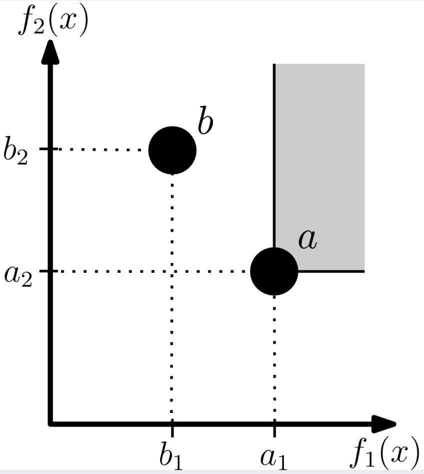

Bi-Objective Optimisation

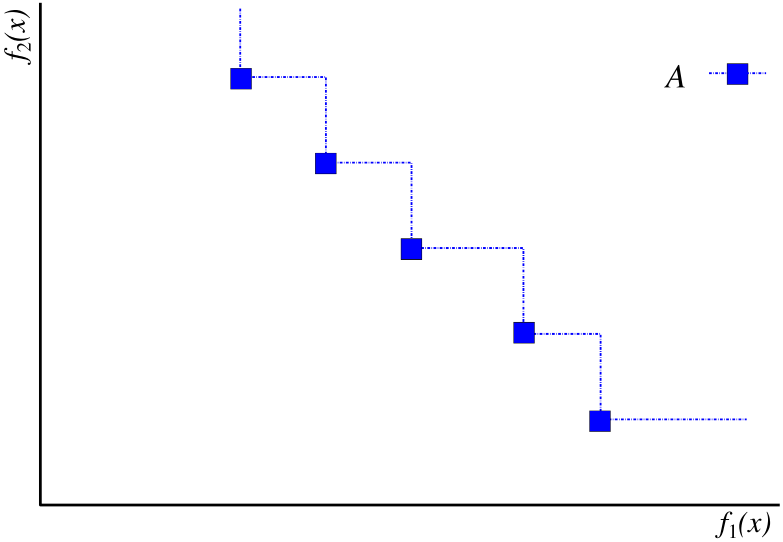

Bi-Objective Optimisation

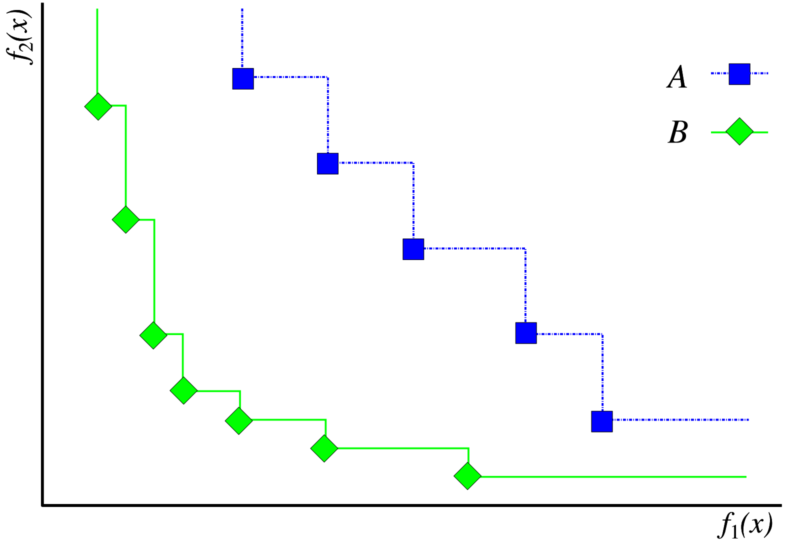

Bi-Objective Optimisation

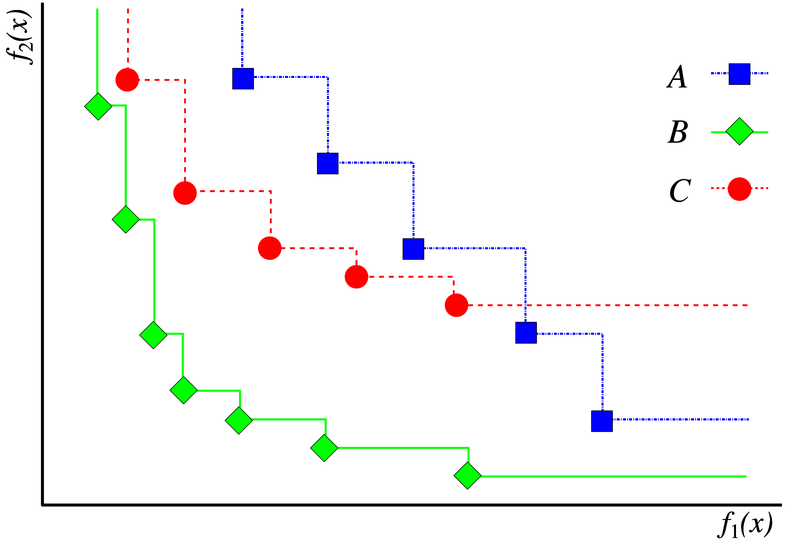

Stochastic Multi-criteria Data

Expensive ⟶ Cheap
ICE **EC, IC** **RE,RB**
Fast ⟶ *Slow*

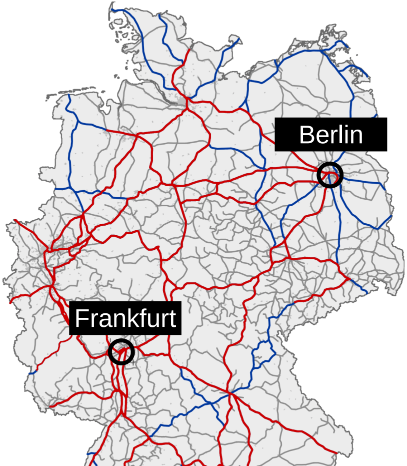

Fastest and cheapest train trip from Frankfurt to Berlin?

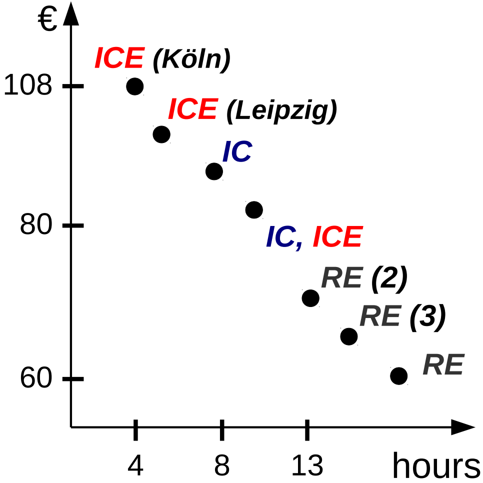

Stochastic Multi-criteria Data

Expensive ——► Cheap

ICE **EC, IC** **RE,RB**

Fast ——► *Slow*

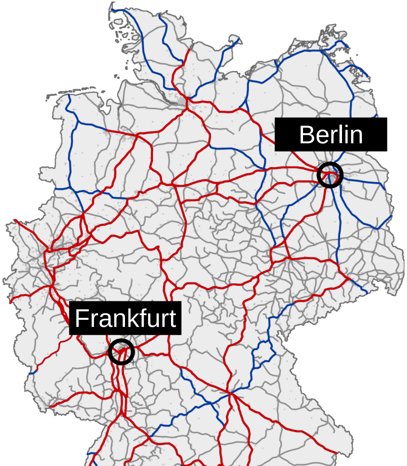

Fastest and cheapest train trip from Frankfurt to Berlin?

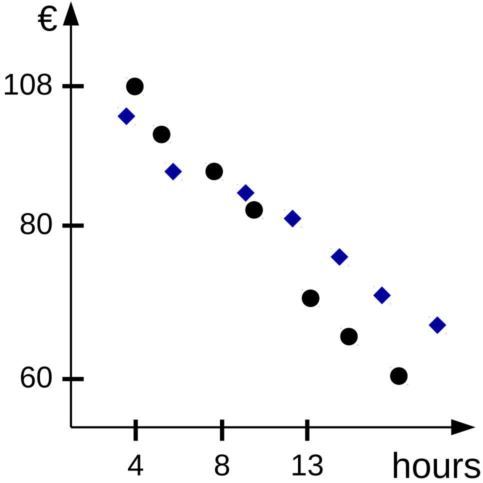

Stochastic Multi-criteria Data

Expensive → Cheap

ICE ***EC, IC*** ***RE,RB***

Fast → *Slow*

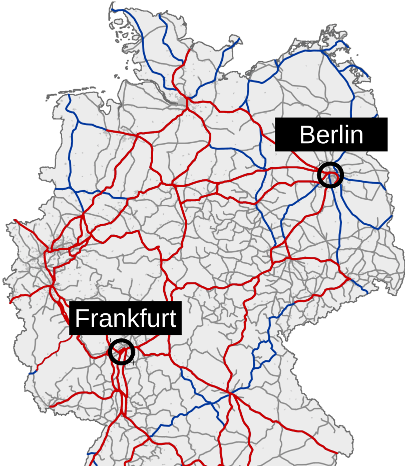

Fastest and cheapest train trip from Frankfurt to Berlin?

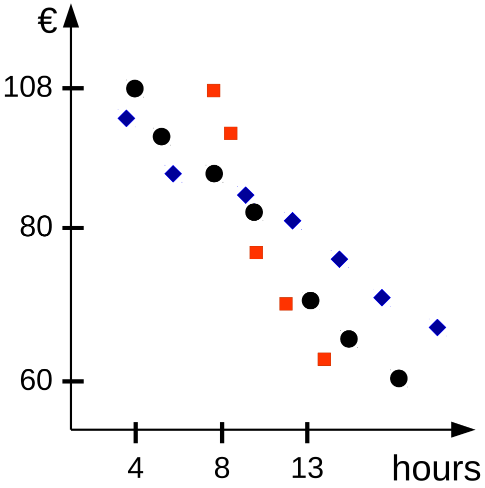

Stochastic Nondominated Sets

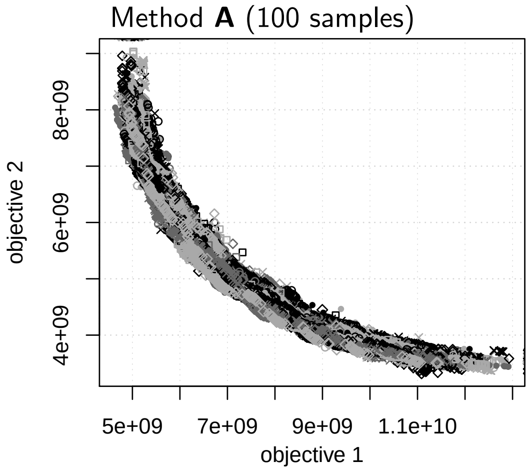

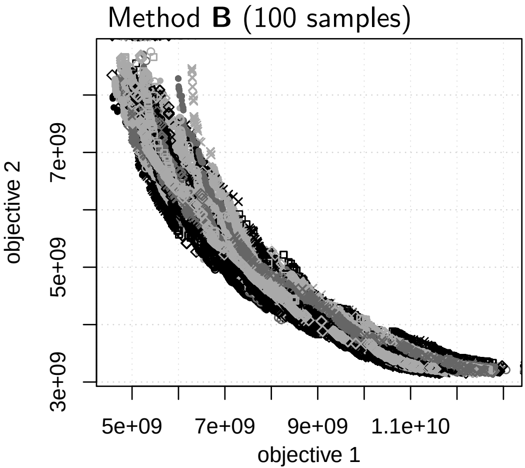

Stochastic Nondominated Sets

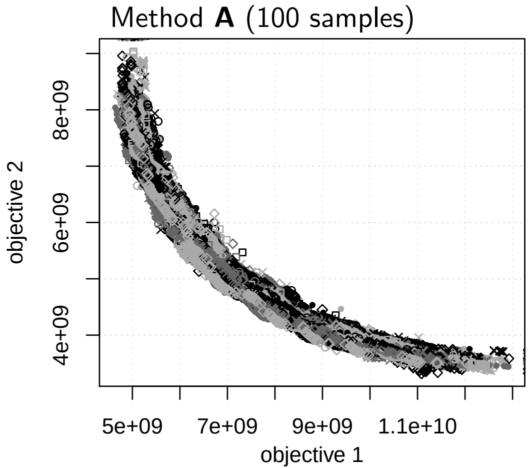

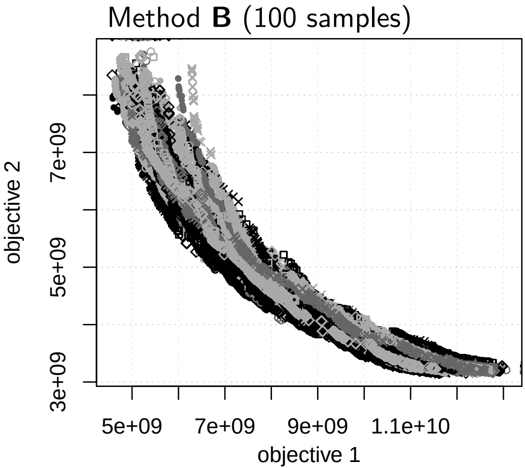

Which one is better?

Hypervolume metric

[Knowles & Corne, 2003; Zitzler et al., 2003]

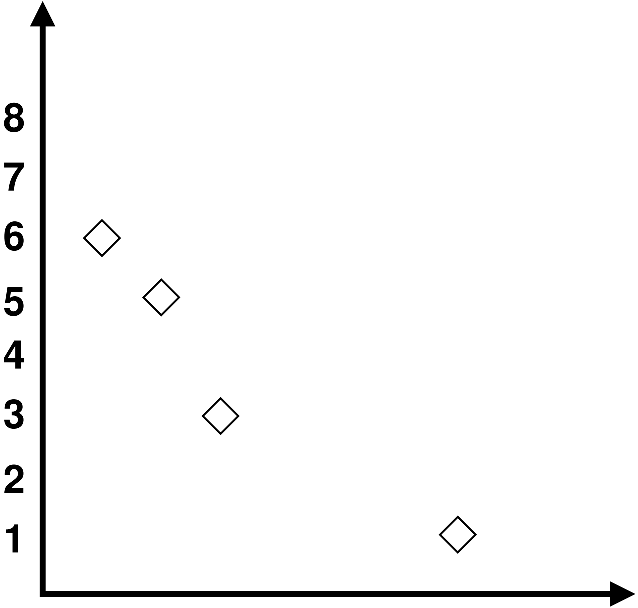

Hypervolume metric

[Knowles & Corne, 2003; Zitzler et al., 2003]

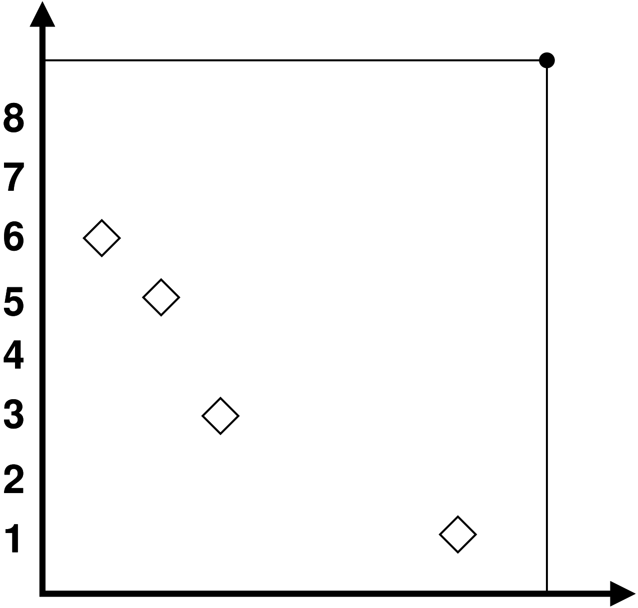

Hypervolume metric

[Knowles & Corne, 2003; Zitzler et al., 2003]

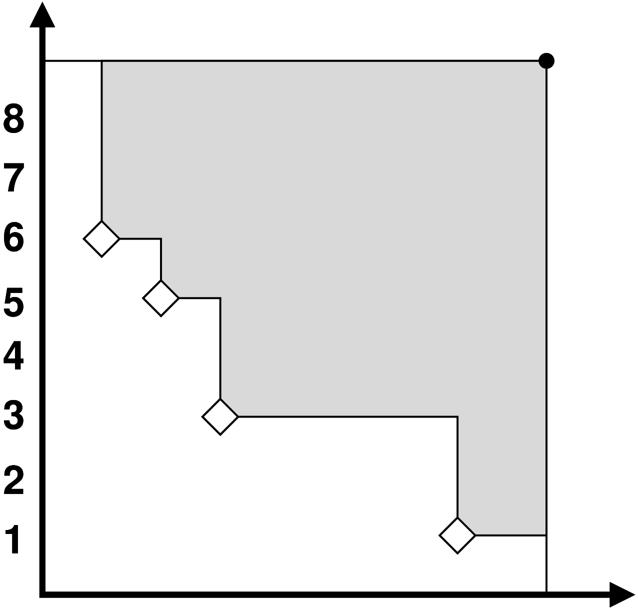

Stochastic Nondominated Sets

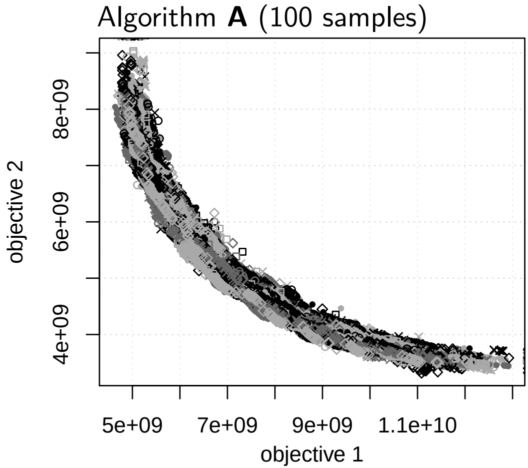

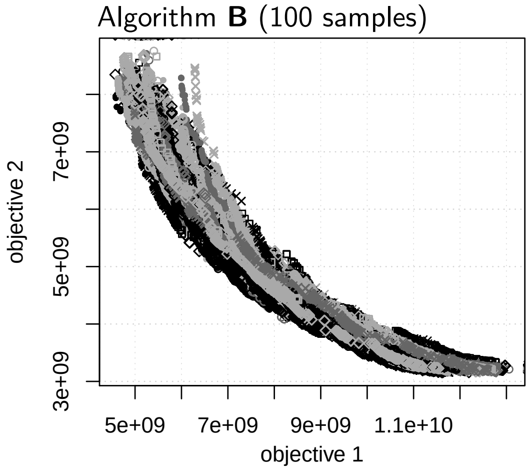

Mean HV = 0.920016

Mean HV = 0.9070087

Automatic Configuration of Multi-Objective Optimizers

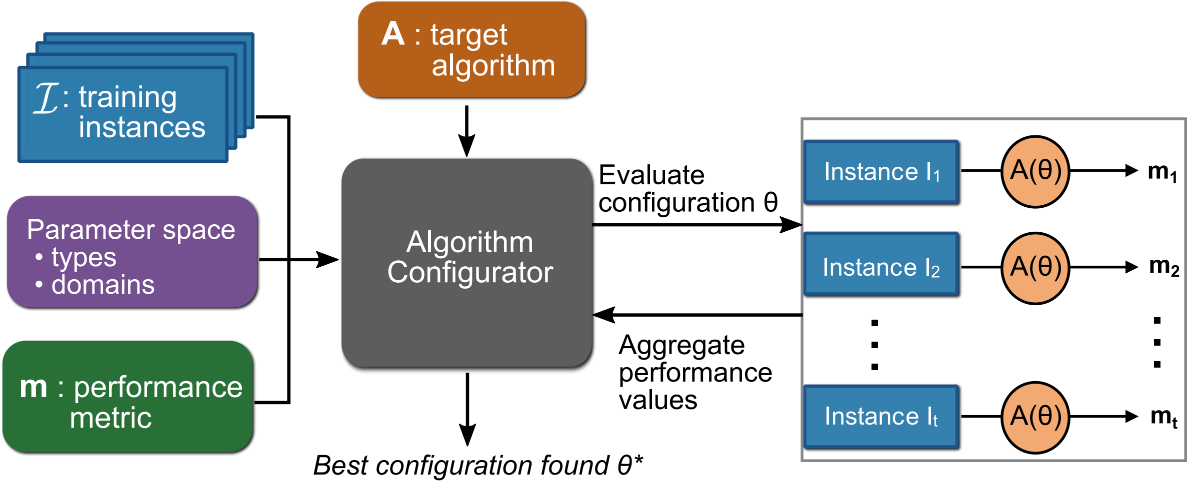

Automatic Configuration of Multi-Objective Optimizers

✘ Multi-objective!

- Algorithm's output is an approximation to the PF
- AC performance metric is scalar

Automatic Configuration of Multi-Objective Optimizers

✘ Multi-objective!

- Algorithm's output is an approximation to the PF
- AC performance metric is scalar

✔ `irace` + hypervolume = automatic configuration of multi-objective solvers!

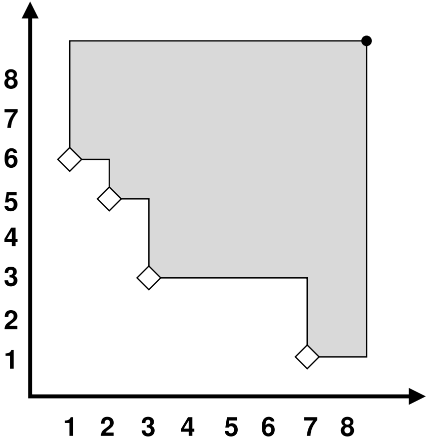

Automatic Configuration of Multi-Objective Optimizers

✘ Multi-objective!

- Algorithm's output is an approximation to the PF
- AC performance metric is scalar

✔ `irace` + hypervolume = automatic configuration of multi-objective solvers!

- TP+PLS: Multi-objective local search framework for PFSP [Dubois-Lacoste et al., 2011]
- Multi/Many-objective Evolutionary Algorithms (AutoMOEA) [Bezerra et al., 2016, 2020a]
- Anytime behavior of exact and heuristic algorithms [López-Ibáñez & Stützle, 2014]

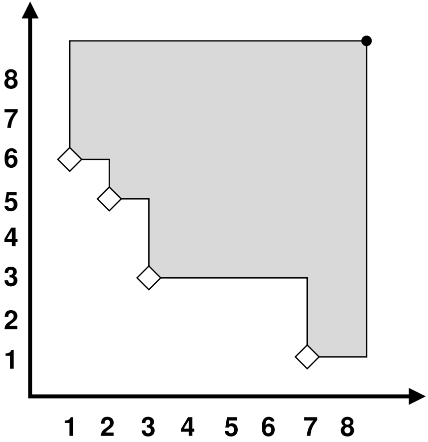

Is HV enough?

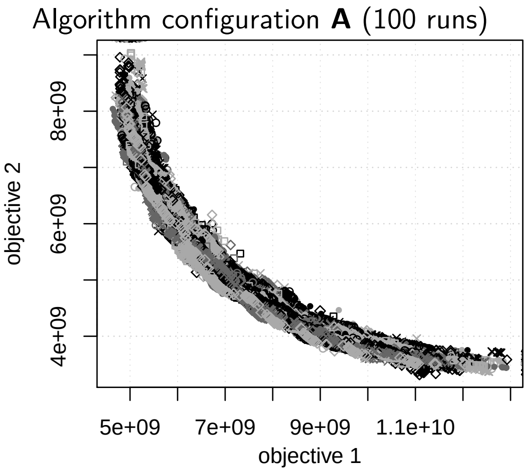

Mean HV = 0.920016

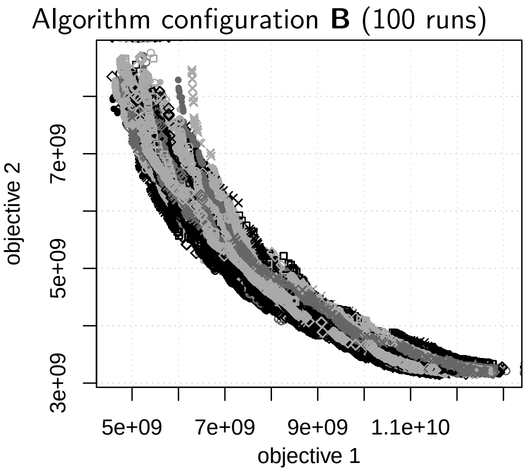

Mean HV = 0.9070087

What are the differences between these two methods?

Is HV enough?

Expensive ⟶ Cheap

ICE **EC, IC** **RE,RB**

Fast ⟶ *Slow*

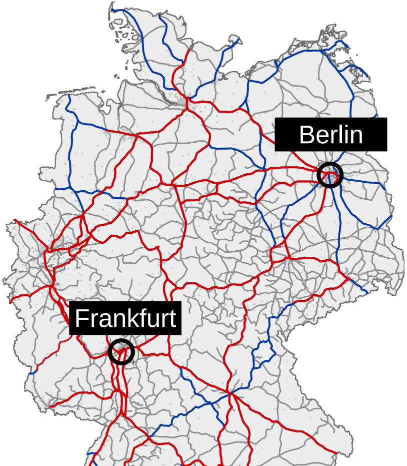

Fastest and cheapest train trip from Frankfurt to Berlin?

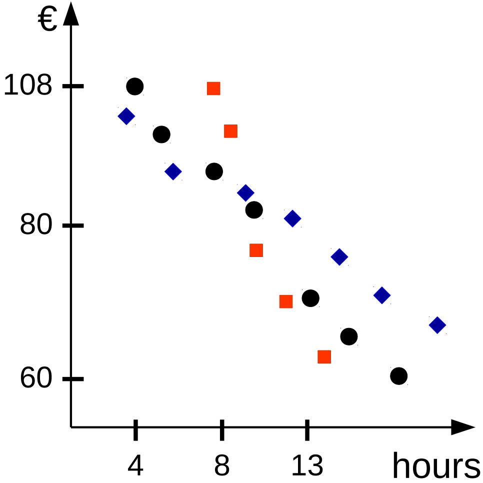

Is HV enough?

Expensive ⟶ Cheap

ICE ***EC, IC*** ***RE,RB***

Fast ⟶ *Slow*

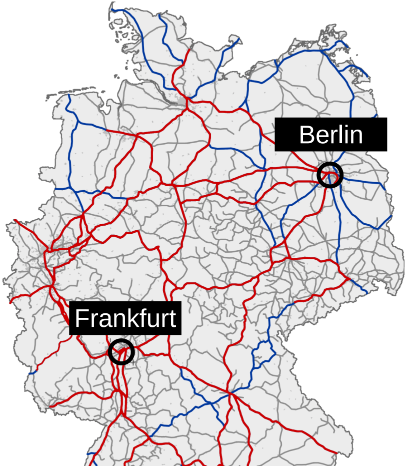

Fastest and cheapest train trip from Frankfurt to Berlin?

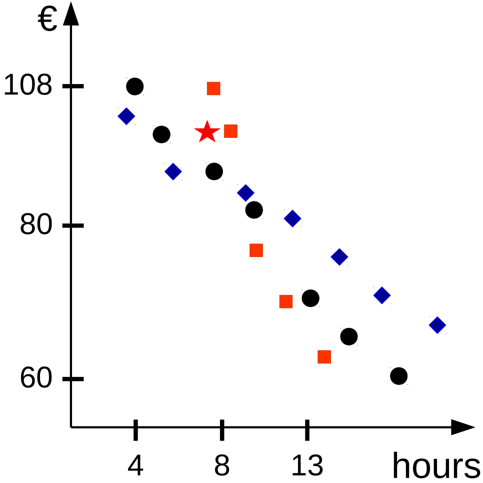

The Attainment Function

[Grunert da Fonseca et al., 2001]

- Probability of a random set $\mathcal{A}$ *attaining* (weakly dominating) a particular point $\vec{z} = \vec{f}(x)$ in the objective space

$$\alpha_{\mathcal{A}}(\vec{z}) = \Pr\{\mathcal{A} \preceq \vec{z}\}$$

- Statistical distribution of the output $(\mathcal{A})$ of stochastic multi-objective optimizers

✘ The real attainment function is *unknown* but...

The Attainment Function

[Grunert da Fonseca et al., 2001]

- Probability of a random set $\mathcal{A}$ *attaining* (weakly dominating) a particular point $\vec{z} = \vec{f}(x)$ in the objective space

$$\alpha_{\mathcal{A}}(\vec{z}) = \Pr\{\mathcal{A} \preceq \vec{z}\}$$

- Statistical distribution of the output ($\mathcal{A}$) of stochastic multi-objective optimizers

✘ The real attainment function is *unknown* but...

✔ ...we can *estimate* it given $\{A_1, A_2, \ldots\} \sim \mathcal{A}$:

Empirical attainment function (EAF)

The Empirical Attainment Function

[Grunert da Fonseca et al., 2001]

$$\alpha_{\mathcal{A}}(\vec{z}) = \Pr\{\mathcal{A} \preceq \vec{z}\}$$

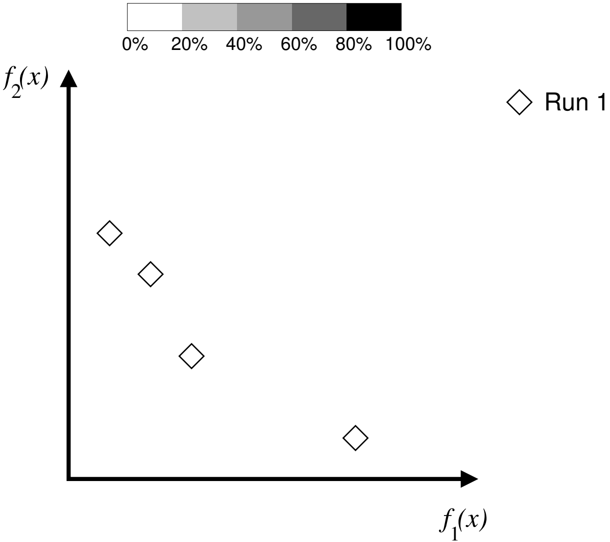

The Empirical Attainment Function [Grunert da Fonseca et al., 2001]

$$\alpha_{\mathcal{A}}(\vec{z}) = \Pr\{\mathcal{A} \preceq \vec{z}\}$$

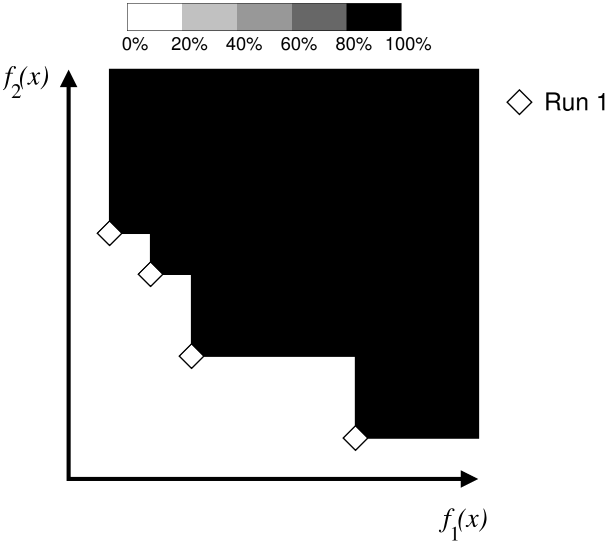

The Empirical Attainment Function

[Grunert da Fonseca et al., 2001]

$$\alpha_{\mathcal{A}}(\vec{z}) = \Pr\{\mathcal{A} \preceq \vec{z}\}$$

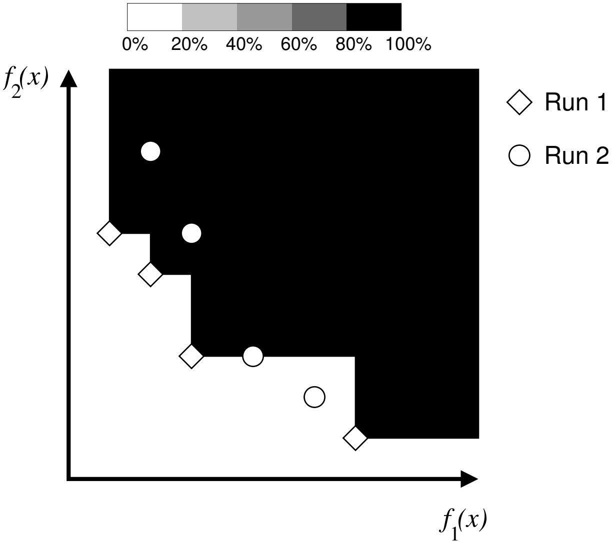

The Empirical Attainment Function

$$\alpha_{\mathcal{A}}(\vec{z}) = \Pr\{\mathcal{A} \preceq \vec{z}\}$$

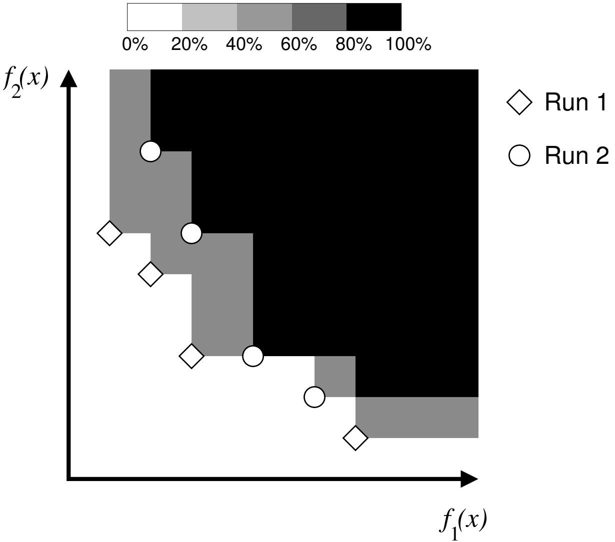

The Empirical Attainment Function

[Grunert da Fonseca et al., 2001]

$$\alpha_{\mathcal{A}}(\vec{z}) = \Pr\{\mathcal{A} \preceq \vec{z}\}$$

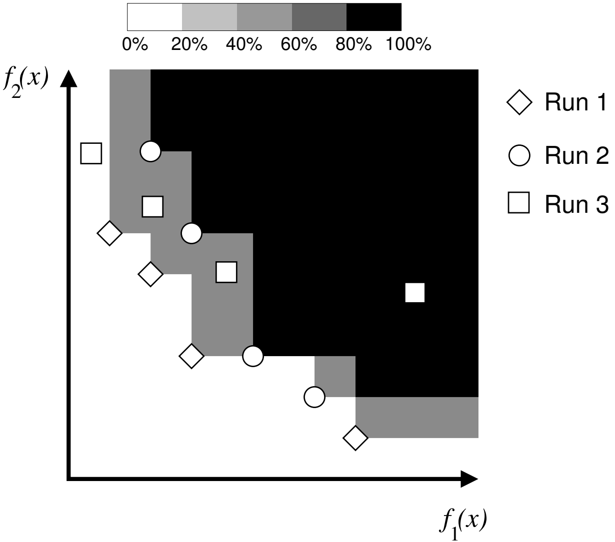

The Empirical Attainment Function

[Grunert da Fonseca et al., 2001]

$$\alpha_{\mathcal{A}}(\vec{z}) = \Pr\{\mathcal{A} \preceq \vec{z}\}$$

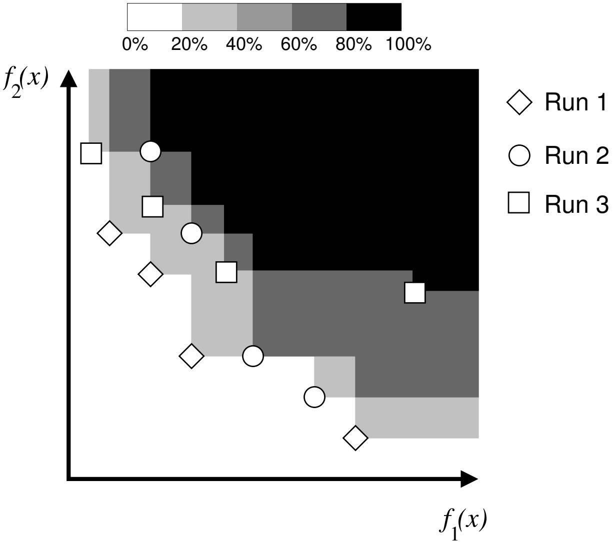

Stochastic Nondominated Sets

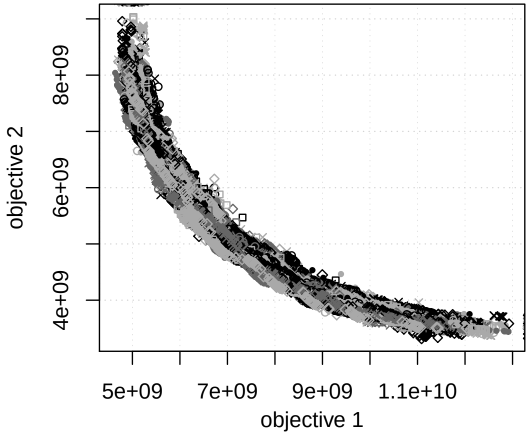

Algorithm **A** (100 runs)

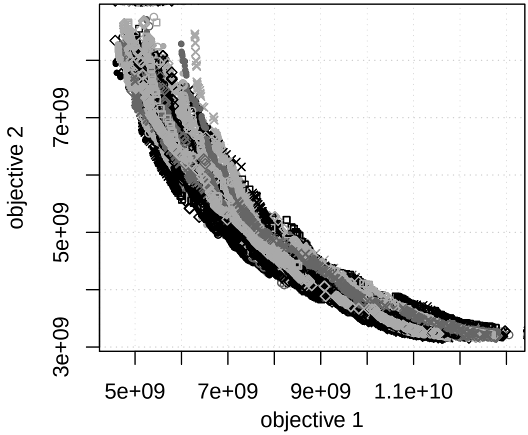

Algorithm **B** (100 runs)

Comparing two algorithms: Side-by-side EAFs

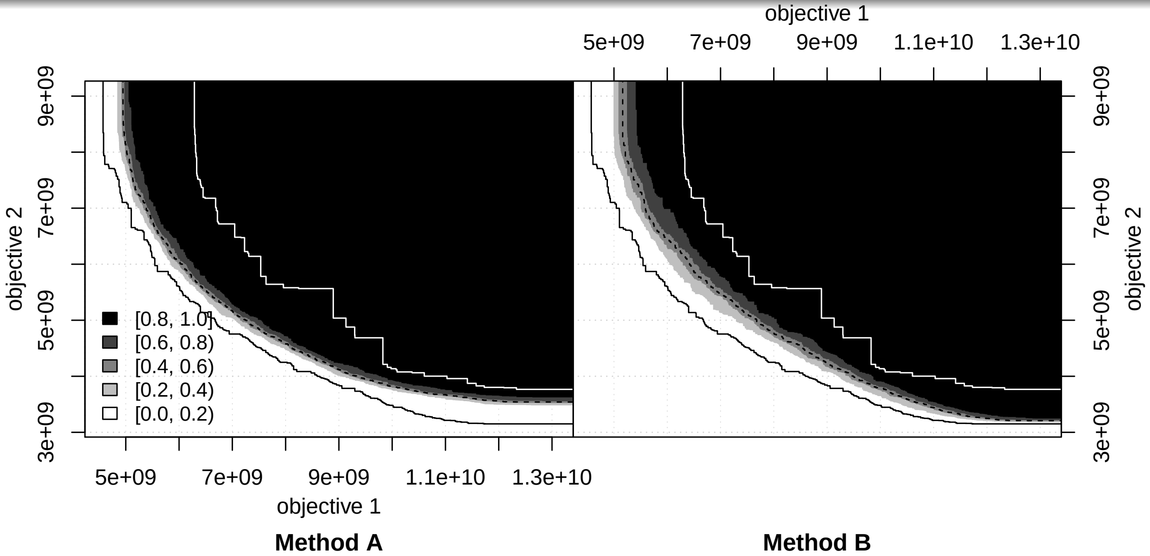

$\alpha_A(\vec{z})$ $\alpha_B(\vec{z})$

Comparing two algorithms: EAF differences

[López-Ibáñez et al., 2010]

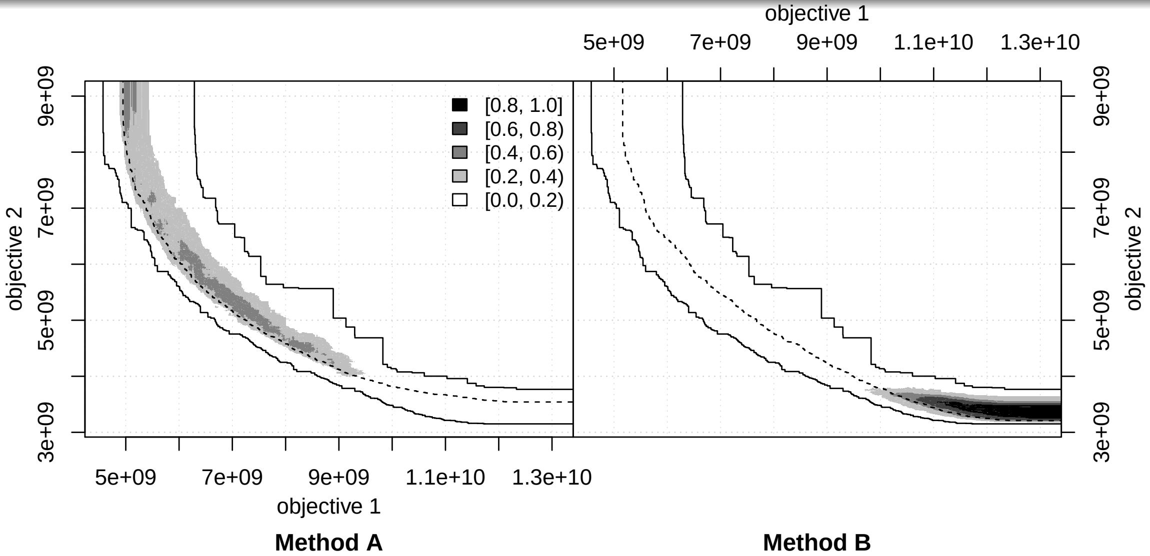

$$\max\{\alpha_A(\vec{z}) - \alpha_B(\vec{z}), 0\}$$

$$\max\{\alpha_B(\vec{z}) - \alpha_A(\vec{z}), 0\}$$

- Scalarize the bQAP using a weight $\lambda \in [0, 1]$

$$f_\lambda(x) = \lambda \cdot f_1(x) + (1 - \lambda) \cdot f_2(x)$$

- Solve each f_λ with RoTS [Taillard, 1991] stopping after $\ell \cdot n$ iterations

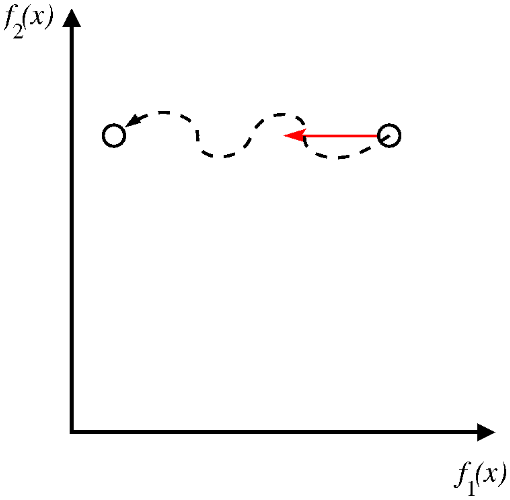

W-RoTS for the bi-objective QAP [López-Ibáñez et al., 2006]

- Scalarize the bQAP using a weight $\lambda \in [0, 1]$

$$f_\lambda(x) = \lambda \cdot f_1(x) + (1 - \lambda) \cdot f_2(x)$$

- Solve each f_λ with RoTS [Taillard, 1991] stopping after $\ell \cdot n$ iterations
- Each RoTS run starts from a random solution

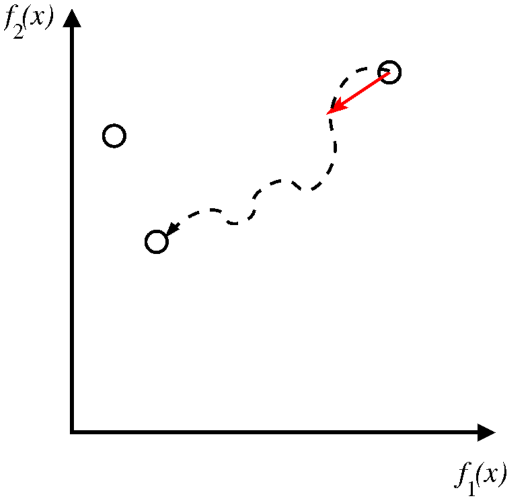

W-RoTS for the bi-objective QAP

[López-Ibáñez et al., 2006]

- Scalarize the bQAP using a weight $\lambda \in [0,1]$

$$f_\lambda(x) = \lambda \cdot f_1(x) + (1 - \lambda) \cdot f_2(x)$$

- Solve each f_λ with RoTS [Taillard, 1991] stopping after $\ell \cdot n$ iterations
- Each RoTS run starts from a random solution
- Maximum time limited to 1s: higher $\ell \Rightarrow$ fewer λ

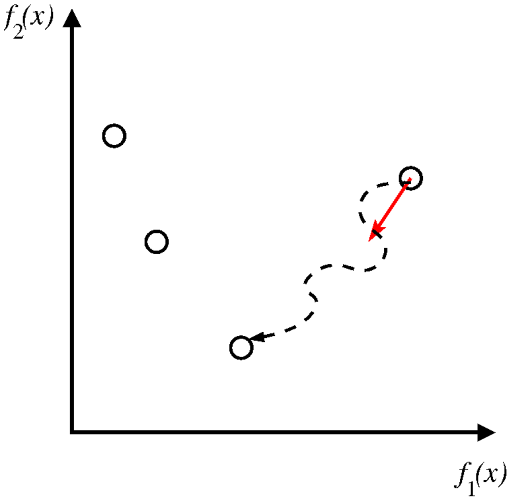

W-RoTS for the bi-objective QAP

[López-Ibáñez et al., 2006]

- Scalarize the bQAP using a weight $\lambda \in [0,1]$

$$f_\lambda(x) = \lambda \cdot f_1(x) + (1-\lambda) \cdot f_2(x)$$

- Solve each f_λ with RoTS [Taillard, 1991] stopping after $\ell \cdot n$ iterations
- Each RoTS run starts from a random solution
- Maximum time limited to 1s: higher $\ell \Rightarrow$ fewer λ

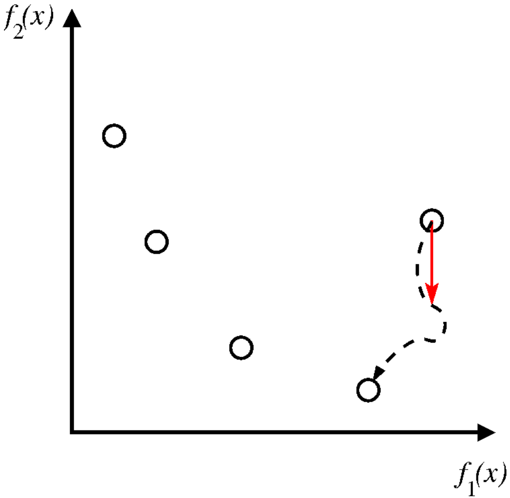

W-RoTS for the bi-objective QAP

[López-Ibáñez et al., 2006]

- Scalarize the bQAP using a weight $\lambda \in [0, 1]$

$$f_\lambda(x) = \lambda \cdot f_1(x) + (1 - \lambda) \cdot f_2(x)$$

- Solve each f_λ with RoTS [Taillard, 1991] stopping after $\ell \cdot n$ iterations
- Each RoTS run starts from a random solution
- Maximum time limited to 1s: higher $\ell \Rightarrow$ fewer λ

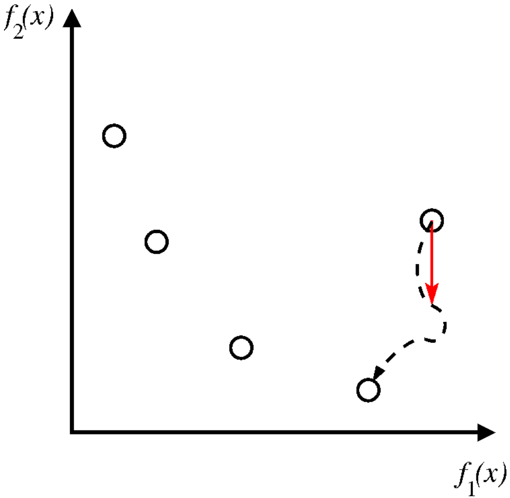

Best value of ℓ ?

Tuning the value of ℓ in W-RoTS

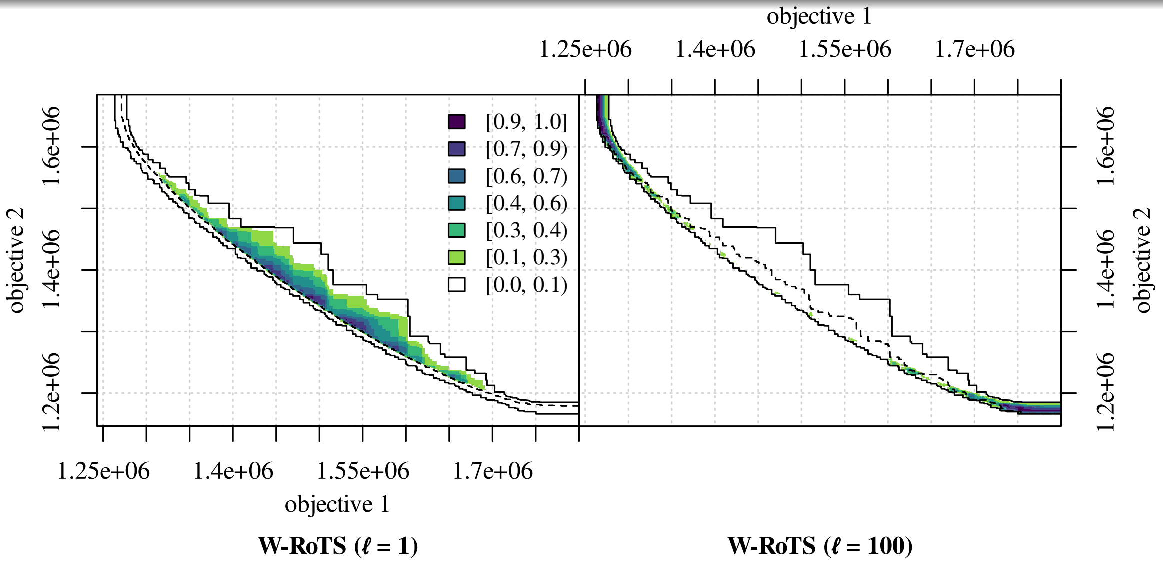

Tuning the value of ℓ in W-RoTS

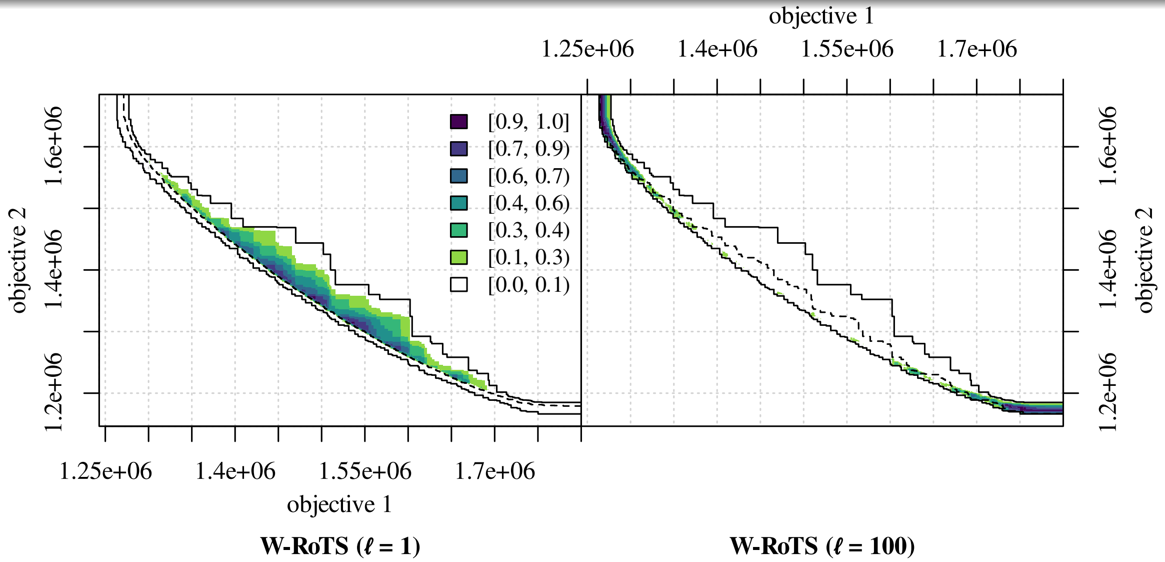

HV assumes a preference that may be different from the DM's !

Weighted Hypervolume

[Zitzler et al., 2007]

Bias the computation of the HV by assigning weights to regions of the objective space

Weighted Hypervolume

[Zitzler et al., 2007]

Bias the computation of the HV by assigning weights to regions of the objective space

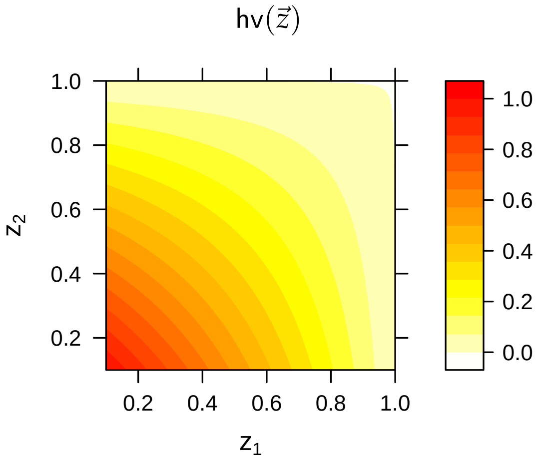

Weighted Hypervolume

Bias the computation of the HV by assigning weights to regions of the objective space

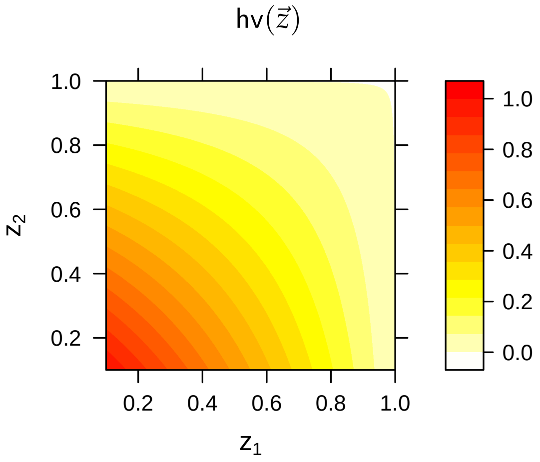

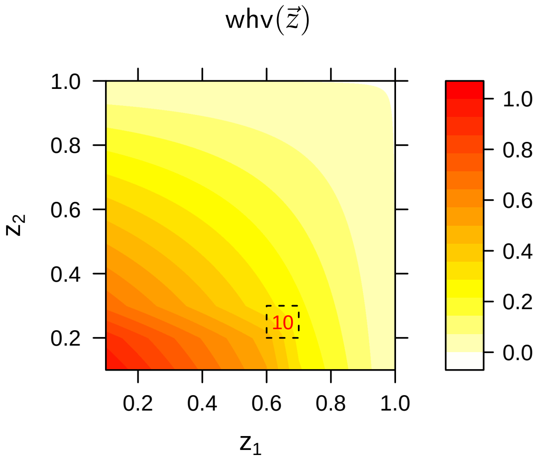

Interactively Tuning the value of ℓ in W-RoTS

Convert EAF differences into weighted hypervolume

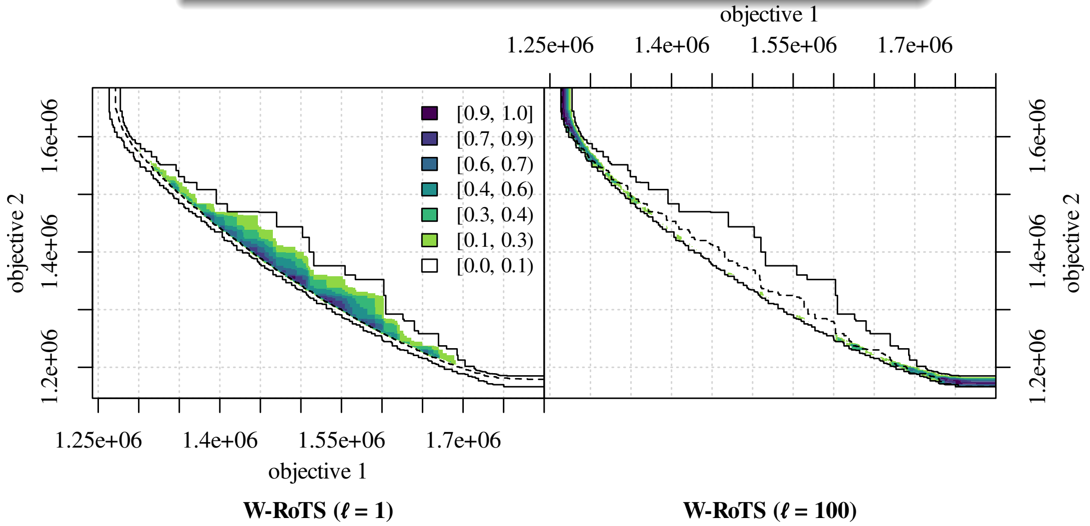

Interactively Tuning the value of ℓ in W-RoTS

Use this weighted hypervolume to guide
the automatic algorithm configuration of a bi-objective optimizer

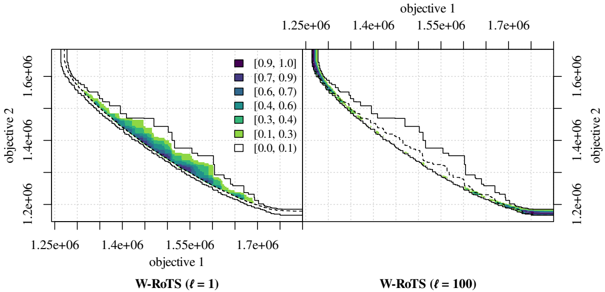

(1) The DM chooses one side, e.g., $\ell = 1$

(2) Compute regions $\mathcal{R}$ in favour

(2) Create $HV^{\mathrm{w}}_{\mathcal{R}}(\cdot)$ based on positive EAF differences

(3) Tune $\ell \in [1, 200]$ using irace guided by $HV^{\mathrm{w}}_{\mathcal{R}}(\cdot)$ (budget $=$ 500 runs of W-RoTS)

Interactively Tuning the value of ℓ in W-RoTS

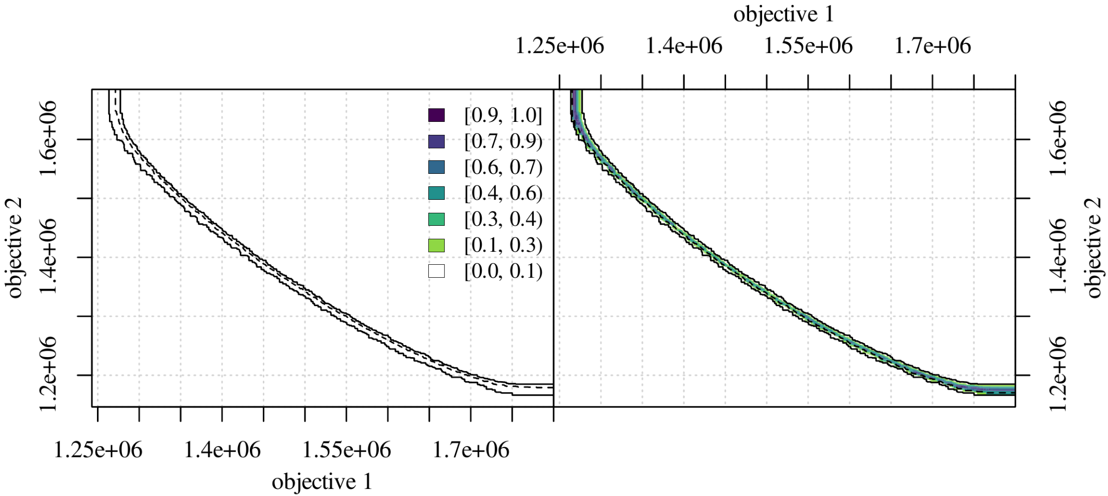

W-RoTS (ℓ = 1) **P1iraceC1 (ℓ = 4)**

Interactively Tuning the value of ℓ in W-RoTS

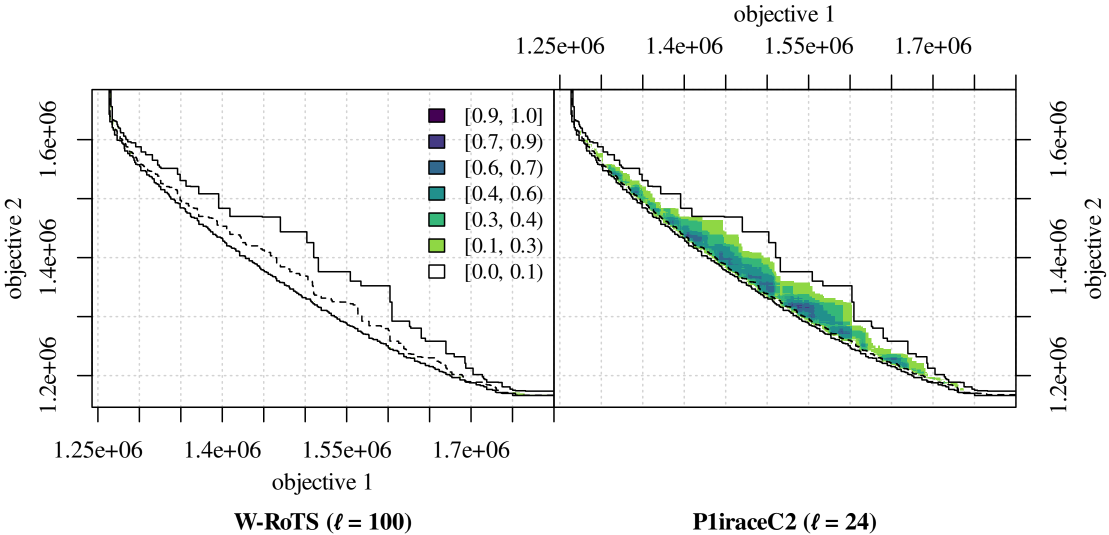

irace is clearly focusing on different regions

A more realistic scenario

[Diaz et al., 2017, 2018]

- Bi-objective multi-product capacitated production planning problem within a failure-prone batch manufacturing system specialised on cleaning products

- Custom MOEA + discrete event simulator

- MOEA has 4 categorical, 2 real and 2 integer parameters

A more realistic scenario

- Bi-objective multi-product capacitated production planning problem within a failure-prone batch manufacturing system specialised on cleaning products
- Custom MOEA + discrete event simulator
- MOEA has 4 categorical, 2 real and 2 integer parameters
- Tuning procedure:

(1) Run irace guided by *unweighted* HV (budget $=$ 1000 MOEA runs)

(2) Select 2 configurations showing the largest EAF differences

(3) Show EAF differences to DM and ask to choose a side

(4) Compute regions $\mathcal{R}$ in favour of chosen side

(5) Create $HV^{\mathsf{w}}_{\mathcal{R}}(\cdot)$ based on chosen regions

(6) Continue running irace guided by $HV^{\mathsf{w}}_{\mathcal{R}}(\cdot)$ (budget $=$ 1000 MOEA runs)

Interactive Tuning Procedure

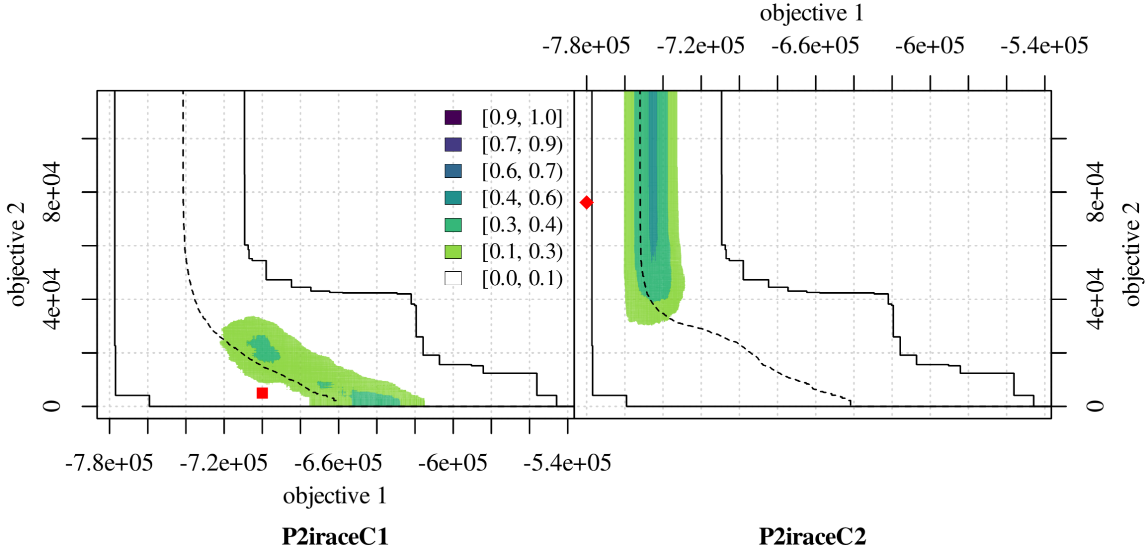

irace is clearly focusing on different regions

Not Only Bi-objective Algorithms

- SCIP (single-objective MIP solver, 200 parameters)
- **Goal:** improve anytime behavior using irace

[López-Ibáñez & Stützle, 2014]

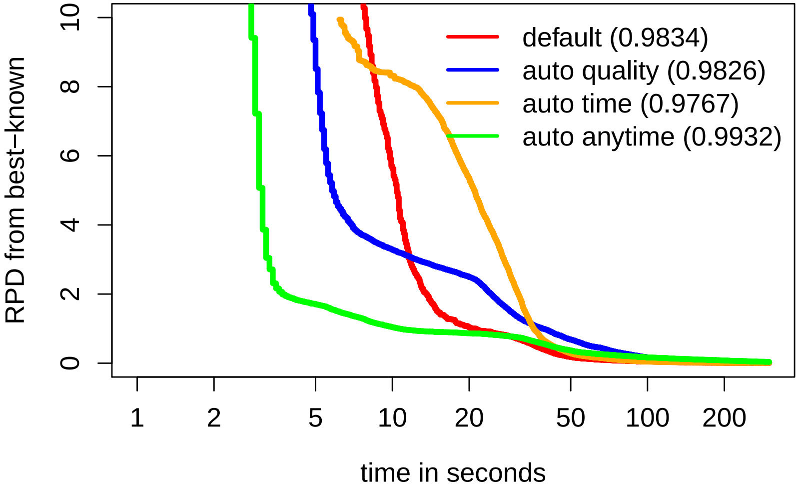

Conclusions and Future Work

- Automatic configuration and design of multi-objective algorithms
  ☞ This is the past! Use it now! [Bezerra et al., 2020b]
- irace + EAF diffs + weighted hypervolume
  = *Interactive configuration and design of bi-objective optimizers*
  ☞ This is the present !

✔ More control of the DM over the resulting algorithm design

Conclusions and Future Work

- Automatic configuration and design of multi-objective algorithms
  ☞ This is the past! Use it now! [Bezerra et al., 2020b]
- irace + EAF diffs + weighted hypervolume
  = *Interactive configuration and design of bi-objective optimizers*
  ☞ This is the present !

✔ More control of the DM over the resulting algorithm design

📋 TODO list:

? How to aggregate over subsequent DM interactions?

? How to aggregate preferences over very different problem instances?

? How to interact with ≥ 3 objectives?

Multi-objective Bayesian Optimization

M. Binois, D. Ginsbourger, & O. Roustant.
Quantifying uncertainty on Pareto fronts with Gaussian process conditional simulations.
European Journal of Operational Research, 243(2):386–394, 2015

☞ Conditional Pareto fronts = Nondominated sets sampled from the GPs

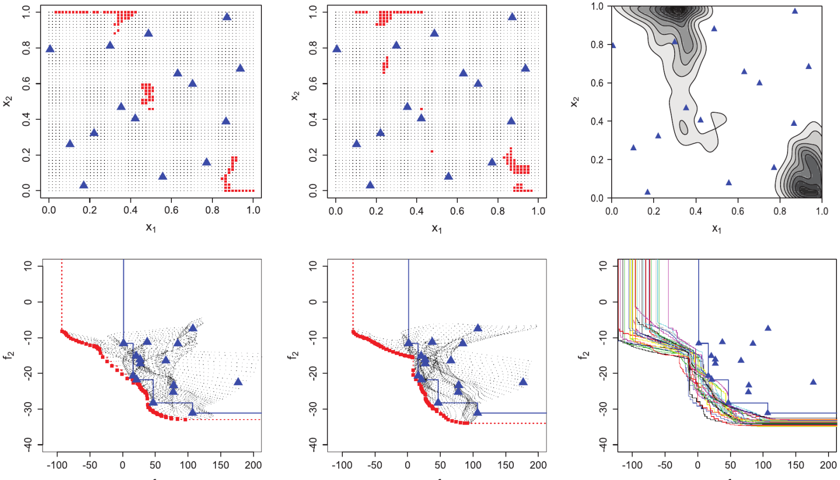

Multi-objective Bayesian Optimization

M. Binois, D. Ginsbourger, & O. Roustant.
Quantifying uncertainty on Pareto fronts with Gaussian process conditional simulations.
European Journal of Operational Research, 243(2):386–394, 2015

☞ Conditional Pareto fronts ⇒ EAF

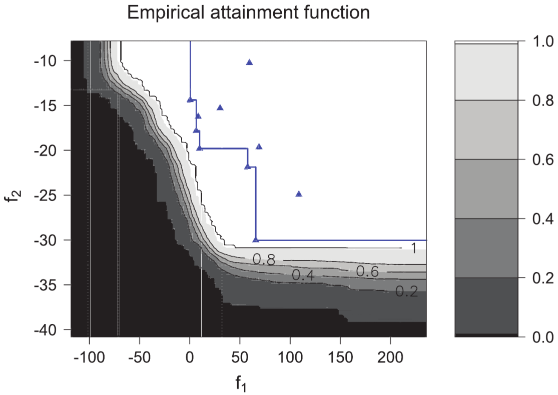

Interactive Hyper-parameter Tuning of Multi-objective Bayesian Optimizers

- Conditional Pareto fronts ⇒ EAF
- Alternative hyper-parameter settings: kernels, AF, optimizer, ...
- EAF differences ⇒ Let the DM choose which uncertainty is preferable

  ⇒ Tune accordingly

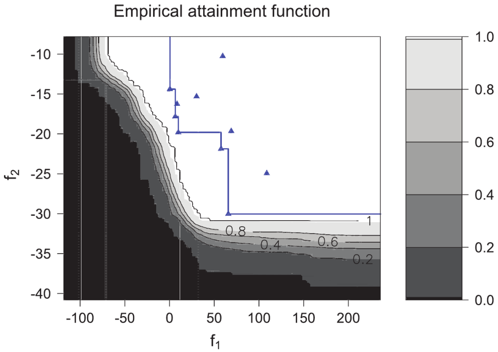

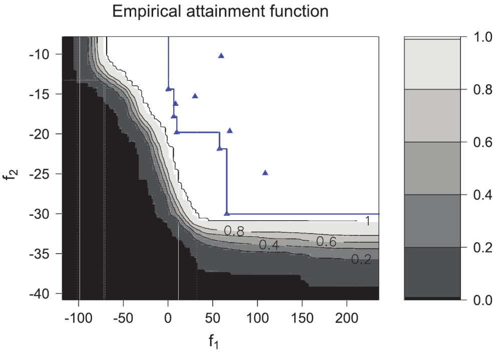


QUIT LIVING IN THE PAST
THE FUTURE IS TODAY

Incorporating Decision-Maker's Preferences into the Automatic Configuration of Bi-objective Optimisation Algorithms

Manuel López-Ibáñez

manuel.lopez-ibanez@manchester.ac.uk
http://lopez-ibanez.eu

The Joint Lectures on Evolutionary Algorithms (JoLEA)

15 February 2023

The University of Manchester
Alliance Manchester Business School

References I

T. Bartz-Beielstein, C. Lasarczyk, and M. Preuss. Sequential parameter optimization. In *Proceedings of the 2005 Congress on Evolutionary Computation (CEC 2005)*, pages 773–780, Piscataway, NJ, Sept. 2005. IEEE Press.

T. Bartz-Beielstein, M. Chiarandini, L. Paquete, and M. Preuss, editors. *Experimental Methods for the Analysis of Optimization Algorithms*. Springer, Berlin, Germany, 2010.

L. C. T. Bezerra, M. López-Ibáñez, and T. Stützle. Automatic component-wise design of multi-objective evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 20(3):403–417, 2016. doi: 10.1109/TEVC.2015.2474158.

L. C. T. Bezerra, M. López-Ibáñez, and T. Stützle. Automatically designing state-of-the-art multi- and many-objective evolutionary algorithms. *Evolutionary Computation*, 28(2):195–226, 2020a. doi: 10.1162/evco_a_00263.

L. C. T. Bezerra, M. López-Ibáñez, and T. Stützle. Automatic configuration of multi-objective optimizers and multi-objective configuration. In T. Bartz-Beielstein, B. Filipič, P. Korošec, and E.-G. Talbi, editors, *High-Performance Simulation-Based Optimization*, pages 69–92. Springer International Publishing, Cham, Switzerland, 2020b. doi: 10.1007/978-3-030-18764-4_4.

B. Bischl, M. Lang, J. Bossek, L. Judt, J. Richter, T. Kuehn, and E. Studerus. *mlr: Machine Learning in R*, 2013. URL http://cran.r-project.org/package=mlr. R package.

B. Bischl, M. Lang, L. Kotthoff, J. Schiffner, J. Richter, E. Studerus, G. Casalicchio, and Z. M. Jones. mlr: Machine learning in R. *Journal of Machine Learning Research*, 17(170):1–5, 2016.

N. Dang and C. Doerr. Hyper-parameter tuning for the $(1 + (\lambda, \lambda))$ GA. In M. López-Ibáñez, A. Auger, and T. Stützle, editors, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2019*, pages 889–897. ACM Press, New York, NY, 2019. ISBN 978-1-4503-6111-8. doi: 10.1145/3321707.3321725.

J. E. Diaz, J. Handl, and D.-L. Xu. Evolutionary robust optimization in production planning: interactions between number of objectives, sample size and choice of robustness measure. *Computers & Operations Research*, 79:266–278, 2017. doi: 10.1016/j.cor.2016.06.020.

J. E. Diaz, J. Handl, and D.-L. Xu. Integrating meta-heuristics, simulation and exact techniques for production planning of a failure-prone manufacturing system. *European Journal of Operational Research*, 266(3):976–989, 2018. ISSN 0377-2217. doi: 10.1016/j.ejor.2017.10.062.

J. Dubois-Lacoste, M. López-Ibáñez, and T. Stützle. A hybrid TP+PLS algorithm for bi-objective flow-shop scheduling problems. *Computers & Operations Research*, 38(8):1219–1236, 2011. doi: 10.1016/j.cor.2010.10.008.

References II

C. Fawcett, M. Helmert, H. H. Hoos, E. Karpas, G. Röger, and J. Seipp. FD-Autotune: Domain-specific configuration using fast-downward. In Karpas et al. [2011].

M. Feurer, A. Klein, K. Eggensperger, J. Springenberg, M. Blum, and F. Hutter. Efficient and robust automated machine learning. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NIPS 28)*, pages 2962–2970, 2015. URL http://papers.nips.cc/book/advances-in-neural-information-processing-systems-28-2015.

G. Francesca, M. Brambilla, A. Brutschy, L. Garattoni, R. Miletitch, G. Podevijn, A. Reina, T. Soleymani, M. Salvaro, C. Pinciroli, F. Mascia, V. Trianni, and M. Birattari. AutoMoDe-Chocolate: Automatic design of control software for robot swarms. *Swarm Intelligence*, 2015. doi: 10.1007/s11721-015-0107-9.

T. Friedrich, F. Quinzan, and M. Wagner. Escaping large deceptive basins of attraction with heavy-tailed mutation operators. In H. E. Aguirre and K. Takadama, editors, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 293–300. ACM Press, New York, NY, 2018. doi: 10.1145/3205455.3205515.

V. Grunert da Fonseca, C. M. Fonseca, and A. O. Hall. Inferential performance assessment of stochastic optimisers and the attainment function. In E. Zitzler, K. Deb, L. Thiele, C. A. Coello Coello, and D. Corne, editors, *Evolutionary Multi-criterion Optimization, EMO 2001*, volume 1993 of *Lecture Notes in Computer Science*, pages 213–225. Springer, Heidelberg, 2001. doi: 10.1007/3-540-44719-9_15.

H. H. Hoos. Programming by optimization. *Communications of the ACM*, 55(2):70–80, Feb. 2012. doi: 10.1145/2076450.2076469.

F. Hutter, H. H. Hoos, and T. Stützle. Automatic algorithm configuration based on local search. In R. C. Holte and A. Howe, editors, *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1152–1157. AAAI Press/MIT Press, Menlo Park, CA, 2007.

F. Hutter, H. H. Hoos, and K. Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In C. A. Coello Coello, editor, *Learning and Intelligent Optimization, 5th International Conference, LION 5*, volume 6683 of *Lecture Notes in Computer Science*, pages 507–523. Springer, Heidelberg, 2011. doi: 10.1007/978-3-642-25566-3_40.

E. Karpas, S. Jiménez Celorrio, and S. Kambhampati, editors. *Proceedings of the 3rd Workshop on Learning and Planning, collocated with the 21st International Conference on Automated Planning and Scheduling (ICAPS-PAL'11)*, 2011.

J. D. Knowles and D. Corne. Properties of an adaptive archiving algorithm for storing nondominated vectors. *IEEE Transactions on Evolutionary Computation*, 7(2):100–116, Apr. 2003.

References III

L. Kotthoff, C. Thornton, H. H. Hoos, F. Hutter, and K. Leyton-Brown. Auto-WEKA 2.0: Automatic model selection and hyperparameter optimization in WEKA. *Journal of Machine Learning Research*, 17:1–5, 2016.

M. Lang, H. Kotthaus, P. Marwedel, C. Weihs, J. Rahnenführer, and B. Bischl. Automatic model selection for high-dimensional survival analysis. *Journal of Statistical Computation and Simulation*, 85(1):62–76, 2014. doi: 10.1080/00949655.2014.929131.

M. López-Ibáñez and J. D. Knowles. Machine decision makers as a laboratory for interactive EMO. In A. Gaspar-Cunha, C. H. Antunes, and C. A. Coello Coello, editors, *Evolutionary Multi-criterion Optimization, EMO 2015 Part II*, volume 9019 of *Lecture Notes in Computer Science*, pages 295–309. Springer, Heidelberg, 2015. doi: 10.1007/978-3-319-15892-1_20.

M. López-Ibáñez and T. Stützle. Automatically improving the anytime behaviour of optimisation algorithms. *European Journal of Operational Research*, 235(3):569–582, 2014. doi: 10.1016/j.ejor.2013.10.043.

M. López-Ibáñez, L. Paquete, and T. Stützle. Hybrid population-based algorithms for the bi-objective quadratic assignment problem. *Journal of Mathematical Modelling and Algorithms*, 5(1):111–137, 2006. doi: 10.1007/s10852-005-9034-x.

M. López-Ibáñez, L. Paquete, and T. Stützle. Exploratory analysis of stochastic local search algorithms in biobjective optimization. In Bartz-Beielstein et al. [2010], pages 209–222. doi: 10.1007/978-3-642-02538-9_9.

M. López-Ibáñez, J. Dubois-Lacoste, T. Stützle, and M. Birattari. The irace package, iterated race for automatic algorithm configuration. Technical Report TR/IRIDIA/2011-004, IRIDIA, Université Libre de Bruxelles, Belgium, 2011. URL http://iridia.ulb.ac.be/IridiaTrSeries/link/IridiaTr2011-004.pdf. Published in Operations Research Perspectives López-Ibáñez et al. [2016].

M. López-Ibáñez, J. Dubois-Lacoste, L. Pérez Cáceres, T. Stützle, and M. Birattari. The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives*, 3:43–58, 2016. doi: 10.1016/j.orp.2016.09.002.

M. López-Ibáñez, J. Branke, and L. Paquete. Reproducibility in evolutionary computation. *ACM Transactions on Evolutionary Learning and Optimization*, 1(4):1–21, 2021. doi: 10.1145/3466624.

L. Pérez Cáceres, F. Pagnozzi, A. Franzin, and T. Stützle. Automatic configuration of GCC using irace: Supplementary material. http://iridia.ulb.ac.be/supp/IridiaSupp2017-009/, 2017.

É. D. Taillard. Robust taboo search for the quadratic assignment problem. *Parallel Computing*, 17(4-5):443–455, 1991.

References IV

M. Vallati, C. Fawcett, A. E. Gerevini, H. H. Hoos, and A. Saetti. Generating fast domain-optimized planners by automatically configuring a generic parameterised planner. In Karpas et al. [2011].

E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. Grunert da Fonseca. Performance assessment of multiobjective optimizers: an analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, 2003. doi: 10.1109/TEVC.2003.810758.

E. Zitzler, D. Brockhoff, and L. Thiele. The hypervolume indicator revisited: On the design of Pareto-compliant indicators via weighted integration. In S. Obayashi et al., editors, *Evolutionary Multi-criterion Optimization, EMO 2007*, volume 4403 of *Lecture Notes in Computer Science*, pages 862–876. Springer, Heidelberg, 2007. doi: 10.1007/978-3-540-70928-2_64.

The irace Package

Manuel López-Ibáñez, Jérémie Dubois-Lacoste, Leslie Pérez Cáceres, Thomas Stützle, and Mauro Birattari.
The irace package: Iterated Racing for Automatic Algorithm Configuration.
Operations Research Perspectives, 3:43–58, 2016. doi:10.1016/j.orp.2016.09.002
http://iridia.ulb.ac.be/irace

The irace Package

Manuel López-Ibáñez, Jérémie Dubois-Lacoste, Leslie Pérez Cáceres, Thomas Stützle, and Mauro Birattari.
The irace package: Iterated Racing for Automatic Algorithm Configuration.
Operations Research Perspectives, 3:43–58, 2016. doi:10.1016/j.orp.2016.09.002
http://iridia.ulb.ac.be/irace

- Implementation of (Elitist) Iterated Racing in R

  Goal 1: Flexible

  Goal 2: Easy to use

The irace Package

Manuel López-Ibáñez, Jérémie Dubois-Lacoste, Leslie Pérez Cáceres, Thomas Stützle, and Mauro Birattari.
The irace package: Iterated Racing for Automatic Algorithm Configuration.
Operations Research Perspectives, 3:43–58, 2016. doi:10.1016/j.orp.2016.09.002
http://iridia.ulb.ac.be/irace

- Implementation of (Elitist) Iterated Racing in R

  Goal 1: Flexible

  Goal 2: Easy to use

- R package available at CRAN (GNU/Linux, Windows, OSX)

The irace Package

Manuel López-Ibáñez, Jérémie Dubois-Lacoste, Leslie Pérez Cáceres, Thomas Stützle, and Mauro Birattari.
The irace package: Iterated Racing for Automatic Algorithm Configuration.
Operations Research Perspectives, 3:43–58, 2016. doi:10.1016/j.orp.2016.09.002
http://iridia.ulb.ac.be/irace

- Implementation of (Elitist) Iterated Racing in R

  Goal 1: Flexible

  Goal 2: Easy to use

- R package available at CRAN (GNU/Linux, Windows, OSX)

- Use it through the command-line: (see `irace --help`)

```
irace --max-experiments 1000 --param-file parameters.txt
```

✔ No knowledge of R needed

- Google group: https://groups.google.com/forum/#!forum/irace-package

An overview of applications of irace

- Parameter tuning
  - Exact MIP solvers: CPLEX, SCIP [López-Ibáñez & Stützle, 2014]
  - Single-objective optimization metaheuristics
  - Multi-objective optimization metaheuristics
  - Anytime optimization (improve time-quality trade-offs)
  - Command-line flags of GCC compiler [Pérez Cáceres et al., 2017]
- Automatic algorithm design
  - From a flexible framework of algorithm components
  - From a grammar description
- Machine learning (AutoML)
  - Automatic model selection for survival analysis [Lang et al., 2014]
  - **mlr** R package [Bischl et al., 2013, 2016]
- Design of control software for robots [Francesca et al., 2015]
- *Theoretical research* *[Friedrich et al., 2018; Dang & Doerr, 2019]*

1197 citations in Google Scholar, 120 000 downloads

EAF differences $\Rightarrow$ Weighted HV

Use EAF differences to define the weighted HV

$\mathcal{R}$: "regions" with EAF differences in favor of the choice made by the DM

A : a new approximation front

$$HV^{\mathsf{w}}(A) = HV(A) + HV^{\mathsf{w}}_{\mathcal{R}}(A) \cdot \beta$$

where $\qquad \beta = (r_1 - f_1^{\min}) \cdot (r_2 - f_2^{\min}) \cdot \varphi$

where $\qquad \varphi \in [0, 1]$ is a scaling factor

Computation of regions from EAF differences

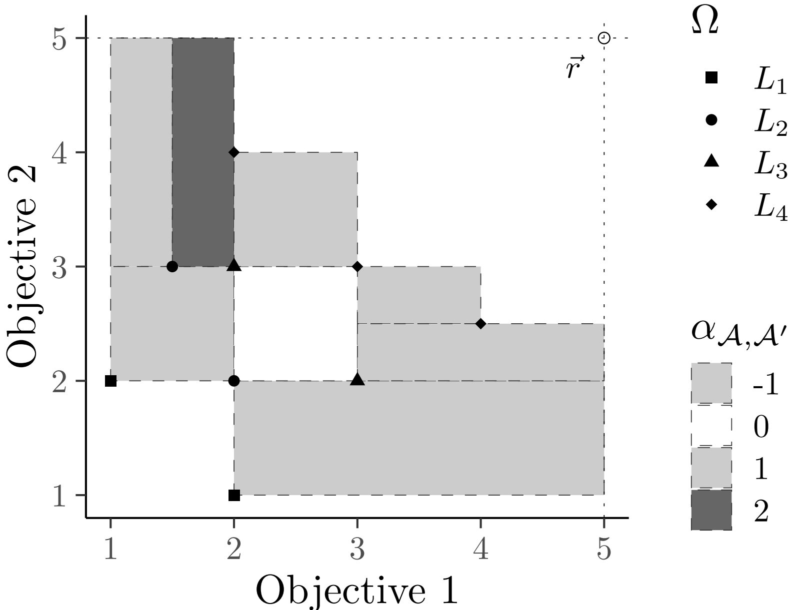

From EAF differences ($\mathcal{R}$) to $HV^{\mathsf{w}}_{\mathcal{R}}(A)$

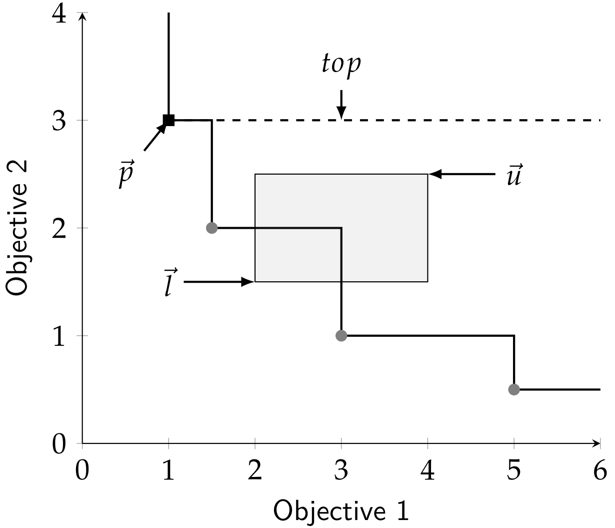

Case #1

From EAF differences ($\mathcal{R}$) to $HV^{\mathsf{w}}_{\mathcal{R}}(A)$

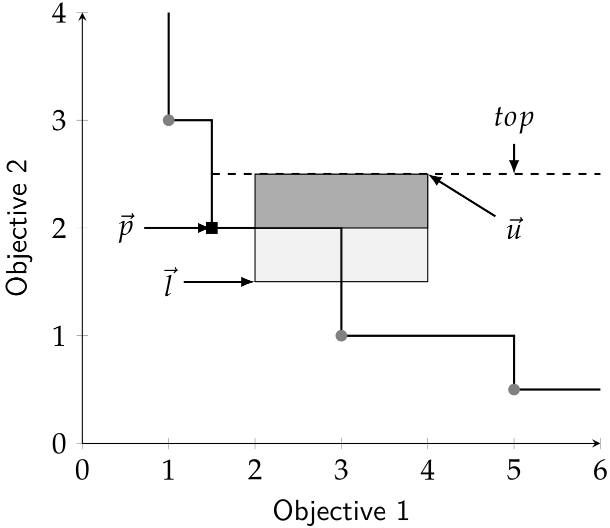

Case #2 bounded by u_2

From EAF differences ($\mathcal{R}$) to $HV^{\mathrm{w}}_{\mathcal{R}}(A)$

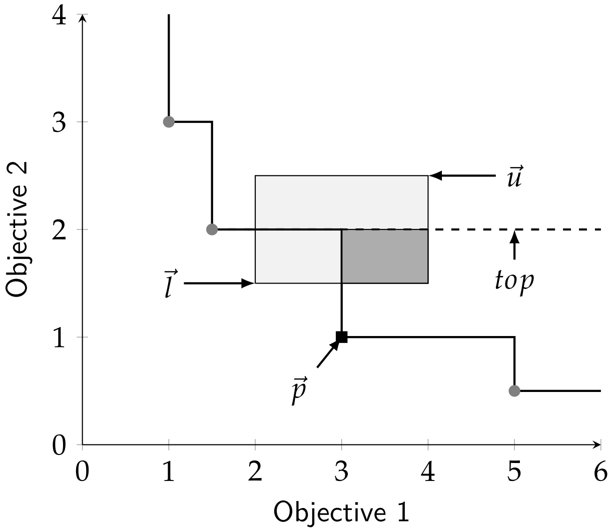

Case #2 bounded by top

From EAF differences ($\mathcal{R}$) to $HV_{\mathcal{R}}^{\mathrm{w}}(A)$

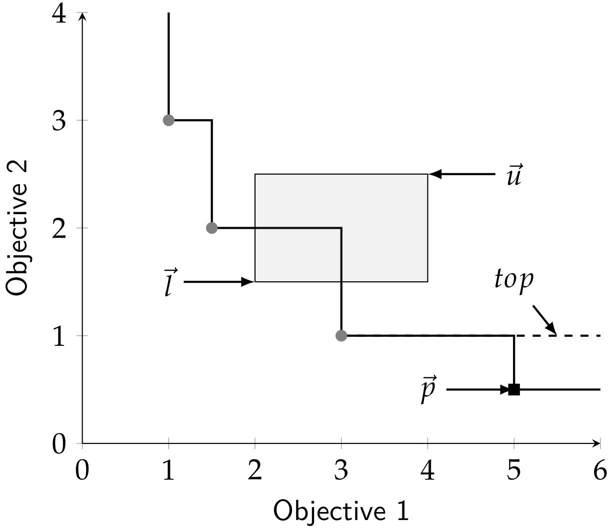

Case #3

Example of the effect of β

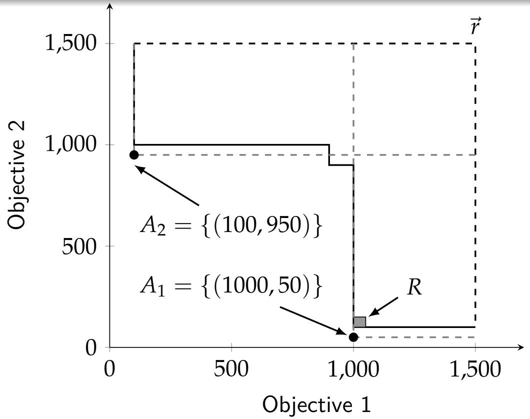

- If $\beta = 1$,

  then $HV^{\mathsf{w}}(A_1) = 730000 < HV^{\mathsf{w}}(A_2) = 770000$

- If $\beta = 0.5 \cdot (r_1 - f_1^{\mathsf{min}}) \cdot (r_2 - f_2^{\mathsf{min}})$

  then $HV^{\mathsf{w}}(A_1) = 2813225000 > HV^{\mathsf{w}}(A_2) = 770000$

Largest EAF differences

For each unique pair $\{\mathcal{A}, \mathcal{A}'\}$, where $\mathcal{A}$ and $\mathcal{A}'$ are the multiset of approximation fronts produced by two different algorithm configurations

1. Compute the EAF difference regions in favour of either side $\mathcal{R}_\mathcal{A}$ for $\Delta\alpha_{\mathcal{A},\mathcal{A}'}(\cdot)$ and $\mathcal{R}_{\mathcal{A}'}$ for $-\Delta\alpha_{\mathcal{A},\mathcal{A}'}(\cdot)$
2. Compute the total area of each set of regions $\mathcal{R}_\mathcal{A}$ and $\mathcal{R}_{\mathcal{A}'}$ by calculating $HV^{\mathrm{w}}_{\mathcal{R}_\mathcal{A}}(A^T)$ and $HV^{\mathrm{w}}_{\mathcal{R}_{\mathcal{A}'}}(A^T)$,
where $A^T = \{(f_1^{\min}, f_2^{\min})\}$
3. The DM is shown the EAF differences of the pair $\{\mathcal{A}, \mathcal{A}'\}$:

$$\max_{\forall\{\mathcal{A},\mathcal{A}'\}} \{\min\{HV^{\mathrm{w}}_{\mathcal{R}_\mathcal{A}}(A^T), HV^{\mathrm{w}}_{\mathcal{R}_{\mathcal{A}'}}(A^T)\}\}$$

4. If the pair selected is not interesting to the DM, or the DM does not feel confident in choosing either side, then the second best pair that satisfies the previous requirement (if such a pair exists) could be shown to the DM.

Boxplots

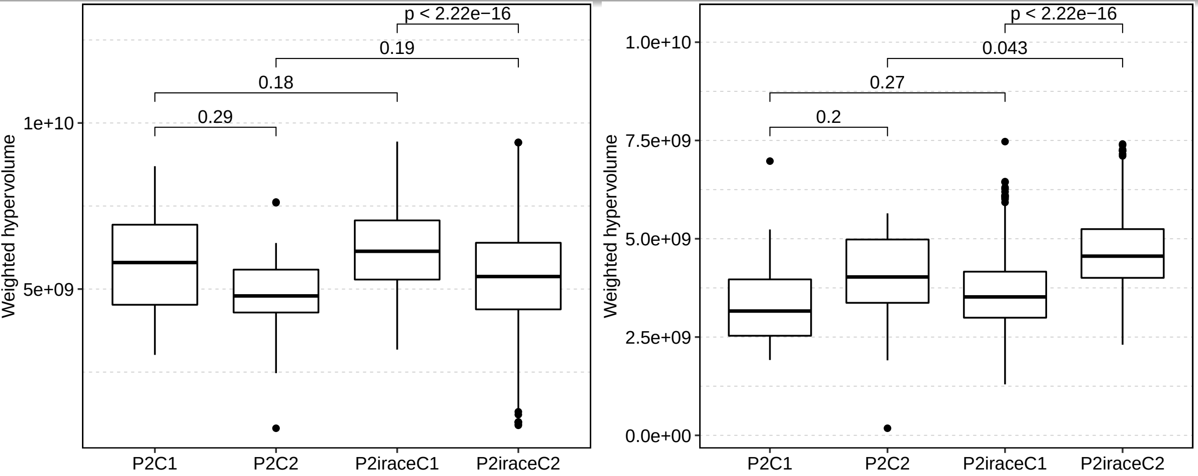

(a) $w_{\text{goal,C1}}$ with $\vec{g} = (-700\,000, 5000)$ (b) $w_{\text{goal,C2}}$ with $\vec{g} = (-780\,000, 76\,200)$

- The same approximation fronts are evaluated on both subfigures.
- Sets P2C1 and P2C2 contain 16 fronts each, i.e., 16 runs of one configuration of the MOEA.
- Sets P2iraceC1 and P2iraceC2 contain 900 fronts each, i.e., 30 runs of the best configuration of the MOEA returned by

A more realistic scenario

[Diaz et al., 2017, 2018]

- Bi-objective multi-product capacitated production planning problem within a failure-prone batch manufacturing system specialised on cleaning products
- Custom MOEA + discrete event simulator
- MOEA has 4 categorical , 2 real and 2 integer parameters

Interactive Tuning Procedure

(1) Run irace guided by *unweighted* HV (budget = 1000 MOEA runs)

(2) Select 2 configurations showing the largest EAF differences

(3) Show EAF differences between 2 configurations
and ask DM to choose a side

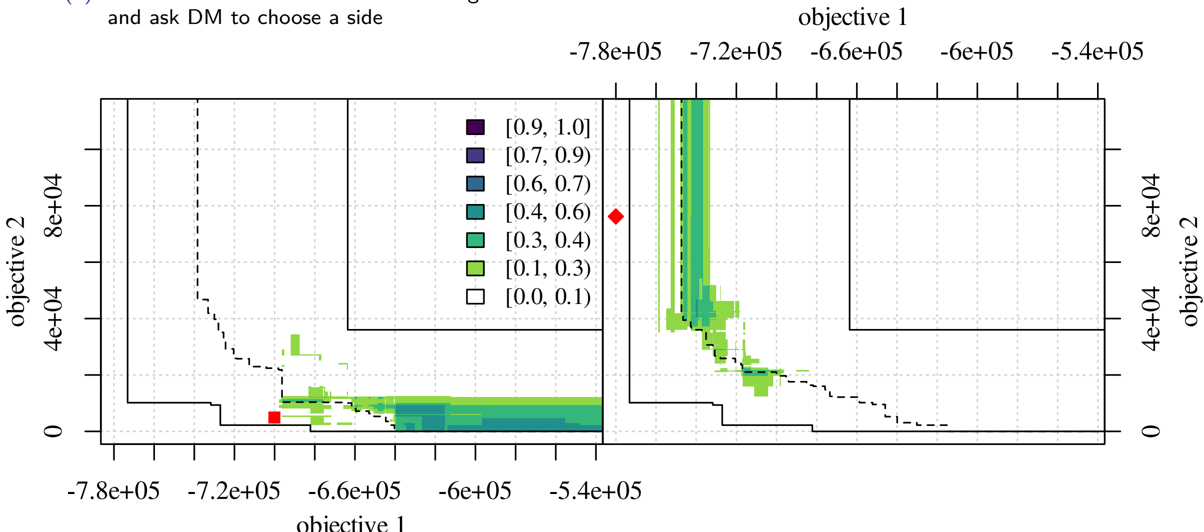

Interactive Tuning Procedure

(4) Compute regions $\mathcal{R}$ in favour of chosen side

(5) Create $HV^{\mathrm{w}}_{\mathcal{R}}(\cdot)$ based on chosen regions

(6) Continue running irace guided by $HV^{\mathrm{w}}_{\mathcal{R}}(\cdot)$
(budget = 1000 MOEA runs)

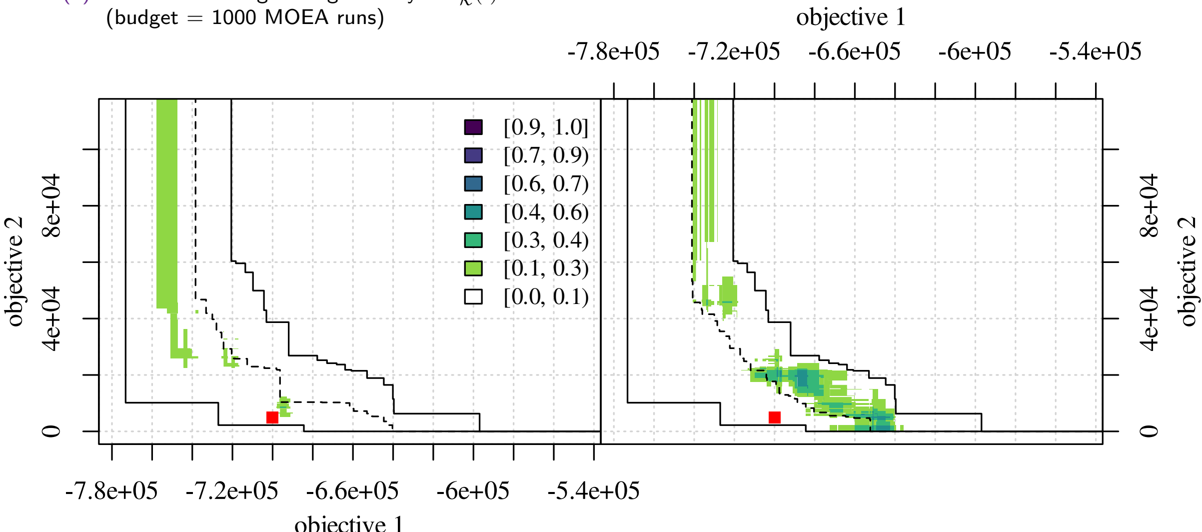

Interactive Tuning Procedure

(4) Compute regions $\mathcal{R}$ in favour of chosen side

(5) Create $HV^{\mathrm{w}}_{\mathcal{R}}(\cdot)$ based on chosen regions

(6) Continue running irace guided by $HV^{\mathrm{w}}_{\mathcal{R}}(\cdot)$

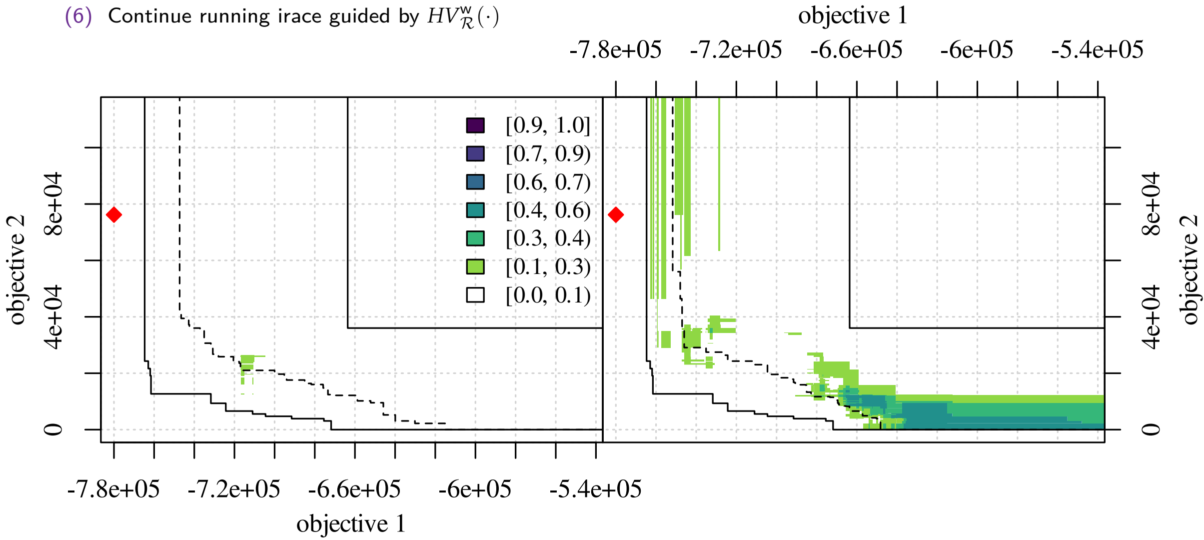

Interactive Tuning Procedure