GP REGRESSION WITH INEQUALITY CONSTRAINTS

ADAPTIVE STRATEGIES

—

Sébastien Da Veiga, Safran Tech
Joint work with Amandine Marrel (CEA)
ETICS Research School, Oct. 2020

SAFRAN

Introduction

Surrogate models are now commonly used for emulating complex computer codes

- UQ, GSA, optimization, …

Very often, computer codes simulate real physical phenomena, which usually have specific properties

- Symmetries
- Bound constraints (e.g. concentrations between 0 and 1, …)
- Monotonicity w.r.t. some input variables
- Solutions of PDEs (e.g. null Laplacian, divergence or curl free, …)

It is of great interest to incorporate such constraints in the proxy model

- Physics and expected behavior are respected (engineers like that !)
- Predictions and robustness may be improved

Introduction

Incorporation of bounds and monotonicity constraints have already been studied in nonparametric regression

- 1D setting
  - Ramsay 2005, Bigot and Gadat 2010
- Kernel regression
  - Dette and Scheder 2006
  - Constraints on weights: Hall and Huang 2001, Racine et al. 2009

Here, we focus on the GP regression framework

- Several old and very recent papers on the topic …

Introduction

The GP regression framework is very powerful when considering linear equality constraints

- Gaussianity + linear constraints make it possible to design adapted covariance functions (kernels)
  - This produces trajectories that intrinsically respect the constraints

- This « simple » remark gave rise to several interesting examples

- General theory recently studied (Ginsbourger et al. 2013)

Introduction

Sample paths of a GP with kernels designed for spatial symmetries

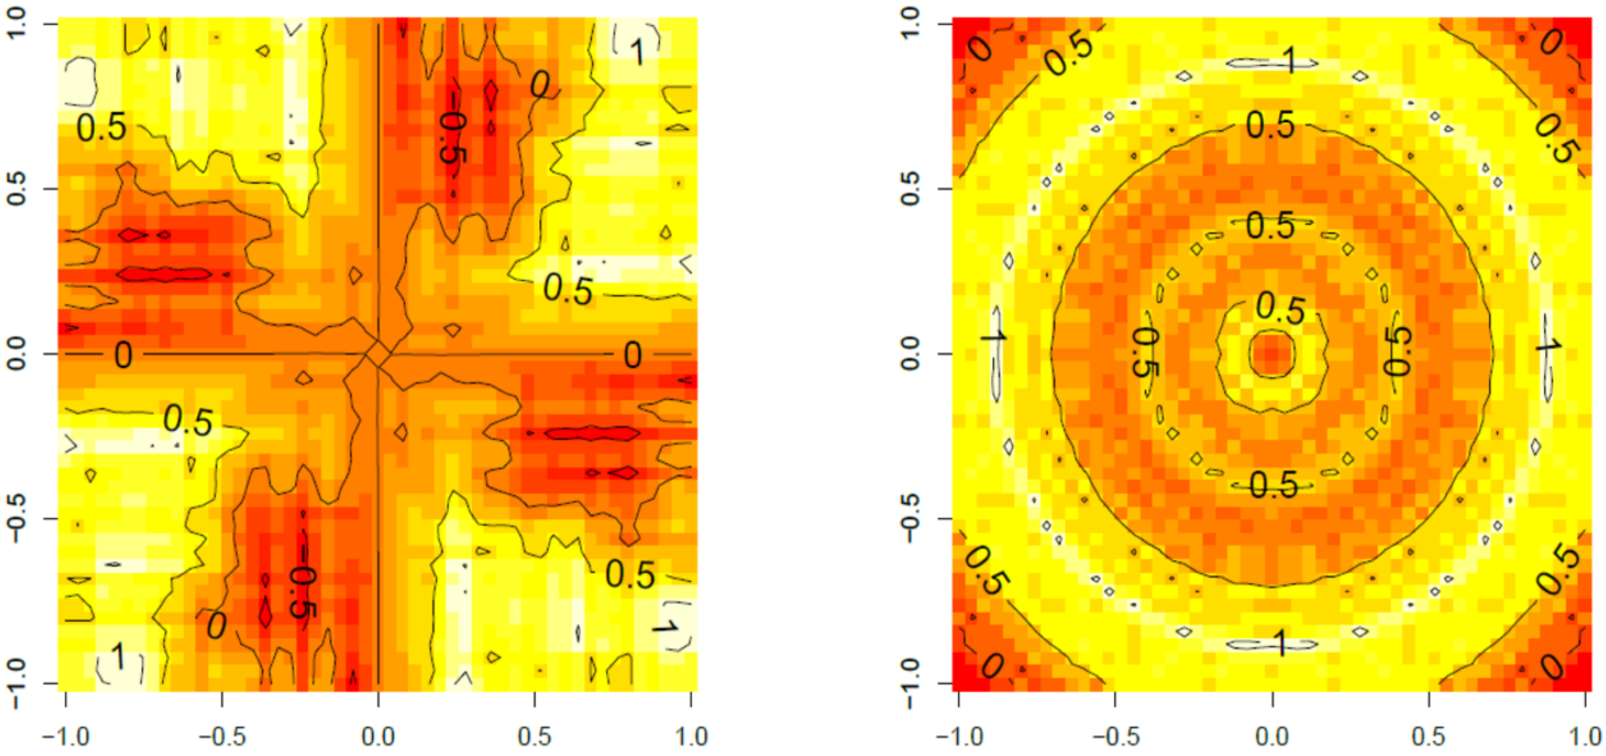

Ginsbourger et al. 2013

Introduction

Sample paths of a GP with kernels designed for specific constraints (null integral, solution of ODE and null Laplacian)

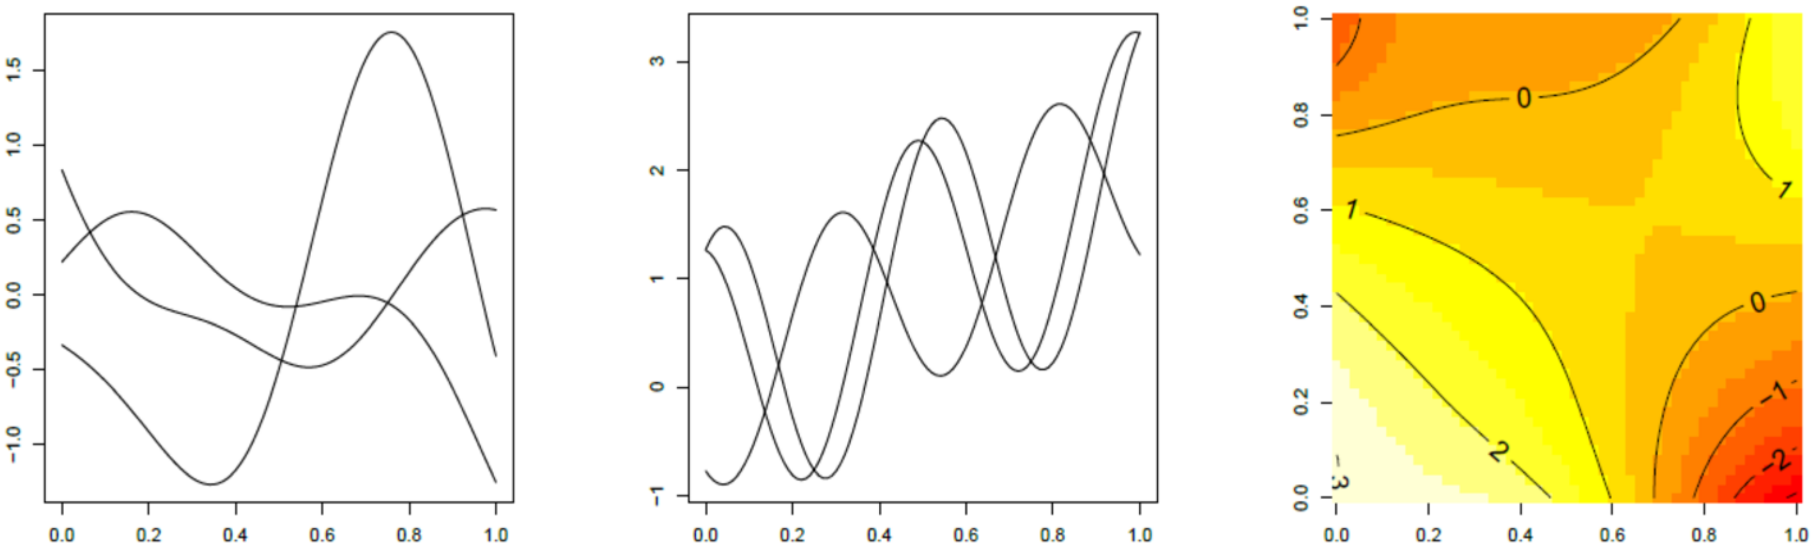

Ginsbourger et al. 2013

Introduction

Sample paths of a 2D-GP with kernels for curl-free and divergence free fields

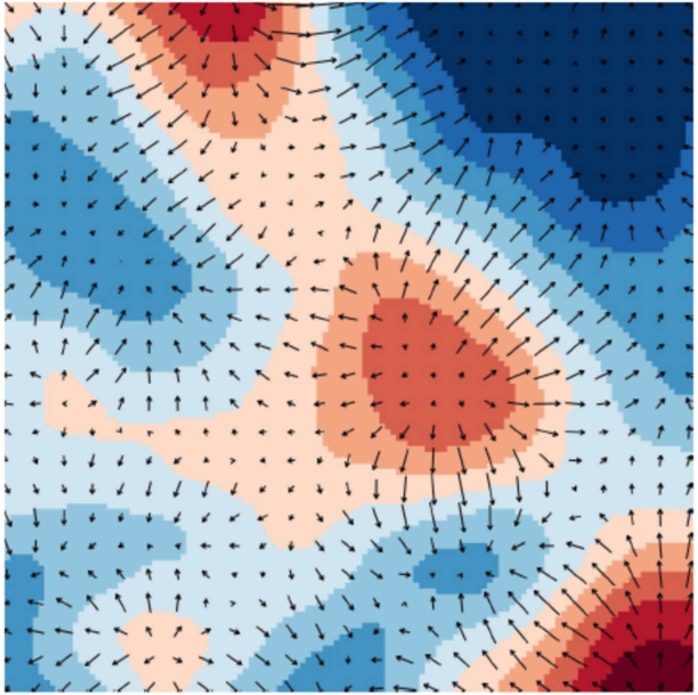

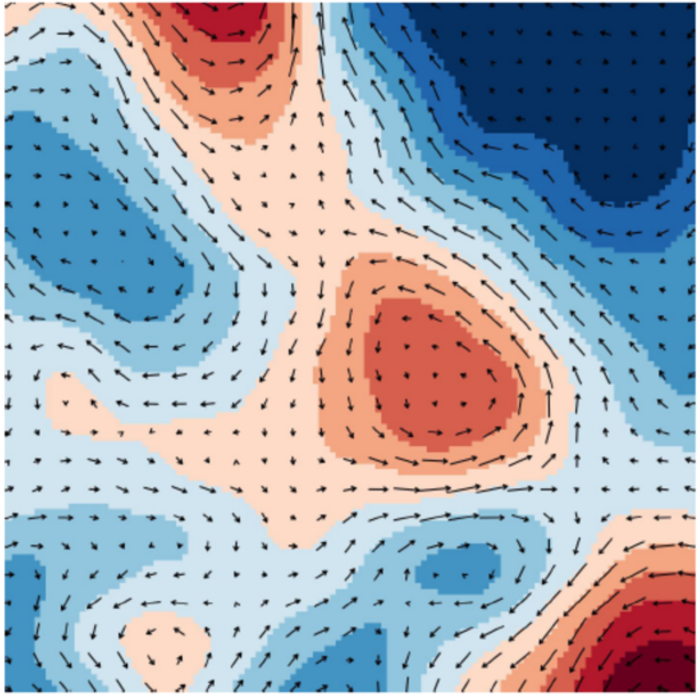

Scheuerer and Schlather 2012

Introduction

The GP regression framework is very powerful when considering linear equality constraints

However, inequality constraints cannot be handled so easily

- This includes bound and monotonicity constraints
- But also bounds on integrals or divergence/curl

Previous work on GP regression with inequality constraints

- Monotonicity
  - Data-augmentation: Abrahamsen and Benth 2001
  - Weights: Yoo and Kyriadis 2006
  - Sampling: Michalak 2008, Kleijnen and van Beers 2010
  - Constrained posterior distribution: Riihimaki and Vehtari 2010, Wang and Berger 2011
  - Expansion on a dedicated basis + constraints on weights: Mattouk & Bay 2017, Lopez-Lopera et al. 2018
- Any linear inequality constraints
  - Expectation of truncated normal distributions: Da Veiga and Marrel 2012, 2019
  - Sampling of truncated normal distributions: Agrell 2019

Standard GP regression

Notations

- Computer code $g : \mathbb{R}^D \to \mathbb{R}$
- Inputs $\mathbf{x} = (x^1, \ldots, x^D)$
- Output $y = g(\mathbf{x})$
- Observations $(\mathbf{x}_i, y_i)_{i=1,\ldots,n}$ $X_s = \left[\mathbf{x}_1^T, \ldots, \mathbf{x}_n^T\right]^T$ $Y_s = [y_1, \ldots, y_n]^T$

Model: Output seen as realization of stationary Gaussian process

$$Y(\mathbf{x}) = g_0(\mathbf{x}) + U(\mathbf{x})$$

$$g_0(\mathbf{x}) = \sum_{j=1}^J \beta_j g_j(\mathbf{x}) = G(\mathbf{x})\beta$$

$$C(\mathbf{x}, \mathbf{x}') = \sigma^2 R(\mathbf{x}, \mathbf{x}')$$

Conditioning on the observations

$$\tilde{Y}(\mathbf{x}^*) = [Y(\mathbf{x}^*) | Y(X_s) = Y_s]$$

Standard GP regression

Final predictor is a Gaussian with mean and covariance given by

$$\tilde{\mu}(\mathbf{x}^*) = \mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)\right)$$
$$= G(\mathbf{x}^*)\hat{\beta} + k(\mathbf{x}^*)^T \Sigma_s^{-1}\left(Y_s - G_s\hat{\beta}\right)$$

$$\tilde{C}(\mathbf{x},\mathbf{x}') = C(\mathbf{x},\mathbf{x}') - k(\mathbf{x})^T \Sigma_s^{-1} k(\mathbf{x}')$$

> MLE estimates of hyperparameters

$$\widehat{\psi} = \arg\min_{\psi} \widehat{\sigma^2}\, |R_{\psi,s}|^{\frac{1}{n}}$$

$$\widehat{\sigma^2} = \frac{1}{n}(Y_s - G_s\hat{\beta})^T R_{\psi,s}^{-1}(Y_s - G_s\hat{\beta})$$

$$\hat{\beta} = (G_s^T R_{\psi,S}^{-1} G_s)^{-1} G_s^T R_{\psi,s}^{-1} Y_s$$

GP regression with inequality constraints

To incorporate the constraints, we propose to keep the conditional expectation framework

- Predictions are equal to the expectation of the GP (conditioned at the observations) given that it respects the inequality constraints

For example, the corresponding predictor for bound or monotonicity constraints may be

$$\mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)|\forall \mathbf{x} \in I, a \leq \tilde{Y}(\mathbf{x}) \leq b\right) \qquad \mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)|\forall \mathbf{x} \in I, \frac{\partial \tilde{Y}}{\partial x^j}(\mathbf{x}) \geq 0\right)$$

- Note the link with with extrema of random fields …

$$\mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)|\min_{\mathbf{x}\in I} \tilde{Y}(\mathbf{x}) \geqslant a, \max_{\mathbf{x}\in I} \tilde{Y}(\mathbf{x}) \leqslant b\right)$$

- … but no tractable formula exists for joint distributions in the general case

GP regression with inequality constraints

We thus propose a discrete-location approximation:

$$\mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)|\forall i=1,\ldots,N, a \leq \tilde{Y}(\mathbf{x}_i) \leq b\right) \quad \mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)|\forall i=1,\ldots,N, \frac{\partial \tilde{Y}}{\partial x^j}(\mathbf{x}_i) \geq 0\right)$$

> Same approximation in Riihimaki and Vehtari 2010, Wang and Berger 2011

This generalizes easily to other constraints

$$\mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)|\ \forall k=1,\ldots,K,\ \forall i=1,\ldots,N_k,\ a_i^{(k)} \leq Z^{(k)}(\mathbf{x}_i^{(k)}) \leq b_i^{(k)}\right)$$

$$Z^{(k)} = \mathcal{L}^{(k)}\left[\tilde{Y}\right]$$

GP regression with inequality constraints

We thus propose a discrete-location approximation:

$$\mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)|\forall i = 1, \ldots, N, a \leq \tilde{Y}(\mathbf{x}_i) \leq b\right) \quad \mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)|\forall i = 1, \ldots, N, \frac{\partial \tilde{Y}}{\partial x^j}(\mathbf{x}_i) \geq 0\right)$$

> Same approximation in Riihimaki and Vehtari 2010, Wang and Berger 2011

This generalizes easily to other constraints

$$\mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)|\mathbf{a} \leq \mathbf{Z} \leq \mathbf{b}\right)$$

GP regression with inequality constraints

Standard framework:

- Take all trajectories which interpolate the observations
- Compute the average to get the kriging predictor
- (If desired, the variance yields a measure of accuracy)

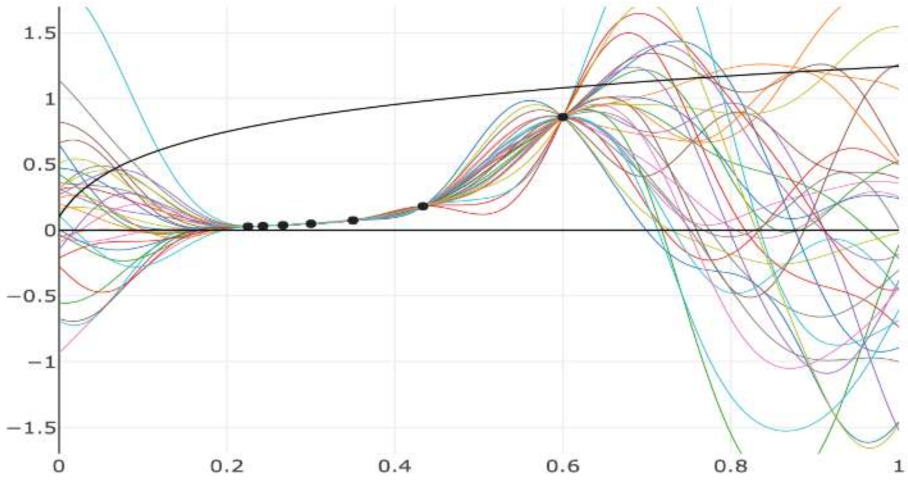

Agrell 2019

GP regression with inequality constraints

Standard framework:

- Take all trajectories which interpolate the observations
- Compute the average to get the kriging predictor
- (If desired, the variance yields a measure of accuracy)

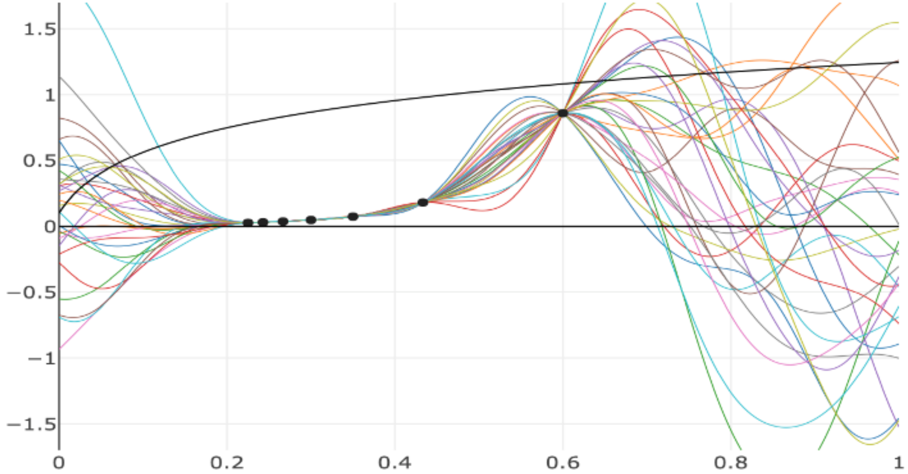

Here:

- Take all trajectories which interpolate the observations
- Select those which respect the constraints of bounds, monotonicity, …
- Compute the average to get the new kriging predictor
- (If desired, the variance yields a measure of accuracy)

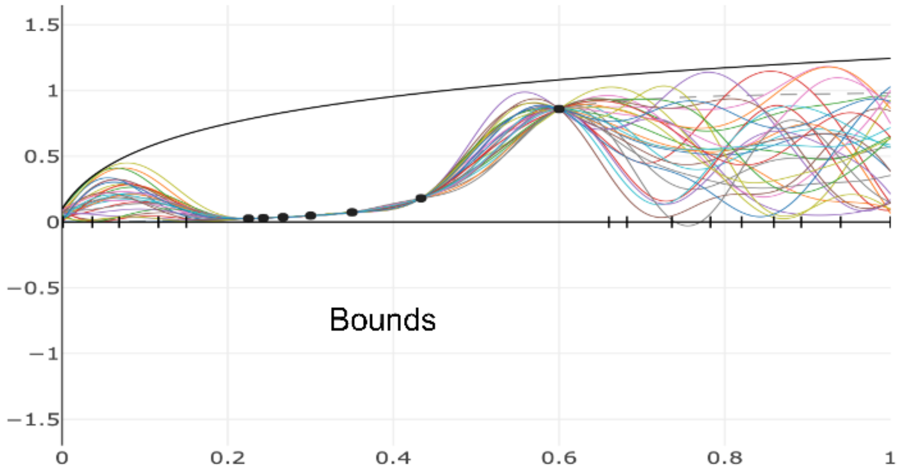

Agrell 2019

GP regression with inequality constraints

Standard framework:

- Take all trajectories which interpolate the observations
- Compute the average to get the kriging predictor
- (If desired, the variance yields a measure of accuracy)

Here:

- Take all trajectories which interpolate the observations
- Select those which respect the constraints of bounds, monotonicity, …
- Compute the average to get the new kriging predictor
- (If desired, the variance yields a measure of accuracy)

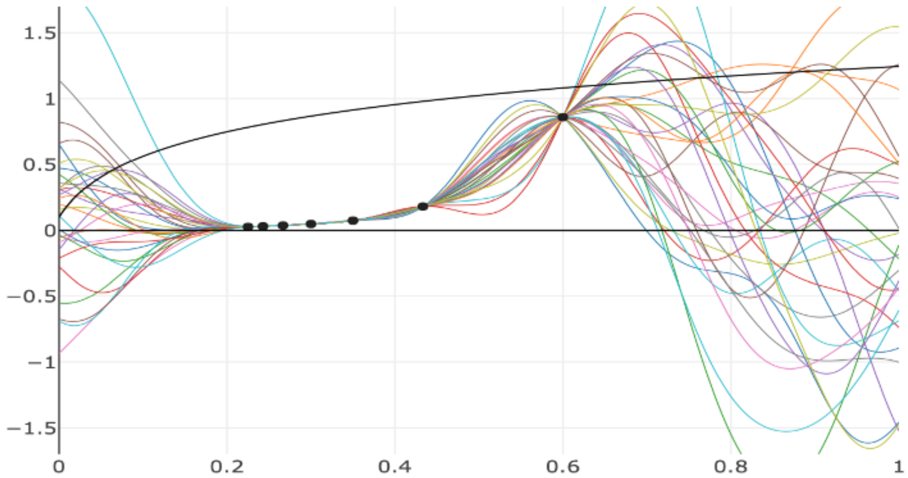

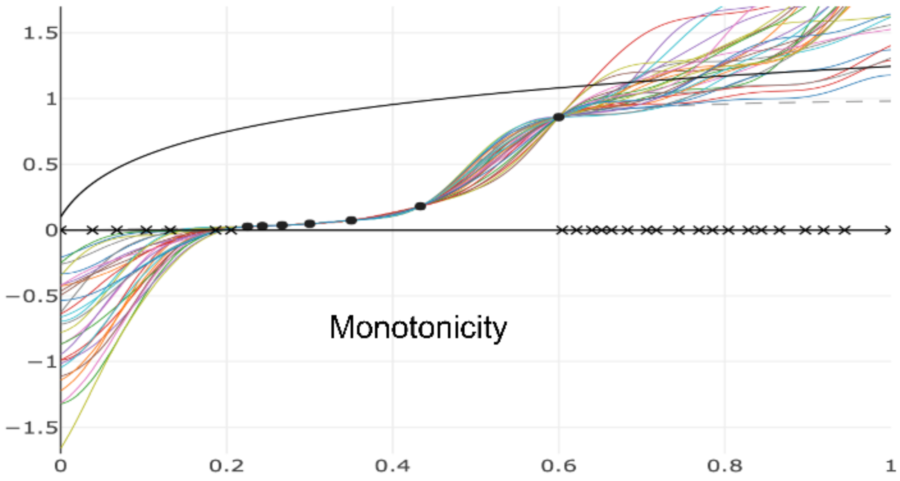

Agrell 2019

GP regression with inequality constraints

Standard framework:

- Take all trajectories which interpolate the observations
- Compute the average to get the kriging predictor
- (If desired, the variance yields a measure of accuracy)

Here:

- Take all trajectories which interpolate the observations
- Select those which respect the constraints of bounds, monotonicity, …
- Compute the average to get the new kriging predictor
- (If desired, the variance yields a measure of accuracy)

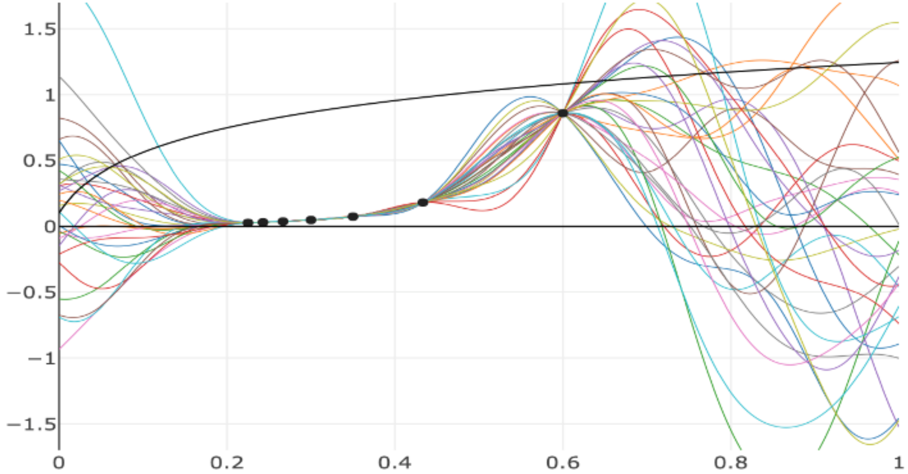

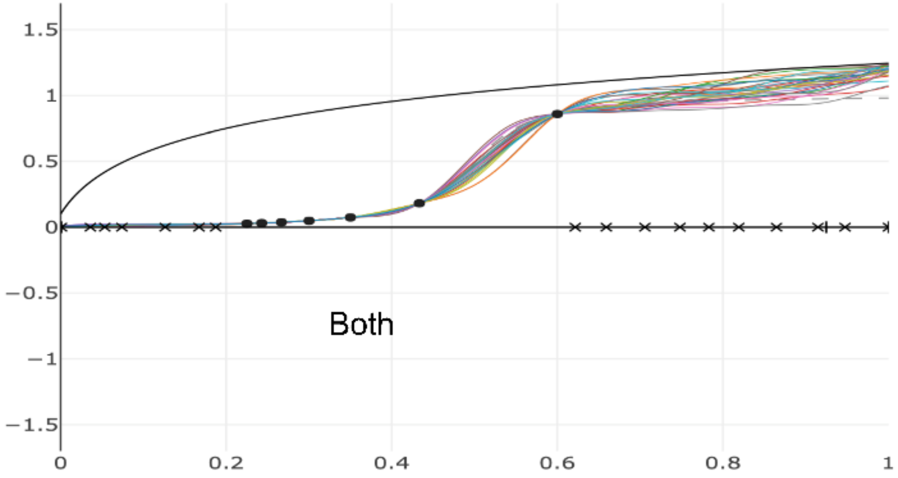

Agrell 2019

GP regression with inequality constraints

But how can we compute such expectations ?

This is where the linearity assumption comes into play

- Bounds, monotonicity, integral, divergence/curl constraints are linear w.r.t. the output
- The GP obtained by stacking the output and the quantities related to the constraints is then a GP too

$$\mathbf{W} = \left(\tilde{Y}(\mathbf{x}^*), \mathbf{Z}\right) \sim \mathcal{N}\left(\begin{bmatrix} \tilde{\mu}(\mathbf{x}^*) \\ \mu_{\mathbf{Z}} \end{bmatrix}, \begin{bmatrix} \tilde{C}(\mathbf{x}^*, \mathbf{x}^*) & \Sigma_{\tilde{Y}\mathbf{Z}} \\ \Sigma_{\tilde{Y}\mathbf{Z}}^T & \Sigma_{\mathbf{Z}} \end{bmatrix}\right)$$

Kotz et al. 2010

$$\mathbb{E}\left(\tilde{Y}(\mathbf{x}^*) | \mathbf{a} \leq \mathbf{Z} \leq \mathbf{b}\right) = \tilde{\mu}(\mathbf{x}^*) + \Sigma_{\tilde{Y}\mathbf{Z}} \Sigma_{\mathbf{Z}}^{-1} \left(\nu_{\mathbf{Z}} - \mu_{\mathbf{Z}}\right)$$

$$\mathrm{Var}\left(\tilde{Y}(\mathbf{x}^*) | \mathbf{a} \leq \mathbf{Z} \leq \mathbf{b}\right) = \tilde{C}(\mathbf{x}^*, \mathbf{x}^*) - \Sigma_{\tilde{Y}\mathbf{Z}} \left(\Sigma_{\mathbf{Z}}^{-1} - \Sigma_{\mathbf{Z}}^{-1} \Gamma_{\mathbf{Z}} \Sigma_{\mathbf{Z}}^{-1}\right) \Sigma_{\tilde{Y}\mathbf{Z}}^T$$

GP regression with inequality constraints

$$\mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)|\mathbf{a} \le \mathbf{Z} \le \mathbf{b}\right) = \tilde{\mu}(\mathbf{x}^*) + \Sigma_{\tilde{Y}\mathbf{Z}}\Sigma_{\mathbf{Z}}^{-1}\left(\nu_{\mathbf{Z}} - \mu_{\mathbf{Z}}\right)$$

$$\mathrm{Var}\left(\tilde{Y}(\mathbf{x}^*)|\mathbf{a} \le \mathbf{Z} \le \mathbf{b}\right) = \tilde{C}(\mathbf{x}^*, \mathbf{x}^*) - \Sigma_{\tilde{Y}\mathbf{Z}}\left(\Sigma_{\mathbf{Z}}^{-1} - \Sigma_{\mathbf{Z}}^{-1}\Gamma_{\mathbf{Z}}\Sigma_{\mathbf{Z}}^{-1}\right)\Sigma_{\tilde{Y}\mathbf{Z}}^T$$

$$\nu_{\mathbf{Z}} = \mathbb{E}(\mathbf{Z}|\mathbf{a} \le \mathbf{Z} \le \mathbf{b})$$

$$\Gamma_{\mathbf{Z}} = \mathrm{Var}(\mathbf{Z}|\mathbf{a} \le \mathbf{Z} \le \mathbf{b})$$

> The problem reduces to compute **moments of a multivariate normal vector subject to linear inequality constraints**

> Key object is then the **truncated normal distribution**

The truncated multivariate normal distribution

Start with a classical multivariate normal vector

$$\phi_{\mu_{\mathbf{Z}},\Sigma_{\mathbf{Z}}}(\mathbf{z}) = \frac{1}{(2\pi)^{N/2}|\Sigma_{\mathbf{Z}}|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{z}-\mu_{\mathbf{Z}})^T\Sigma_{\mathbf{Z}}^{-1}(\mathbf{z}-\mu_{\mathbf{Z}})\right), \mathbf{z} \in \mathbb{R}^N$$

Its truncated version is defined via the following pdf

$$\phi_{\mu,\Sigma,\mathbf{a},\mathbf{b}}(\mathbf{z}) = \begin{cases} \frac{\phi_{\mu,\Sigma}(\mathbf{z})}{\mathbb{P}(\mathbf{a}\leq\mathbf{Z}\leq\mathbf{b})} & \text{for } \mathbf{a} \leq \mathbf{z} \leq \mathbf{b}, \\ 0 & \text{otherwise,} \end{cases}$$

- Expectation writes

$$\mathbb{E}(Z_i|\mathbf{a} \leq \mathbf{Z} \leq \mathbf{b}) = \mu_i + \sum_{j=1}^{N} \sigma_{ij}\left(f_j(a_j) - f_j(b_j)\right)$$

$$f_i(z) = \int_{a_1}^{b_1} \cdots \int_{a_{i-1}}^{b_{i-1}} \int_{a_{i+1}}^{b_{i+1}} \cdots \int_{a_N}^{b_N} \phi_{\mu,\Sigma,\mathbf{a},\mathbf{b}}(z_1,\ldots,z_{i-1},z,z_{i+1},\ldots,z_N)d\mathbf{z}_{-i}$$

- Other formulas for the covariance, linear and elliptical constraints available since the 60's (Tallis 61, Tallis 63, Tallis 65)

The truncated multivariate normal distribution

The expectation and variance are our goal here

Available formulas involve Gaussian integrals with dimensionality equal to the number of points where we impose the constraints

We thus need efficient approximations when this number is large (as it should be !)

- Genz numerical approximation of Gaussian integrals (Genz 1992)
- Sampling from a truncated Gaussian
- Correlation-free formula (« crude » covariance tapering)

$$
\begin{aligned}
\mathbb{E}(Z_i|\mathbf{a} \leq \mathbf{Z} \leq \mathbf{b}) &\approx \mathbb{E}(Z_i|a_i \leq Z_i \leq b_i) \\
&\approx \mu_i + \sigma_i \frac{\phi(\tilde{a}_i) - \phi(\tilde{b}_i)}{\Phi(\tilde{b}_i) - \Phi(\tilde{a}_i)} \\
\mathrm{Var}(Z_i|\mathbf{a} \leq \mathbf{Z} \leq \mathbf{b}) &\approx \mathrm{Var}(Z_i|a_i \leq Z_i \leq b_i) \\
&\approx \sigma_i^2 \left[1 + \frac{\tilde{a}_i \phi(\tilde{a}_i) - \tilde{b}_i \phi(\tilde{b}_i)}{\Phi(\tilde{b}_i) - \Phi(\tilde{a}_i)} - \left(\frac{\phi(\tilde{a}_i) - \phi(\tilde{b}_i)}{\Phi(\tilde{b}_i) - \Phi(\tilde{a}_i)} \right)^2 \right]
\end{aligned}
$$

GP regression with inequality constraints – Algorithm in practice

Algorithm 1: Constrained GP prediction at $\mathbf{x}^*$

Given X_s, Y_s, $f_0(\mathbf{x})$, $R_\psi(\mathbf{x}-\mathbf{x}')$, $\mathcal{L}^{(1)},\dots,\mathcal{L}^{(K)}$, $\mathbf{X}=\{X^{(1)},\dots,X^{(k)}\}$, $\mathbf{a}$, $\mathbf{b}$

(1) Estimate the GP parameters β, σ and ψ.

(2) Distribution of $\mathbf{W}=\left(\tilde{Y}(\mathbf{x}^*),\mathbf{Z}\right)$ (Equation (14))

(a) Build the Gaussian vector $\left(\tilde{Y}(\mathbf{x}^*),\mathcal{L}^{(1)}\left[Y\right]\left(X^{(1)}\right),\dots,\mathcal{L}^{(K)}\left[Y\right]\left(X^{(K)}\right),Y_s\right)$

(b) By conditioning w.r.t. Y_s, compute $\tilde{\mu}(\mathbf{x}^*)$ and $\tilde{C}(\mathbf{x}^*,\mathbf{x}^*)$, $\mu_\mathbf{Z}$, $\Sigma_\mathbf{Z}$ and $\Sigma_{\tilde{Y}\mathbf{Z}}$ (formulas for conditional moments of Gaussian vector, Equations (3) and (4)).

(3) Compute the truncated moments $\nu_\mathbf{Z}$ and $\Gamma_\mathbf{Z}$ with Tallis formula and Genz's approximation.

(4) Build the final constrained predictor and variance by computing the truncated moments $\mathbb{E}\left(\tilde{Y}(\mathbf{x}^*)|\mathbf{a}\leq\mathbf{Z}\leq\mathbf{b}\right)$ and $\mathrm{Var}\left(\tilde{Y}(\mathbf{x}^*)|\mathbf{a}\leq\mathbf{Z}\leq\mathbf{b}\right)$ with equations (15) and (16).

GP regression with inequality constraints – Simple examples

Incorporation of monotonicity on 100 equally-spaced constraint points

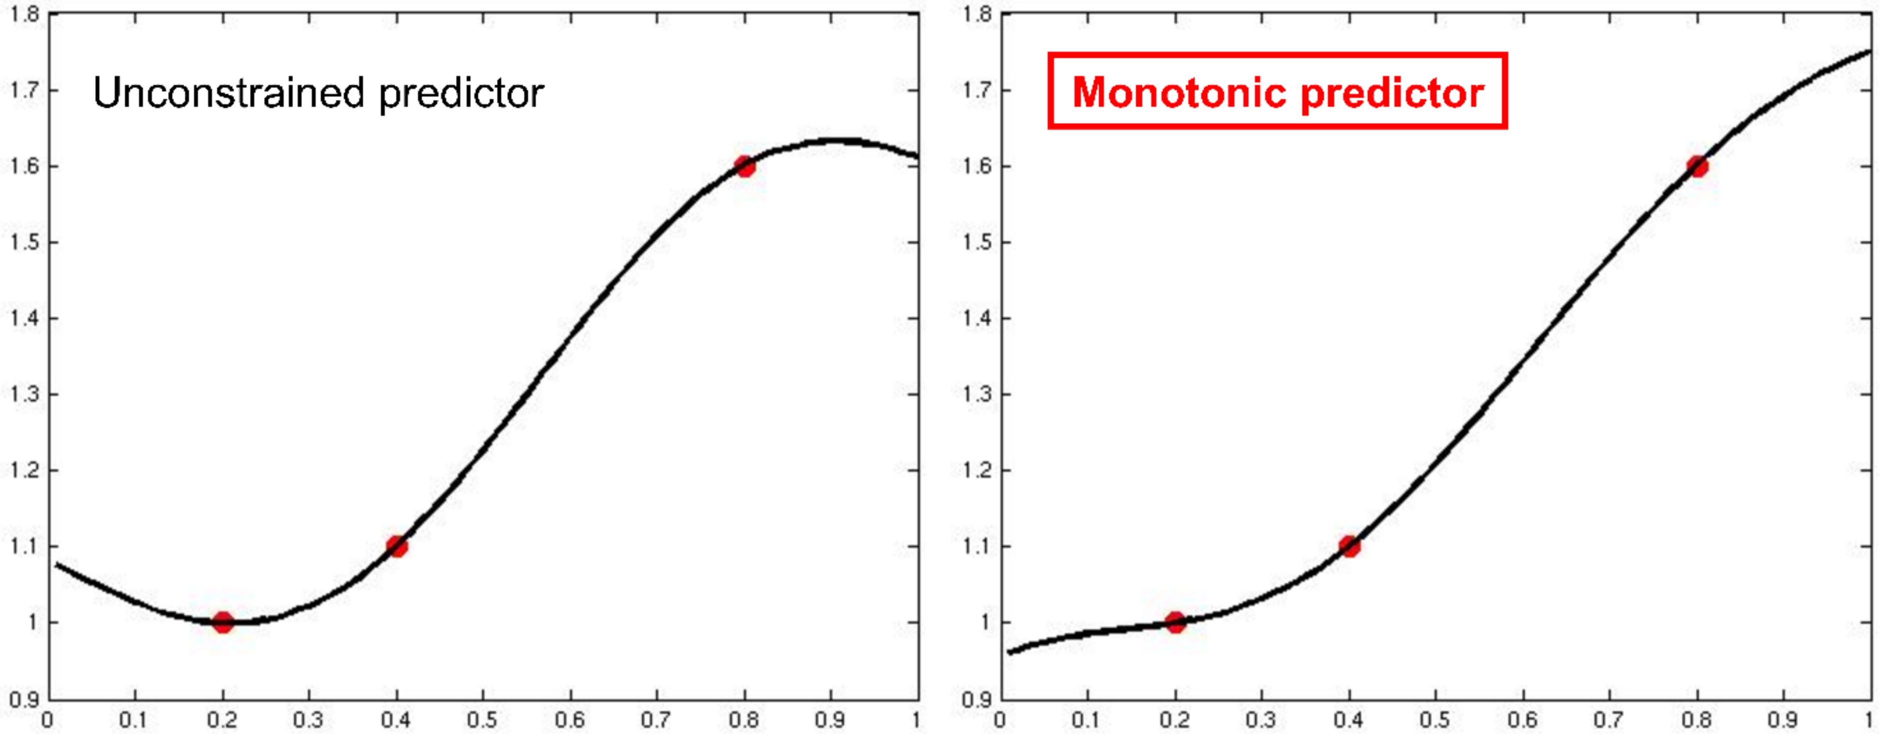

GP regression with inequality constraints – Simple examples

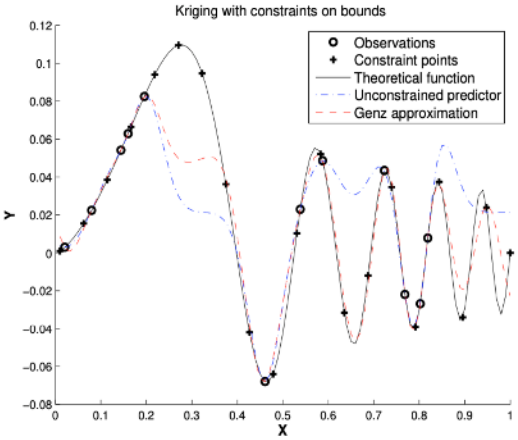

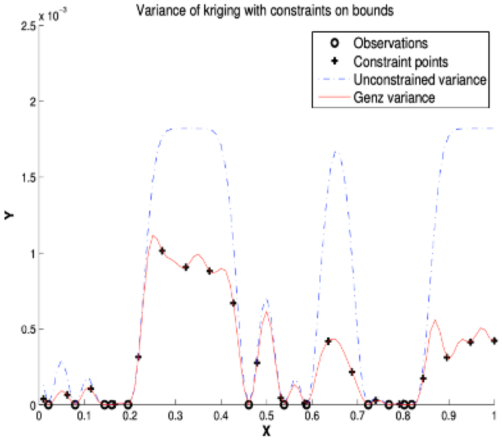

GP regression with inequality constraints – Simple examples

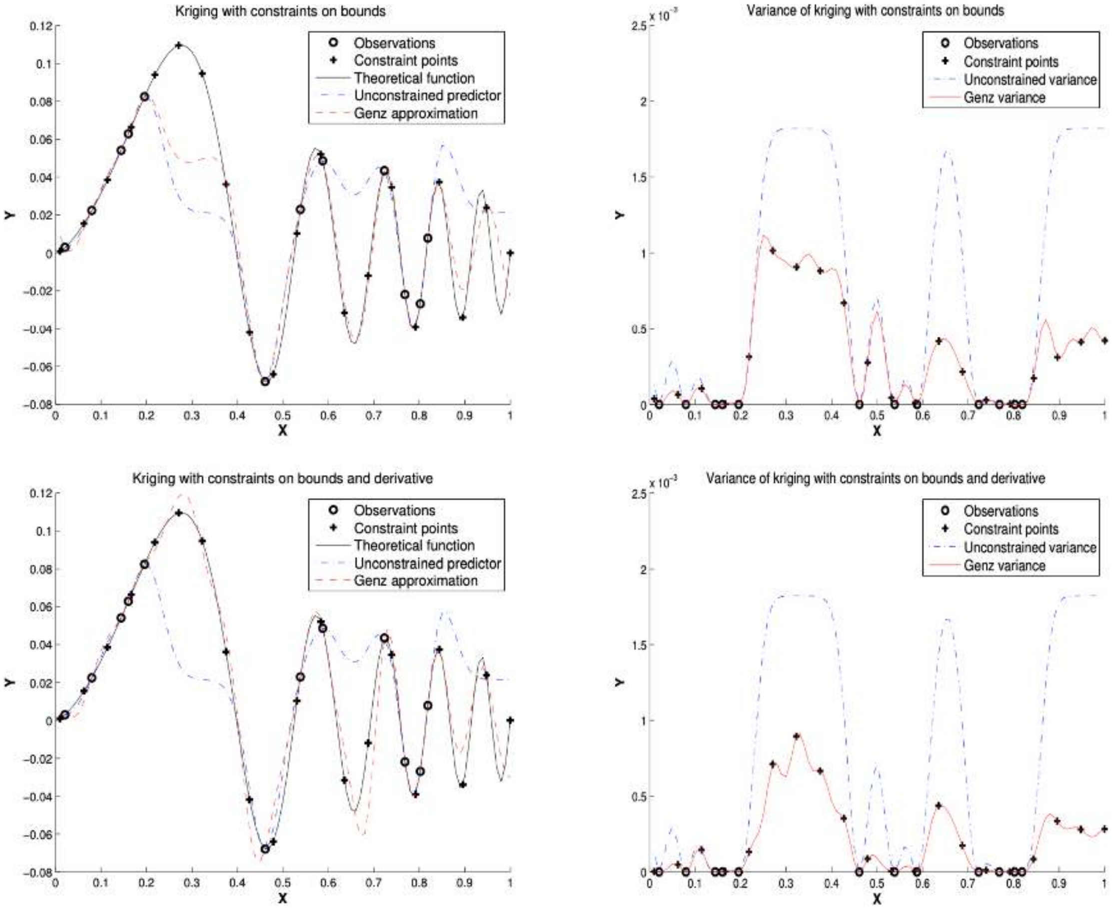

SAFRAN

GP regression with inequality constraints – Things « under the carpet »

Hyperparameters estimation

- Here we estimate it on the initial GP, **so we don't account for the constraints**
- This leads to potential unexpected behaviours for the final constrained predictor
  - In particular we may need to add more locations for constraints than usual
- Recent works on MLE with constraints are promising

Efficient generalization to higher dimensional problems is not so easy

- From a theoretical perspective, no change in the formulas
- However, « spanning » the subset where we impose constraints will necessitate much more constraint points in the discrete-location approximation
  - Genz numerical integration and sampling cannot be used with tens of thousands of constraints
- The idea is to use the correlation induced among the constraint points (and with the observations)
  - **It is not necessary to place constraint points where the predictor has a high probability to respect the constraints (e.g. close to another constraint point, or where the prediction variance is very low)**

Adaptive strategy for the constraint locations

This motivates the design of an adaptive strategy for choosing the constraints locations

The key here is to compute the probability that the constrained predictor does not respect the constraint at any point

- **Obviously this is not a nice and friendly normal distribution as in standard GP regression**
- It involves the CDF of a truncated normal distribution, two ways to handle it:
  - Use a truncated normal sampling algorithms (Agrell 2019)
  - Make a crude but fast normal approximation (Da Veiga & Marrel 2019)

Constraint points are thus added one at a time, at locations where this probability is the highest

SAFRAN

Adaptive strategy for the constraint locations

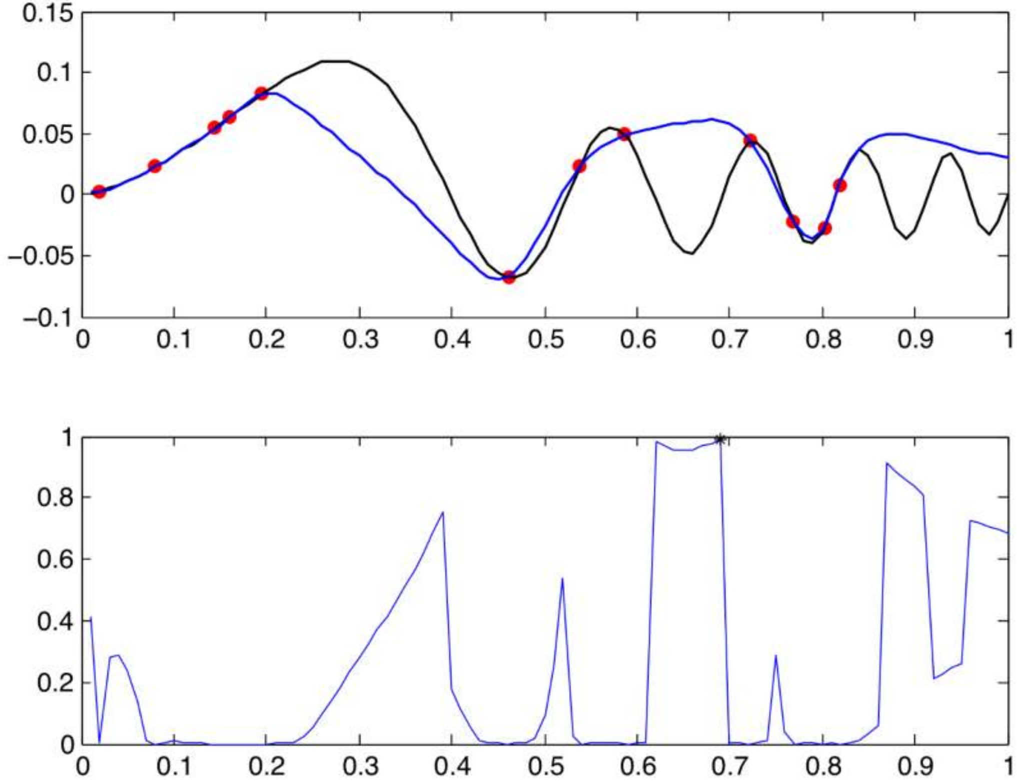

SAFRAN

Adaptive strategy for the constraint locations

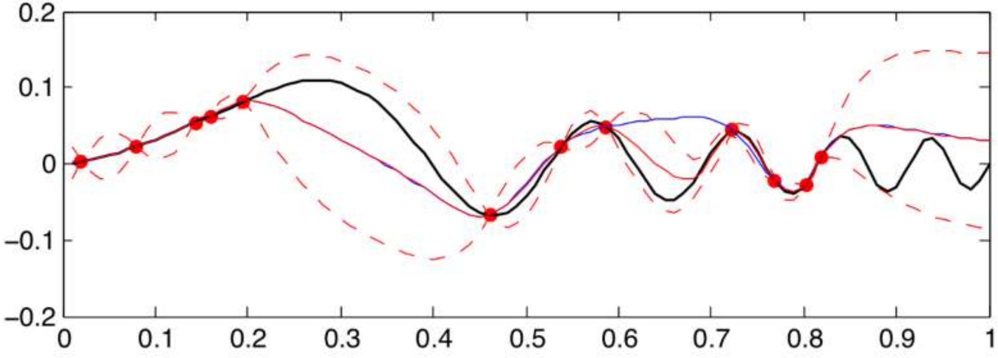

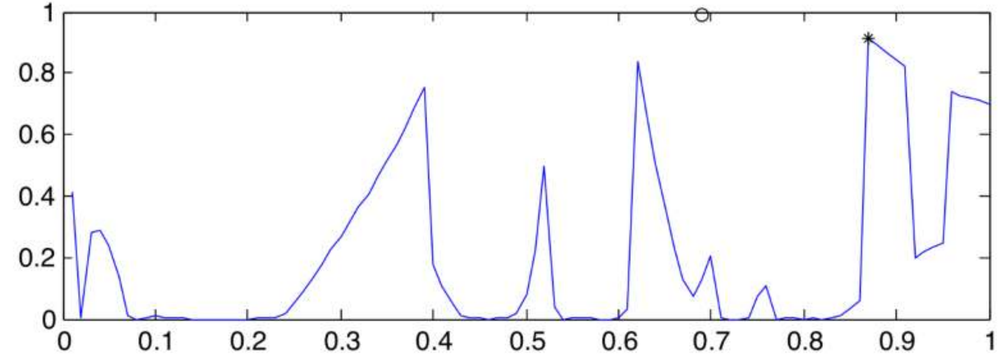

Adaptive strategy for the constraint locations

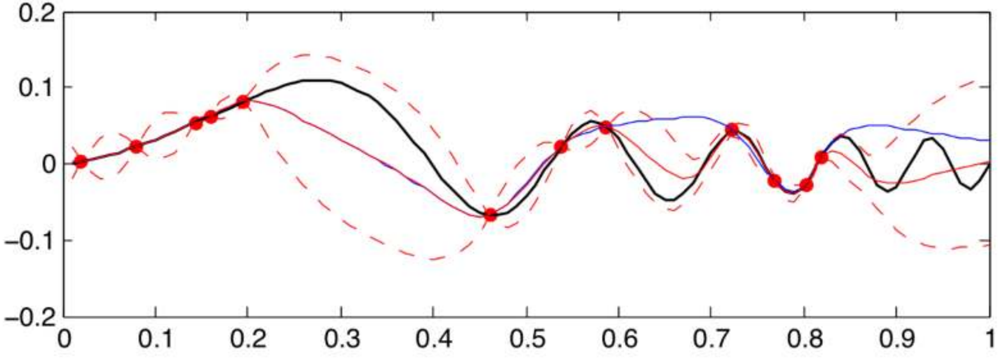

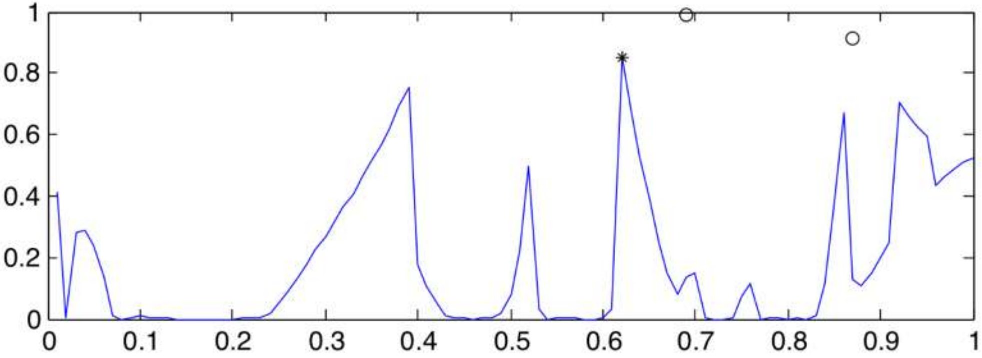

Adaptive strategy for the constraint locations

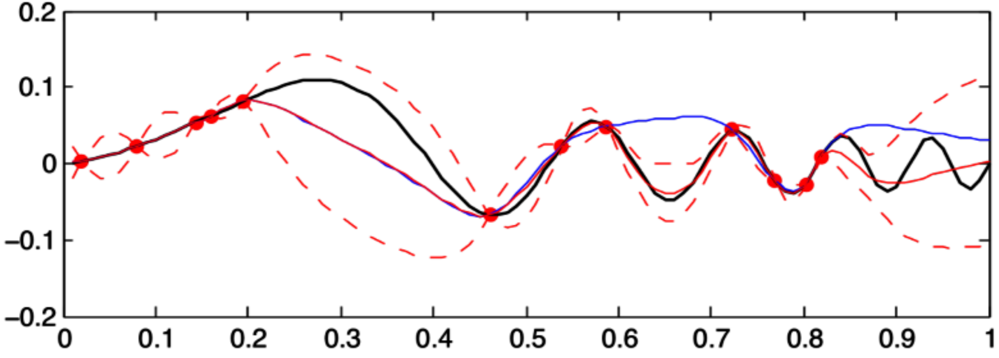

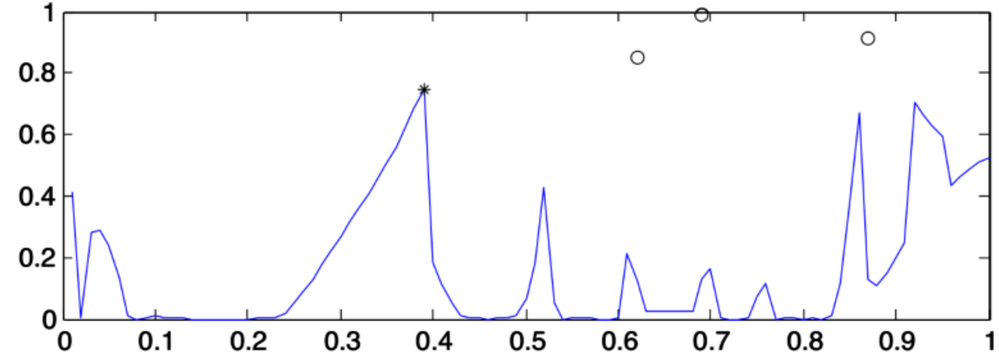

Adaptive strategy for the constraint locations

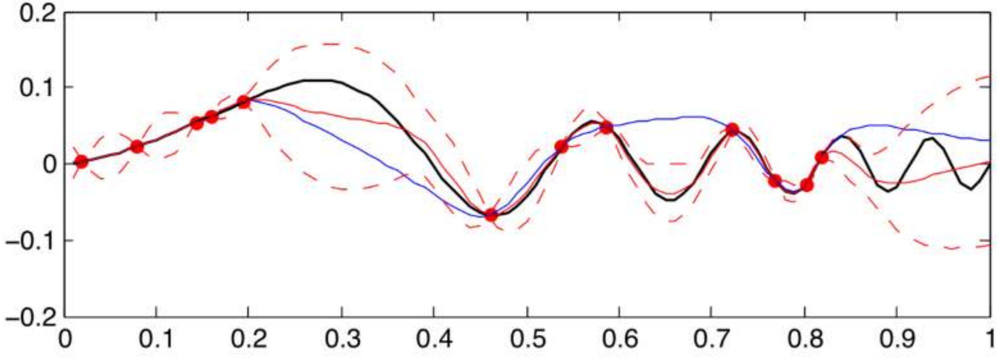

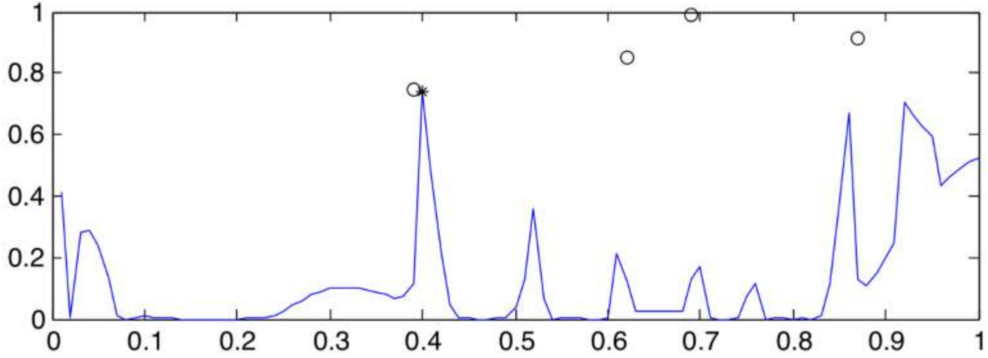

SAFRAN

Adaptive strategy for the constraint locations

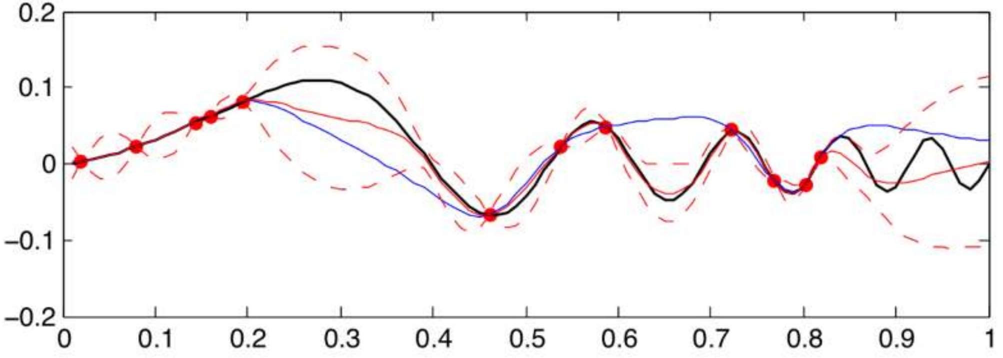

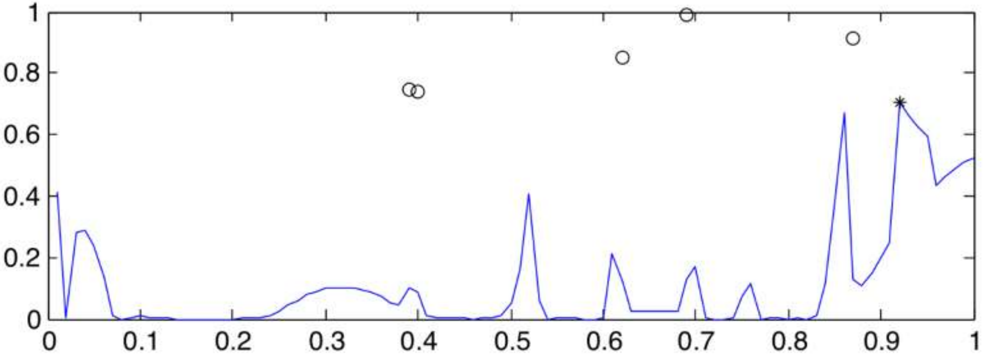

SAFRAN

Adaptive strategy for the constraint locations

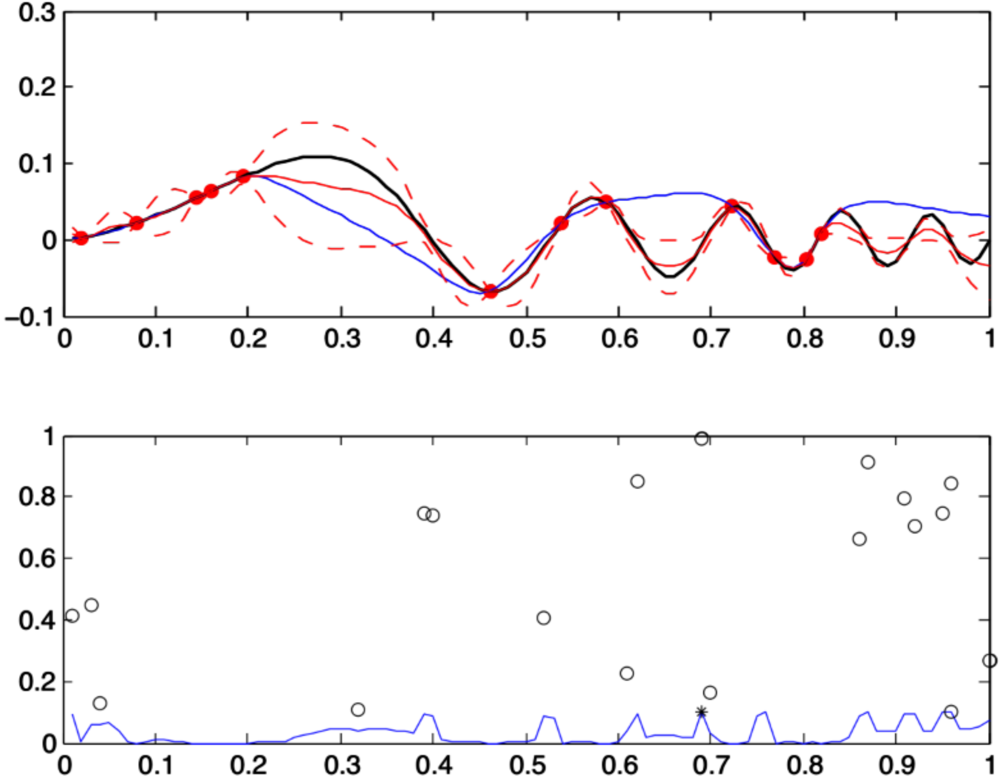

Adaptive strategy for the constraint locations

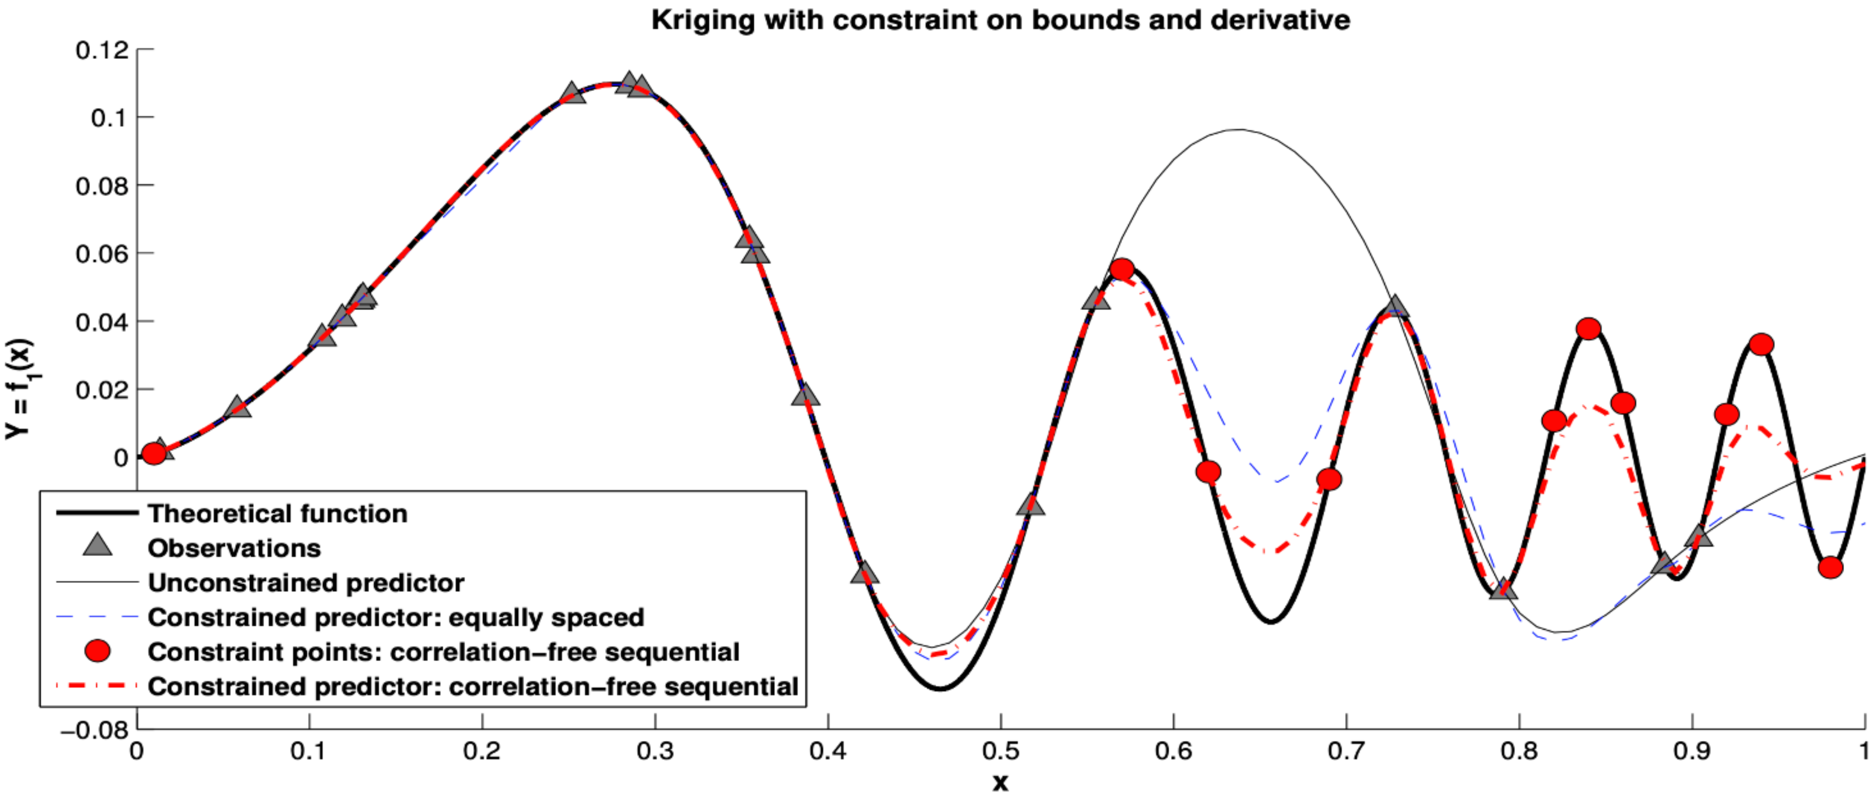

Adaptive strategy for the constraint locations

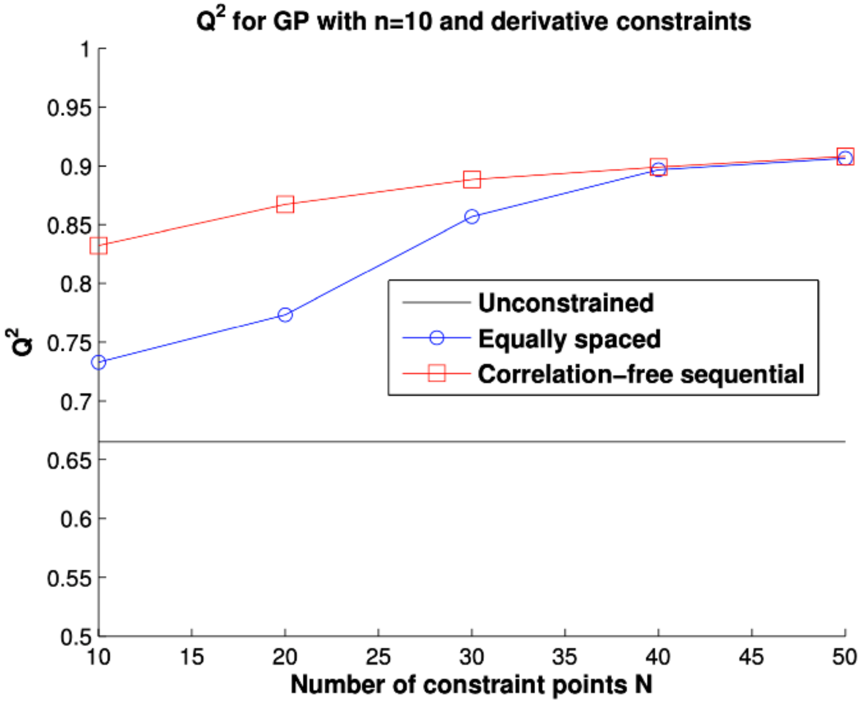

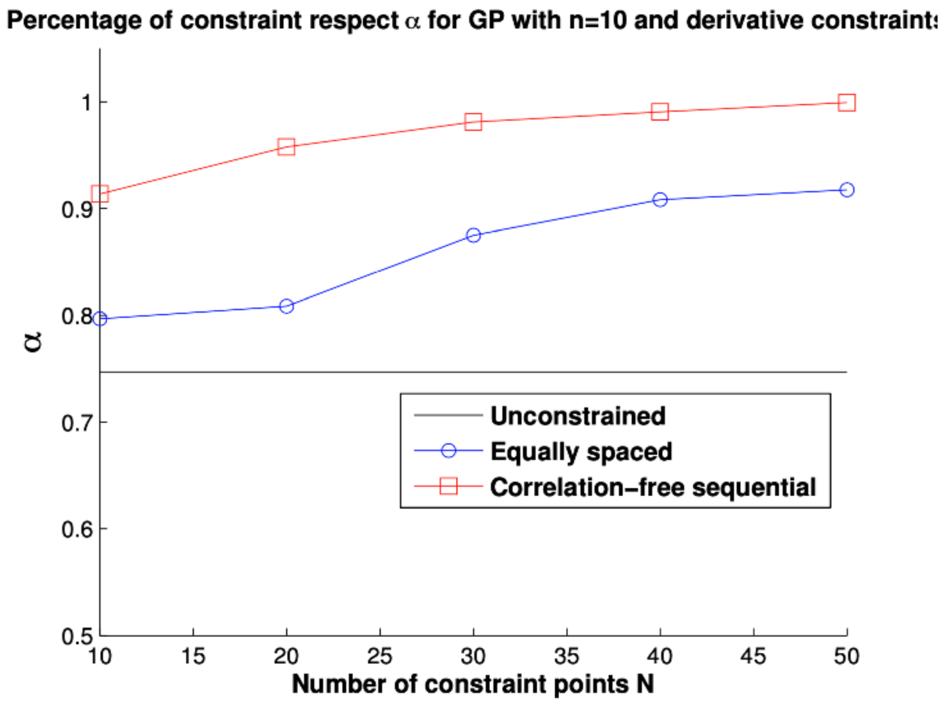

Adaptive strategy for the constraint locations

You can also use a cruder but even faster approximation

- Instead of using Genz n times, find the constraint locations with the correlation-free formula (no cost)
- Once the locations are found, the final prediction is performed with Genz
- **Results seem to indicate that we have almost no lost of prediction accuracy**

Algorithm has been tested on functions depending on up to 20 input variables

Adaptive strategy for the constraint locations

	Virtual obs. likelihood	Inference strategy	Strategy for finding X^v
Agrell (2019)	Indicator + noise	Sampling (Minimax tilting)	Based on estimating $p_c(\mathbf{x})$ from samples
Wang and Berger (2016)	Indicator	Sampling (Gibbs)	Based on estimating $p_c(\mathbf{x})$ from samples
Da Veiga and Marrel (2012, 2015)	Indicator	Moment approximation (Genz)	Based on approximating $p_c(\mathbf{x})$ assuming Gaussian posterior distribution
Riihimki and Vehtari (2010)	Probit	Expectaion Propagation	Based on approximating $p_c(\mathbf{x})$ assuming Gaussian posterior distribution
Golchi et al. (2015)	Probit	SCMC	NA

Table 1: Summary of alternative approaches that make use of virtual observations. The table compares the likelihood used for virtual observations, the method used for inference and to determine the set of virtual observation locations X^v.

Agrell 2019

Conclusion & perspectives

Theoretical framework to incorporate any linear inequality constraints in GP regression

- Truncated normal distribution is central
- Approximation formulas for moments or efficient samplers

From a practical point of view, high-dimensional problems can be accommodated with an adaptive strategy

- Even in low-dimensional examples, it is more efficient to choose the constraint locations sequentially
- The correlation-free trick heavily accelerates the search
- Recent results with an efficient sampler are also promising

Challenges

- **Hyperparameter estimation**
- Advanced computational tools will certainly be necessary - Machine learning methods may be of great help, with adaptation
  - Incomplete Choleshy decomposition (Bach and Jordan 2002)
  - Random Fourier Features (Rahimi and Recht 2007, 2008) & alike

Conclusion & perspectives

Surprinsingly area still very active

- Quite complete survey
  Swiler et al. 2020

A Survey of Constrained Gaussian Process Regression: Approaches and Implementation Challenges

L. Swiler[1], M. Gulian[1], A. Frankel[2], C. Safta[2], and J. Jakeman[1]

[1]Center for Computing Research, Sandia National Laboratories, Albuquerque, NM, 87123
[2]Computational Science and Analysis, Sandia National Laboratories, Livermore, CA, 94550
{lpswile,mgulian,afrank,csafta,jdjakem}@sandia.gov

- New approximation for moments
  Gessner et al. 2020

Integrals over Gaussians under Linear Domain Constraints

Alexandra Gessner
University of Tuebingen and MPI for Intelligent Systems
Tübingen, Germany
agessner@tue.mpg.de

Oindrila Kanjilal
University of Tuebingen and Technical University of Munich
Germany
oindrila.kanjilal@tum.de

Philipp Hennig
University of Tuebingen and MPI for Intelligent Systems
Tübingen, Germany
ph@tue.mpg.de

- Constrained MLE
  Pensoneault et al. 2020

Nonnegativity-Enforced Gaussian Process Regression

Andrew Pensoneault[a], Xiu Yang[b,*], Xueyu Zhu[a,*]

[a]Department of Mathematics, University of Iowa, Iowa City, IA 52246, USA
[b]Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, PA 18015, USA

- RKHS point of view
  Aubin-Frankowski & Szabo 2020

Hard Shape-Constrained Kernel Machines

Pierre-Cyril Aubin-Frankowski
École des Ponts ParisTech and CAS
MINES ParisTech, PSL
France
pierre-cyril.aubin@mines-paristech.fr

Zoltán Szabó
Center of Applied Mathematics, CNRS
École Polytechnique, Institut Polytechnique de Paris
Route de Saclay, Palaiseau, 91128, France
zoltan.szabo@polytechnique.edu

So stay tuned for new advances !