# The `GPareto` and `GPGame` packages for multi and many objective Bayesian optimization

UseR! 2019, Toulouse

Mickael Binois (Inria Sophia) and Victor Picheny (Prowler.io)
mickael.binois@inria.fr

12/07/2019

# Bayesian Optimization context

Black-box model with multiple outputs:

$$f : x \in \mathbf{X} \subset \mathbf{R}^d \mapsto \mathbf{R}^p$$

Working hypotheses: $f$ is expensive to compute, with complex outputs:

- non-convex
- no derivatives available
- possible observation noise
- $2 \leq p \leq 20$
- $\mathbf{X}$ is typically a box of dimension $2 \leq d \leq 100$

Examples:

- engineering design applications
- hyperparameter tuning in Machine Learning

# Multi-objective optimization

$$\text{MOP}: \begin{cases} \min_{x\in\mathbf{X}} & f_1(x) \\ \quad\vdots & \\ \min_{x\in\mathbf{X}} & f_p(x) \end{cases}$$

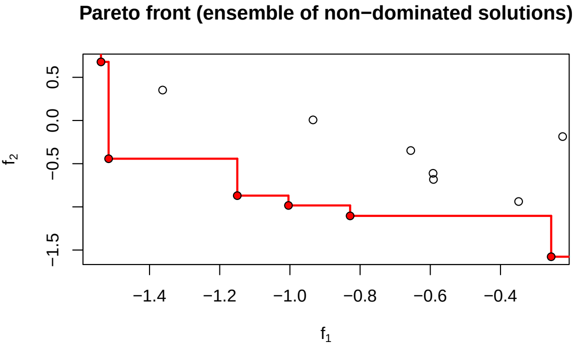


**Classical multi-objective algorithm goal**: obtaining a "good" discrete approximation of the set of non-dominated solutions (Pareto set and front)

# Bayesian optimization (BO) in a nutshell

**BO**: sequential design strategy based on a distribution over functions to define an acquisition function.

Two ingredients:

- fast surrogate (or metamodel) of the objectives
- infill criterion adapted to the problem at hand

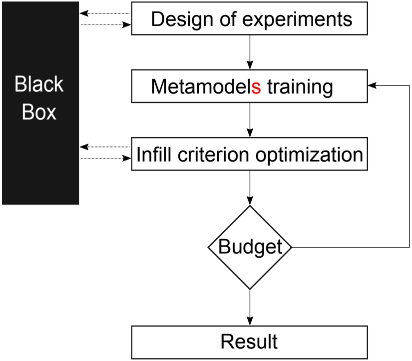

# Gaussian process (GP) regression

GPs make popular surrogates, in particular with their uncertainty quantification and interpolation capabilities.

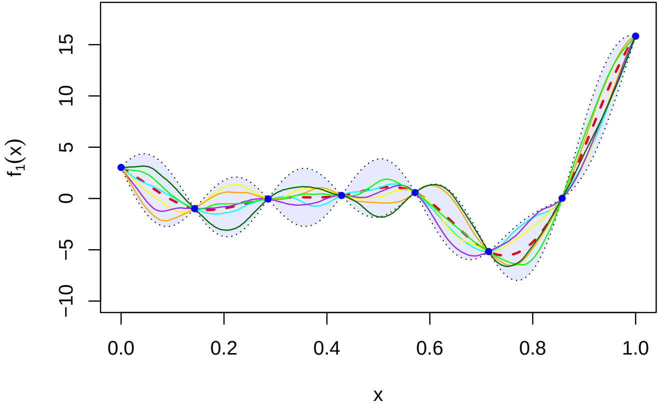


`DiceKriging` is used for GP regression here.

# A word on MO acquisition functions

Some are based on a notion of improvement:

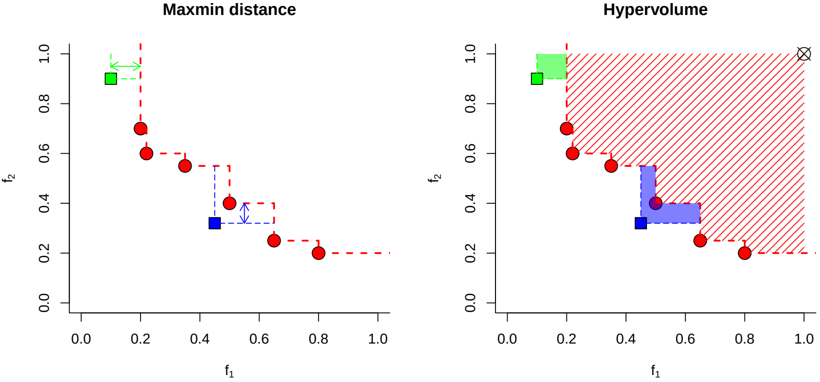


Other on the notion of variance (or entropy) of a given quantity.

# Example 1: bi-objective optimization

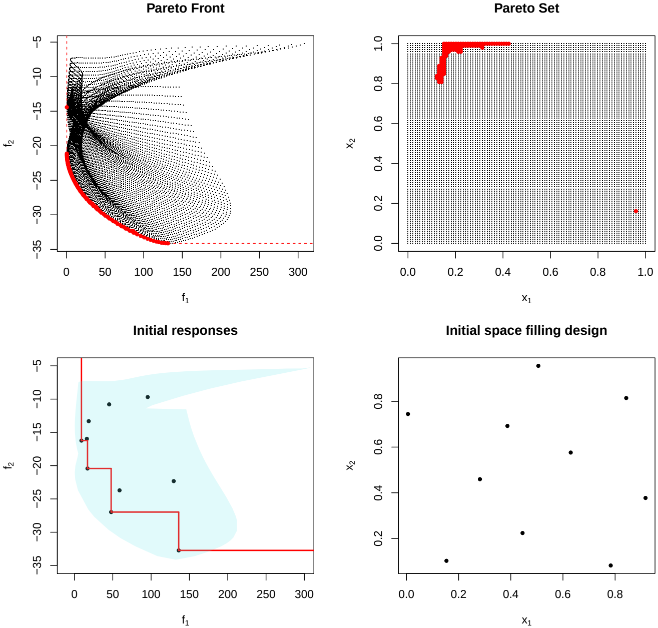

# Example 1: bi-objective optimization (cont'd)

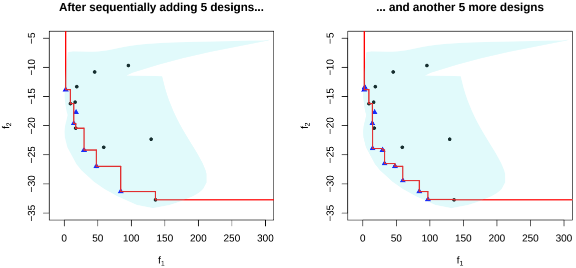


Some postprocessings are available as well:

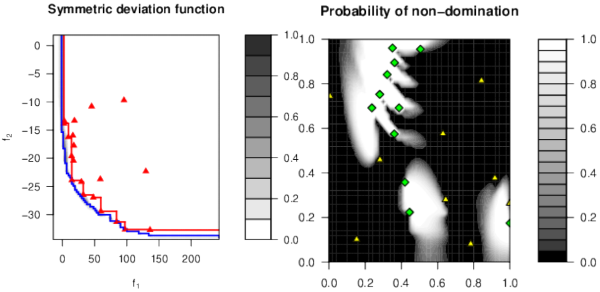

# Example 1: bi-objective optimization (cont'd)

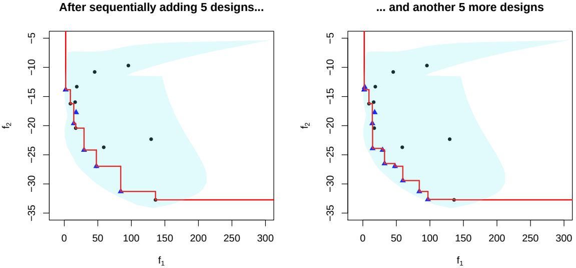


Some postprocessings are available as well:

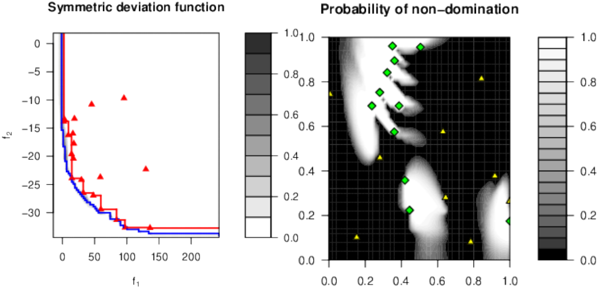

# Test case: rear shock absorber (d = 47)

Objectives: mass, axial and lateral impacts on empty or charged vehicle

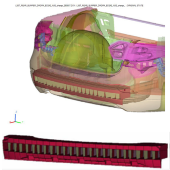

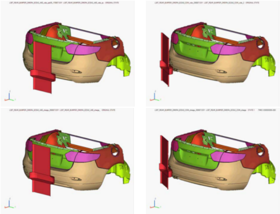

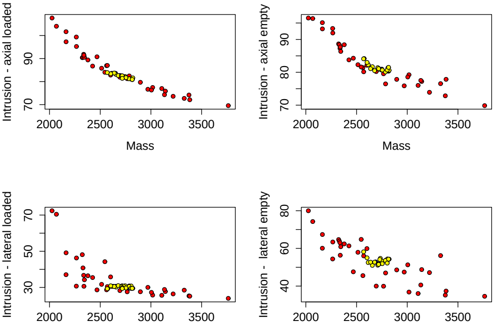

# Many-objective challenges

Taking more objectives is possible, e.g., with 3:

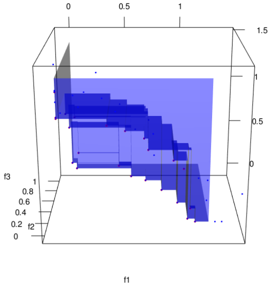


But, as $p$ grows:

- visualization and selection is more complex,
- approximating the Pareto set/front is increasingly difficult,
- the proportion of non-dominated solutions grows quickly.

# Game theoretic perspective on many-objective optimization

With a limited budget, it is more reasonable to focus on a single *good* solution: the Kalai-Smorodinsky (KS) solution.

Objectives are considered as players, aiming to get equal benefit ratios from a disagrement point (e.g., the Nadir point).

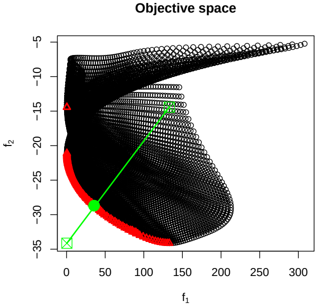


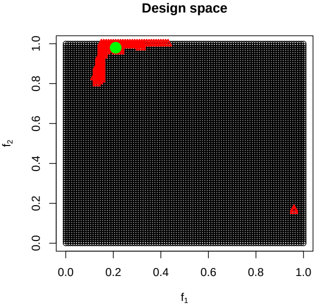

# Calibration of an agent-based behavioral model

13-variable model of behavior of occupants in a building.

The 9 objectives are target values based on record or surveys.

Preferences can be incorporated by defining a custom disagreement point.

Result with 100 initial designs and 100 sequentially added ones:

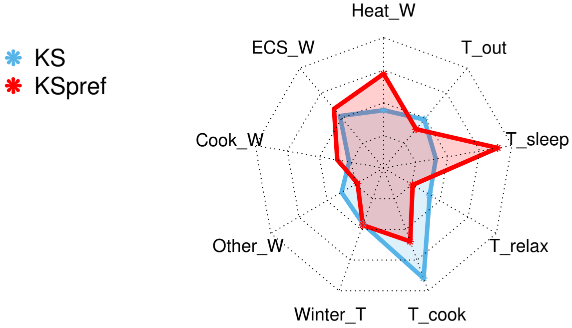

# Calibration of an agent-based behavioral model

13-variable model of behavior of occupants in a building.

The 9 objectives are target values based on record or surveys.

Preferences can be incorporated by defining a custom disagrement point.

Result with 100 initial designs and 100 sequentially added ones:

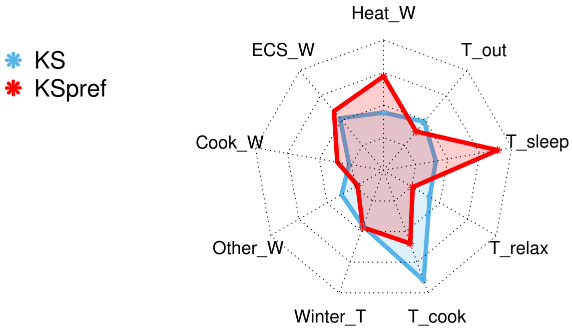

# Thanks

`GPareto` efficiently solves expensive multi-objective optimization problems. Additional post-processing routines are available for further uncertainty quantification on the Pareto front and set.

Complementarily, `GPGame` tackles many-objective optimization with a game theoretic point of view, and can find discrete Nash equilibria.

Both packages are available on CRAN, facilitating everything in this reproducible `Rmarkdown` talk.

- A *Journal of Statistical Software* paper is available for `GPareto`.