Combining game theory and Bayesian optimization to solve many-objective problems

Victor Picheny (*INRA*), Mickael Binois (*Argonne national lab.*),
Abderrahmane Habbal (*Université Côte d'Azur*)

Mascot-Num, Nantes

March 23rd, 2018

Bayesian optimization context

"Black-box" model, multiple outputs

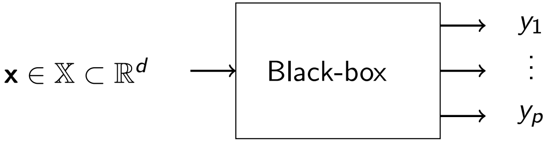

Working hypotheses: expensive to compute, complex y_i's

- non-convex
- no derivatives available
- possibly observed in noise
- $2 \leq p \leq 20$

$\mathbb{X}$: typically a box of dimension $2 \leq d \leq 50$

Multi-objective optimization

$$\begin{cases} \min_{\mathbf{x}\in\mathbb{X}} & y_1(\mathbf{x}) \\ \vdots & \\ \min_{\mathbf{x}\in\mathbb{X}} & y_p(\mathbf{x}) \end{cases}$$

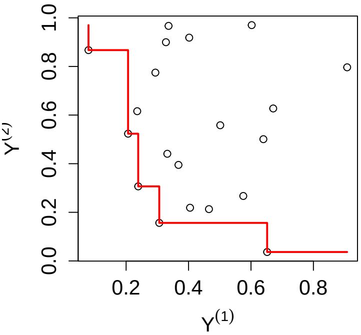

Ensemble of non-dominated solutions

Pareto front (objective space) and Pareto set (design space)

Classical BO algorithm objective

Obtain a good discrete approximation of the Pareto set = non-dominated solutions along the entire Pareto front

Challenging situation: many objectives / restricted budget

Accurate Pareto set may not be attainable

- Many objectives $\Rightarrow$ very large Pareto set
- Not enough budget to cover all the solutions

Accurate Pareto set may not be desirable

Too many solutions to choose from

Our proposition: use **game theory** to seek **compromise solutions**

Game theory literature

Extensive for small discrete problems or convex objectives, scarce for expensive *black-boxes*
$\Rightarrow$ Explore BO alternatives for solving game equilibrium problems

Outline

Game Theory and equilibrium problems

Game theory is the study of mathematical models of conflict and cooperation between intelligent rational decision-makers

- Initially from economics
- Now common in artificial intelligence, engineering or machine learning

Game

- Multiple **decision makers** (*players*) with **antagonistic goals**
- Acceptable **compromise** = game equilibrium

Nash equilibrium problems (NEP)

NEP: canonical antagonist game

Each player tries to solve his optimization problem:

$$(\mathcal{P}_i) \quad \min_{\mathbf{x}_i \in \mathbb{X}_i} y_i(\mathbf{x}), \qquad 1 \leq i \leq p$$

with $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_p] \in \mathbb{X}$. We assume here some **territory splitting:**

$$\mathbb{X} = \mathbb{X}_1 \times \dots \times \mathbb{X}_p$$

Definition: writing $\mathbf{x} = (\mathbf{x}_i, \mathbf{x}_{-i})$ (no actual permutation), $\mathbf{x}^*$ is a NE if:

$$\forall i,\ 1 \leq i \leq p, \quad \mathbf{x}_i^* = \arg\min_{\mathbf{x}_i \in \mathbb{X}_i} y_i(\mathbf{x}_i, \mathbf{x}_{-i}^*)$$

$\Rightarrow$ When everyone plays a NE, then no player has incentive to move from it

An engineering example: multidisciplinary optimization

- Each team makes a decision with antagonistic goals
- Natural territory splitting: each team acts on its own set of variables

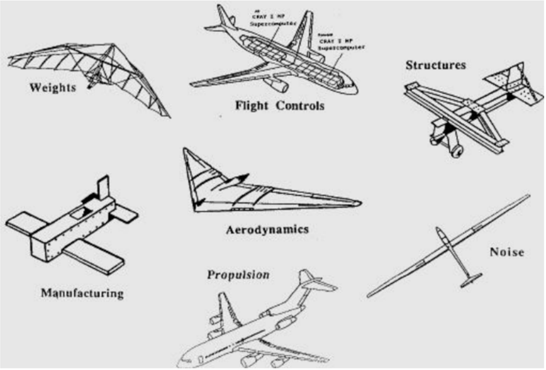

J.A. Désidéri (2012), Cooperation and competition in multidisciplinary optimization, *Comp. Optim. and Applications*

A PDE example: data completion

Steady-state heat equation

$$\begin{cases} \nabla.(k\nabla u) = 0 & \text{in} & \Omega & \\ u = f & \text{on} & \Gamma_c & \text{(Dirichlet)} \\ k\nabla u.\nu = \psi & \text{on} & \Gamma_c & \text{(Neumann)} \end{cases}$$

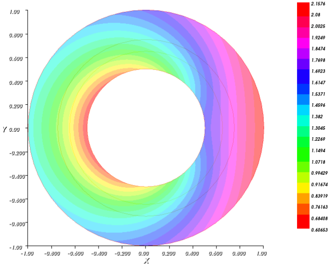

Boundary partially unobservable

- $\partial\Omega = \Gamma_c + \Gamma_i$
- Data to recover: $u_{|\Gamma_i}$ and $k\nabla u.\nu_{|\Gamma_i}$

Nash game: Dirichlet vs. Neumann

A. Habbal and M. Kallel (2013), Neumann-Dirichlet Nash strategies for the solution of elliptic Cauchy problems, *SIAM J. Control Optim*

The Kalai-Smorodinsky solution

Disadvantages of Nash equilibria

- Assumes some *territory splitting*
- In general **not efficient** (in the Pareto sense)

The Kalai-Smorodinsky (KS) idea

- Players start negociation from a *disagreement point*
- Progress towards an *efficient* solution while ensuring *equity of marginal gains*

KS solution

Intersection of the disagreement point - *ideal point* (= utopia or shadow) straight line with the Pareto front

Kalai, Smorodinsky (1975). Other solutions to Nash's bargaining problem, *Econometrica*

Illustration: 2 objectives, 2 variables

y_1

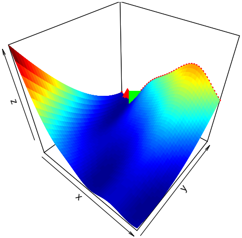

y_2

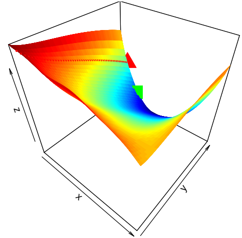

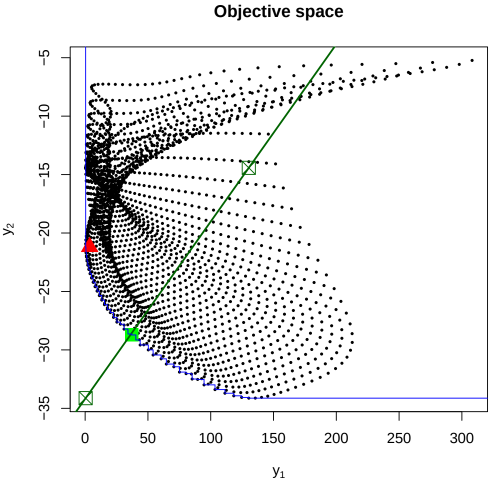

Here: disagreement point = Nadir

Some properties

For 2 players case, KS is proved to be the unique bargaining solution which fulfills the following

Axioms

- Pareto optimality
- Symmetry
- Invariance w.r.t. affine transformations
- Restricted Monotonicity

For N players

- Efficiency, Symmetry, Affine invariance
- Effective selection device even when combined with refinement concepts based on stability with respect to perturbations [1].

[1] De Marco G. and Morgan J. (2010). Kalai-Smorodinsky bargaining solution equilibria. JOTA, 145(3), 429-449.

Going further: robust KS using copula

Motivation: unlike Nash, KS is sensitive to non-linear rescaling

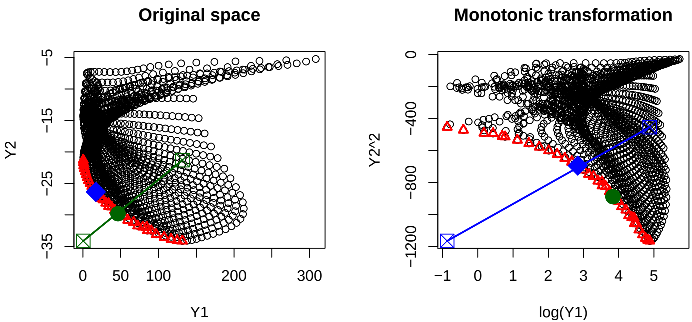

Why copulas?

- Nice link to Pareto optimality (see next slide)
- Insensitive to non-linear transformations on the marginals

Copulas and Pareto optimality

Pareto optimality and multivariate statistics

- Take: $\mathbf{Y} = (y_1(X), \ldots, y_p(X))\,, \qquad X \sim U(\mathbb{X})$
- Pareto front = (part of the) **zero level-line of the multivariate CDF** F_Y
- Multivariate CDF = marginal densities + copulas

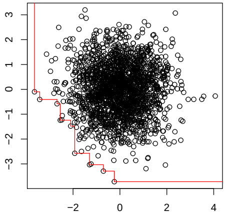

$$F_Y(\mathbf{y}) = C(F_1(y_1), \ldots, F_p(y_p))$$

with $C : \mathbb{R}^p \to \mathbb{R}$ copula, $F_i = \mathbf{P}(Y_i \leq y_i)$ marginal densities

M. Binois, D. Rulliere, D., O. Roustant,
On the estimation of Pareto fronts from the point of view of copula theory
Information Sciences (2015) 324: 270-285.

Our proposition: KS in the copula space (CKS)

Assuming a distribution for X (e.g. $X \sim U(\mathbb{X})$)

We look for the KS of $F_1(y_1(X)), \dots, F_p(y_p(X))$

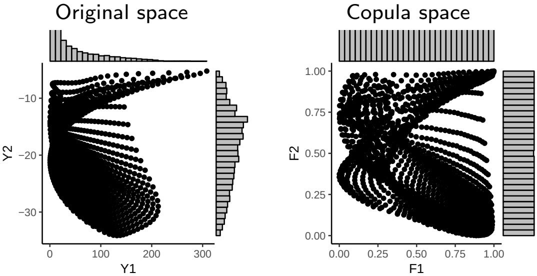

Bonus of copulas: since $F_i\left[\min y_i(X)\right] = 0$, $F_i\left[\max y_i(X)\right] = 1$

- Utopia point is $(0 \dots 0)$
- Worst solution for all objectives is $(1 \dots 1) \Rightarrow$ *new disagreement point*

KS in the copula space (CKS)

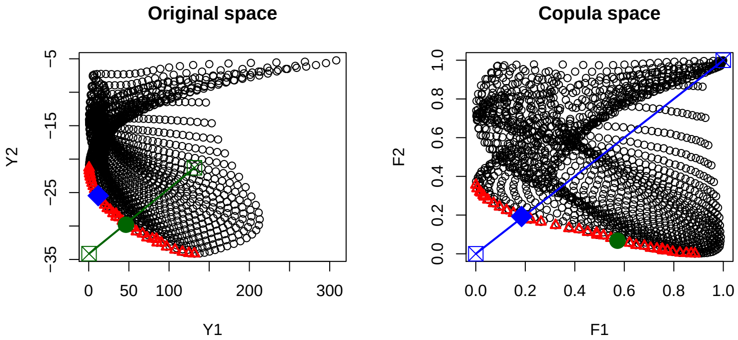

On discrete spaces: rank space

CKS = Pareto optimal solution with **closests ranks among objectives**

Invariant by any (strictly) monotonic transformation of marginals

... but depends on X's distribution

Outline

Bayesian optimization: two ingredients

- Quick-to-evaluate surrogate of the objective: **Gaussian process regression**
- Sequential sampling via maximizing *acquisition functions*

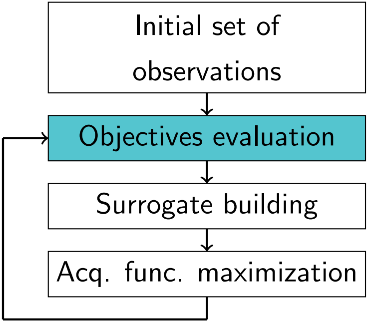

(nested loop)

(nested loop)

Gaussian process regression

Prior π : y realization of a GP Y

Entirely defined by its **mean** $\mathbb{E}_\pi(Y(\mathbf{x}))$ and **covariance** $\text{cov}_\pi(Y(\mathbf{x}), Y(\mathbf{x}'))$

Regression model ($\hat{y}$) = law of the GP conditioned on observations

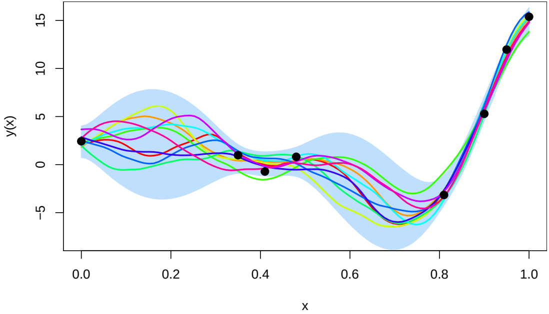

(Parametric form assumed for mean and covariance, estimation by maximum likelihood)

Sequential sampling (EGO algorithm)

The acquisition function balances between *exploration* and *exploitation*

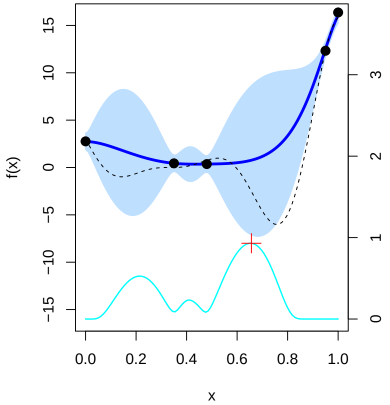

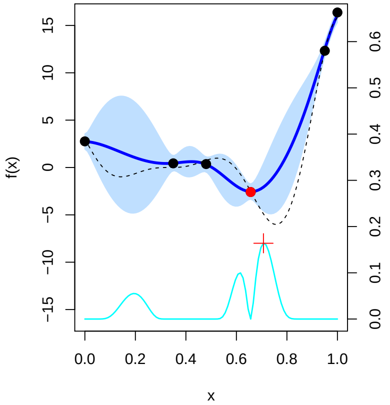

BO and equilibria

Regression: one GP model for each obj y_i (correlation is neglected)

- GP is used to "denoise" the objectives (if needed)
- Multivariate regressor: $\mathbf{Y}(.) \sim \mathcal{GP}\left(\boldsymbol{\mu}(.), \boldsymbol{\Sigma}\left(.,.\right)\right)$ with $\boldsymbol{\Sigma}$ diagonal.

Sequential sampling boils down to finding a "good" acquisition function

$\Rightarrow$ Based on $\mathbf{Y}(.)$, which $\mathbf{x}_{new}$ should I visit next?

Canonical acquisitions: *Expected Improvement*, *Upper Confidence Bound*

Not usable here due to the complex learning task (*measure of progress?*)

- Nash: first-order stationarity
- Kalai-Smorodinsky: part of the Pareto front + Nadir + Shadow
- Copula-Kalai-Smorodinsky: marginals + copula

What we will use: stepwise uncertainty reduction

Villemonteix, Vazquez, Walter (2009) An informational approach to the global optimization of expensive-to-evaluate functions, *J. of Glob. Opt.*

Bect, Ginsbourger, Li, Picheny, Vazquez, (2012) Sequential design of computer experiments for the estimation of a probability of failure, *Statistics and Computing*.

Hernandez-Lobato, Hoffman, Ghahramani, (2014) Predictive entropy search for efficient global optimization of black-box functions, *NIPS*

...

What do we know about the equilibrium?

Computing equilibria on discrete sets is easy

The search for Nash, KS, or CKS can be done by (smart) exhaustive search on a grid.

What we can do

- Discretize $\mathbb{X}$ (grid, LHS...)
- Draw GP samples $\mathcal{Y}_1, \dots, \mathcal{Y}_M$ of $\mathbf{Y}(\mathbb{X}_{\text{disc}})$
- Get (quickly) equilibrium of each $\mathcal{Y}_i$

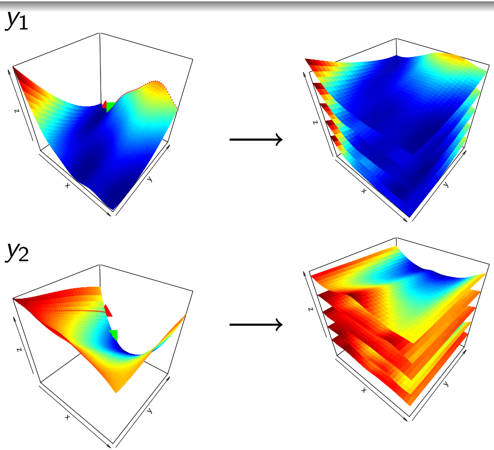

What do we know about the equilibrium?

GP sample ⇒ cloud of potential solutions (+)

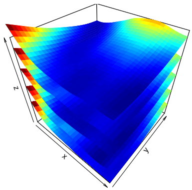

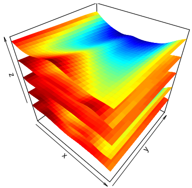

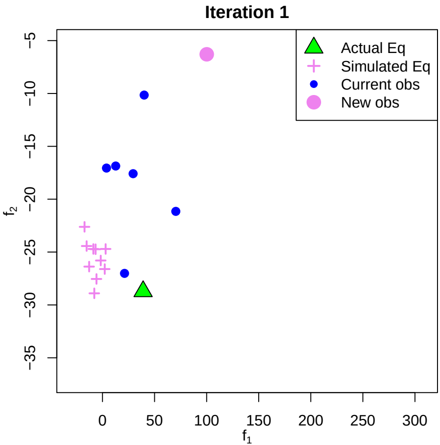

⇒ We want to choose $\mathbf{x}_{new}$ so that the cloud shrinks!

Formalizing: Stepwise uncertainty reduction (SUR)

Random equilibrium

- $\mathbf{y}(\mathbf{x}^*) = \Psi(\mathbf{y})$ (Nash, KS or CKS) equilibrium for cost function $\mathbf{y}$
- $\mathbf{Y}$ random (GP) $\Rightarrow \Psi(\mathbf{Y})$ random ($\rightarrow$ the $+$'s)

Uncertainty measure of equilibrium $\equiv$ "volume" of $\Psi(\mathbf{Y})$:

$$\Gamma(\mathbf{Y}) = \det\left[\mathrm{cov}\left(\Psi(\mathbf{Y})\right)\right]$$

$\Rightarrow \Gamma_k$ if $\mathbf{Y}$ conditioned on observations $(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_k, \mathbf{y}_k)$

Greedy (stepwise) uncertainty reduction

- Optimal choice: $\mathbf{x}_{k+1} = \arg\min \Gamma_{k+1}$
- Out of reach without $\mathbf{y}(\mathbf{x}_{k+1})$!

Acquisition function: **expected** reduction

$$J(\mathbf{x}_{k+1}) = \mathbb{E}_{\mathbf{F}_{k+1}}\left(\Gamma\left[\mathbf{Y} | \mathbf{F}_{k+1} = \mathbf{Y}(\mathbf{x}_{k+1})\right]\right)$$

with $\mathbf{Y}(\mathbf{x}_{k+1}) \sim \mathcal{GP}\left(\boldsymbol{\mu}(\mathbf{x}_{k+1}), \mathbf{\Sigma}\left(\mathbf{x}_{k+1}, \mathbf{x}_{k+1}\right)\right)$

Computing the acquisition function

Acquisition function

$J(\mathbf{x}_{new}) = \mathbb{E}_{\mathbf{F}_{new}} \left(\Gamma \left[\mathbf{Y} | \mathbf{F}_{new} = \mathbf{Y}(\mathbf{x}_{new}) \right] \right)$
with $\mathbf{Y}(\mathbf{x}_{new})$ Gaussian

Two-layer Monte Carlo

- $\mathcal{Y}_1^{(\text{new})}, \dots, \mathcal{Y}_K^{(\text{new})}$ observation drawings at $\mathbf{x}_{new}$
- We condition the paths by these observations: $\boldsymbol{\mathcal{Y}}_m | \mathcal{Y}_k^{(\text{new})}$
- $\hat{J}(\mathbf{x}_{new}) \approx$

$$\frac{1}{K} \sum_{k=1}^{K} \det \left[\operatorname{cov} \left(\boldsymbol{\mathcal{Y}}_1 | \mathcal{Y}_k^{(\text{new})}, \dots, \boldsymbol{\mathcal{Y}}_M | \mathcal{Y}_k^{(\text{new})} \right) \right]$$

$\mathbf{Y}_1, \dots, \mathbf{Y}_K$

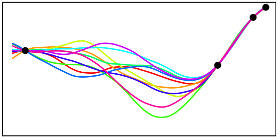

$\mathbf{Y}_1/Y_1^{new}, \dots, \mathbf{Y}_K/Y_1^{new}$

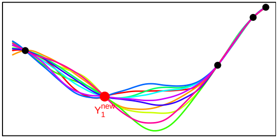

$\mathbf{Y}_1/Y_2^{new}, \dots, \mathbf{Y}_K/Y_2^{new}$

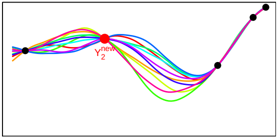

Illustration: iteration 1, 3 and 7 (KS search)

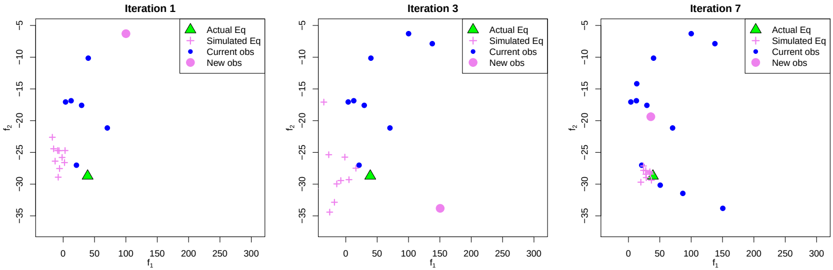

- With only 13 points: KS identified, very small residual uncertainty
- Trade-off between learning the center and the extremities of the Pareto front

In practice: Many numerical tricks are necessary...

- Efficient data structure and discrete games solver
- Use of parallel computation
- Fast algorithm for drawing and updating simulated GP paths

  Chevalier, Emery, Ginsbourger (2015) Fast update of conditional simulation ensembles, *Mathematical Geosciences*

- Small discretized space by choosing only useful points ($\Rightarrow$ shorter paths, smaller games)
- Optimization of the acquisition function only over the discretized space

Outline

li-BIM: an agent-based model for designing buildings

Simulated behavior of occupants in a building

- Numerical modeling of the building: thermal, air quality, lighting, etc.
- Evolved occupational cognitive model (Belief-Desire-Intention)

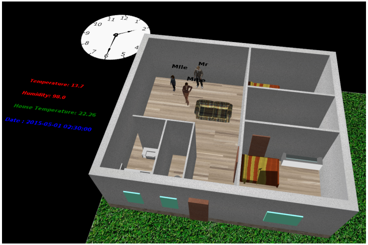

F. Taillandier, A. Micolier, P. Taillandier (2017). Li-BIM (Version 1.0.0), *CoMSES Computational Model Library*

The li-BIM calibration problem

Many parameters to tune

- Behavioral: sensitivity to hot, cold, air quality, hunger, tidiness...
- Appliances characteristics: electrical power, efficiency...

Model should match real data (records or surveys)

- Electrical consumption
- Average temperature, air quality
- Time spent on various activities

A challenging optimization problem

- 11 parameters, 10 objectives
- Expensive, stochastic model (30min / simulation)

Outputs of different nature $\rightarrow$ how to define an acceptable compromise?

Experimental setup

BO configuration

- 100-point LHS + 100 infills
- Space discretization: 1000 points (renewed at each iteration)
- Choosing the next point $\approx$ 2 min
- Objectives: square errors w.r.t. targets (averaged over 8 repetitions)
- KS + CKS (with empirical copula)

Preliminary result

Out of the 300 points, only 5 were dominated

Results - objective space (% error)

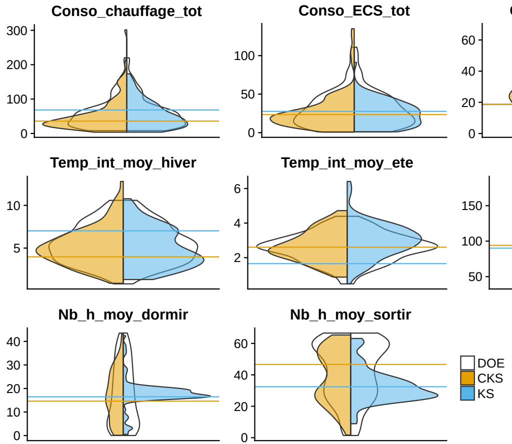

- Apparent trade-off between exploration and exploitation
- CKS more exploratory and sensitive to local mass on marginals

Results - objective space (2D)

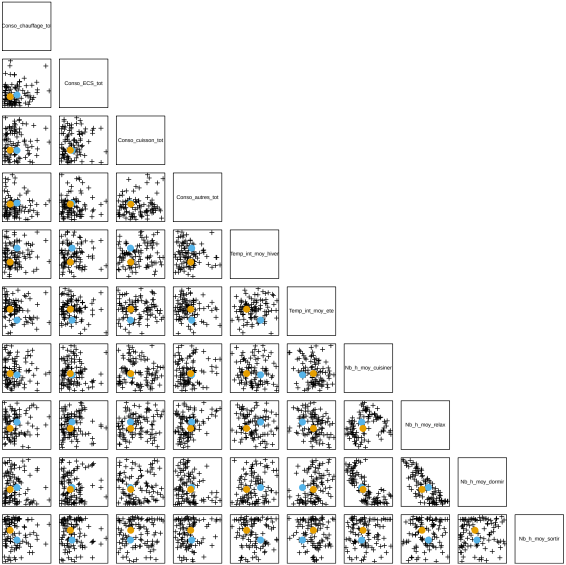

Outline

Summary and future steps

Ingredients

- Game theory concepts to define compromise solutions
- GPs + Stepwise Uncertainty Reduction to find them

$\Rightarrow$ Parsimonious algorithm to tackle automatically many black-box objectives

SUR is very accomodating, so why not also...

- Generalized Nash problems (i.e. with constraints)
- KS with smart disagreement points: e.g. Nash-Kalai-Smorodinsky
- Batch-sequential strategies
- Smart use of repetitions (for stochastic simulators)

References

Want to give it a try? R package GPGame

https://CRAN.R-project.org/package=GPGame

Implements Nash, KS, CKS + Nash-KS (experimental)

Want to know more?

V. Picheny, M. Binois, A. Habbal
A Bayesian optimization approach to find Nash equilibria (2017+)
preprint: https://arxiv.org/abs/1611.02440

M. Binois, V. Picheny, A. Habbal
The Kalai-Smorodinsky solution for many-objective Bayesian optimization (2017+)
NIPS BayesOpt workshop, https://bayesopt.github.io/papers/2017/28.pdf