Guaranteed bounds for stochastic dynamic optimization problems

Georg Ch. Pflug, with Martin Glanzer and Sebastian Maier
(work in progress)

Workshop "Robustness and Resilience in Stochastic Optimization and Statistical Learning: Mathematical Foundations", Erice, May 2022.

The finite sample guarantee problem

Suppose that P is a probability measure and $\gamma(P)$ is some parameter. In Statistics and also in Stochastic Optimization we replace the difficult problem [1] to find $\gamma(P)$ by using a sample. Let $\hat{P}_n$ the empirical distribution and let $\gamma(\hat{P}_n)$ be our empirical estimate. In order to give the quality of this estimate, we use *confidence sets* $\hat{C}_n$ around $\gamma(\hat{P}_n)$ such that $P\{\gamma(P) \notin \hat{C}_n\}$ is small. Given an error level β, confidence sets typically hold this level only asymptotically

$$\lim_{n\to\infty} P\{\gamma(P) \notin \hat{C}_n\} \leq \beta.$$

but noting is said for finite and fixed n. If

$$P\{\gamma(P) \notin \hat{C}_n\} \leq \beta$$

we call this a confidence set with *finite sample guarantee*.

[1] This may be difficult per se (as in stochastic optimization) or P is unavailable (as in statistics)

The basic problem: Stochastic Optimization

$$\max\{\mathbb{E}_P(Q(x,\xi) : x \in \mathbb{X}, x \lhd \mathfrak{F}\}$$

where
$\mathfrak{F} = (\mathcal{F}_0, \mathcal{F}_1, \ldots, \mathcal{F}_T)$ is a filtration on some probability space,
ξ is an $\mathfrak{F}$-adapted stochastic scenario process,
P is the probability law of the scenario process,
x is the $\mathfrak{F}$-adapted stochastic decision process.
We solve such a problem by a sampling approach and are interested in finite sample guaranteed bounds for the optimal value.
Notice that each feasible decision process x gives a lower bound.
There are in principle 3 ways to find good upper bounds

- ambiguity extension
- duality approach
- guaranteed over-estimation

Ambiguity extension

Let dl be the nested distance for adapted stochastic processes introduced by Pflug and Pichler (2012).
If $\hat{P}_n$ is a sampled approximation of the underlying scenario process and if we find a bound ϵ such that

$$P\{\mathrm{dl}(\hat{P}_n, P) \leq \epsilon\} \geq 1 - \beta$$

(n has to be large enough) then solving the extended problem

$$\max_x \max_{dl(\hat{P}_n, P) \leq \epsilon} \{\mathbb{E}_P(Q(x, \xi) : x \in \mathbb{X}, x \lhd \mathfrak{F}\}$$

gives a guaranteed upper bound with error level β.

Duality approach

We illustrate this approach for the optmal stopping problem, where the decision is binary (to stop or not to stop). The goal is to find an $\mathfrak{F}$-adapted stopping time τ such that for a stochastic reward process $(C_0, C_1, \ldots, C_T)$ the expectation $\mathbb{E}(C_\tau)$ of the process stopped at τ process is maximal:

$$V_0 := \sup\{\mathbb{E}(C_\tau) : \tau \lhd \mathfrak{F}\}$$

Notice that if $\bar{\mathfrak{F}} = (\mathcal{F}_T, \ldots, \mathcal{F}_T)$ is the clairvoyant filtration, then

$$\sup\{\mathbb{E}(C_\tau) : \tau \lhd \bar{\mathfrak{F}}\} = \mathbb{E}(\sup_t C_t).$$

Here is what is well known:
Let $\mathcal{M}_{\mathfrak{F}}$ be the family of all $\mathfrak{F}$-martingales. By the optional sampling theorem for $(M_t) \in \mathcal{M}_{\mathfrak{F}}$ and all $\mathfrak{F}$ stopping times τ

$$\mathbb{E}[C_\tau] = M_0 + \mathbb{E}[C_\tau - M_\tau].$$

The duality Theorem

Theorem.

$$V_0 = \inf\{M_0 + \mathbb{E}[\sup_t(C_t - M_t)] : (M_t) \in \mathcal{M}_{\mathfrak{F}}\}.$$

Thus any $(M_t) \in \mathcal{M}_{\mathfrak{F}}$ leads to an upper bound. Moreover, the optimal dual martingale process M_t^* can be explititly given: Let (V_t) be the smallest supermartingale dominating (C_t) (this is called *Snell's envelope*). It is defined in a backward recursive way:

$$\begin{aligned} V_T &= C_T, \\ V_t &= \max(C_t, \mathbb{E}[V_{t+1}|\mathcal{F}_t]). \end{aligned}$$

Then the (typically unavailable) optimal dual process is

$$M_t^* = V_0 + \sum_{s=1}^{t}(V_s - \mathbb{E}(V_s|\mathcal{F}_{s-1})).$$

Information relaxation

If $\mathfrak{G} = (\mathcal{G}_0, \mathcal{G}_1, \ldots, \mathcal{G}_T)$ is a relaxation of the filtration $\mathfrak{F}$, i.e. $\mathcal{G}_t \supseteq \mathcal{F}_t$, then trivially

$$\sup\{\mathbb{E}(C_\tau) : \tau \lhd \mathfrak{F}\} \leq \sup\{\mathbb{E}(C_\tau) : \tau \lhd \mathfrak{G}\} \tag{1}$$

If M is any $\mathfrak{F}$ martingale, then $\sup\{M_0 + \mathbb{E}(C_\tau - M_\tau) : \tau \lhd \mathfrak{G}\}$ is a valid upper bound and $M_0 + \mathbb{E}(\sup_t(C_t - M_t))$ is a typically worse upper bound. Only if M^* is the optimal dual martingale process, then all of the above inequalities become eqalities and the upper bounds coincide, meaning that it is allowed to take even the clairvoyant filtration $\bar{\mathfrak{F}}$ as information structure without increasing the bound.
We call any $\mathfrak{G}$- martingale $(M_t - M_0)$ a *penalty function* for the information relaxation.

Pricing of American Options as a stopping problem

We assume now that the filtration is generated by a Markovian process $S_0, S_1, \ldots, S_T$ such that the filtration $\mathfrak{F}$ is generated by this process: $\mathcal{F}_t = \sigma(S_t)$. The payment process is $C_0 = c_0(S_0), C_1 = c_1(S_1), \ldots, C_T = c_T(S_T)$. As before, we want to solve

$$\sup\{\mathbb{E}(C_\tau) : \tau \lhd \mathfrak{F}\}$$

We form the backward recursion for the value function

$$\begin{aligned} V_T(S_T) &= c_T(S_T) \\ V_t(S_t) &= \max\{c_t(S_t), \mathbb{E}[V_{t+1}(S_{t+1}|S_t)]\} \end{aligned}$$

Then $V_0(S_0)$ is the optimal value of the stopping problem.

The Longstaff-Schwartz sampling method

Let $S_0^{(i)}, \ldots, S_T^{(i)}, \qquad i = 1, \ldots, n$ be n independent samples from the process (S_t).
We form the "approximate" pathwise backward recursion

$$\begin{aligned} \tilde{V}_T(S_T^{(i)}) &= c_T(S_T^{(i)} T) \\ \tilde{V}_t(S_t^{(i)}) &= \begin{cases} c_t(S_t^{(i)}) & \text{if} \quad c_t(S_t^{(i)}) \geq \tilde{\Psi}_t(S_t^{(i)}) \\ \tilde{V}_{t+1}(S_{t+1}^{(i)}) & \text{if} \quad c_t(S_t^{(i)}) < \tilde{\Psi}_t(S_t^{(i)}) \end{cases} \end{aligned}$$

where $\tilde{\Psi}_t(S_t^{(i)})$ is an approximation of $\mathbb{E}(\tilde{V}_{t+1}(S_{t+1})|S_t^{(i)})$. In order to find $\tilde{\Psi}_t(S_t^{(i)})$, we collect the pairs

$$(x_i = S_t^{(i)}, y_i = V_{t+1}(S_{t+1}^{(i)}))$$

and find a parametric or nonparametric regression estimate $\tilde{\Psi}_t \in \mathbb{F}$ which approximates $\mathbb{E}(y|x)$. Here $\mathbb{F}$ is an appropriate class of functions. The final value $\tilde{V}_0(S_0)$ is a lower bound for the true option value.

How to construct martingales for upper bounds?

In principle, for every process Z_t,

$$M_t = Z_0 + \sum_{s=1}^{t} Z_s - \mathbb{E}(Z_s|\mathfrak{G}_{s-1}).$$

is an $\mathfrak{F}$ martingale. Ideally $Z_t = V_t$, and $\mathfrak{F} = \mathfrak{G}$ i.e.

$$M_t^* = V_0 + \sum_{s=1}^{t} V_s - \mathbb{E}(V_s|\mathfrak{F}_{s-1})$$

but one may try to construct an approximant $\tilde{V}$ of V and and find an estimate for $\Psi_t(S_t) \sim \mathbb{E}(\tilde{V}_{t+1}|\mathcal{G}_t)$. The penalty $\mathfrak{G}$-martingale $(\tilde{M}_t)$ is then

$$\tilde{M}_t = \max(C_t, \Psi_{t+1}(S_t) - \mathbb{E}(\max(C_t, \Psi_{t+1}(S_t)|\mathcal{G}_t))$$

However, one needs extra forward samples for each sample $S_t^{(i)}$ which increases much the complexity.

The overestimation method for Longstaff-Schwartz

In order to find an upper bound, we do the very similar recursive procedure

$$\begin{aligned}
\tilde{V}_T(S_T^{(i)}) &= c_T(S_t^{(i)} T) \\
\tilde{V}_t(S_t^{(i)}) &= \begin{cases} c_t(S_t^{(i)}) & \text{if} \quad c_t(S_t^{(i)}) \geq \tilde{\tilde{\Psi}}_t(S_t^{(i)}) \\ \tilde{\tilde{\Psi}}_t(S_t^{(i)}) & \text{if} \quad c_t(S_t^{(i)}) < \tilde{\tilde{\Psi}}_t(S_t^{(i)}) \end{cases}
\end{aligned}$$

where $\tilde{\tilde{\Psi}}_t(x)$ is a guaranteed majorant of $\mathbb{E}(\tilde{V}_{t+1}(S_{t+1})|S_t^{(i)} = x)$. In order to find $\tilde{\tilde{\Psi}}_t$, we collect the pairs

$$(x_i = S_t^{(i)}, y_i = V_{t+1}(S_{t+1}^{(i)}))$$

and find a guaranteed majorant of the regression function, i.e. a $\tilde{\tilde{\Psi}}_t \in \mathbb{F}$ which majorizes $\mathbb{E}(y|x)$ for all x with a given error level β.

Guaranteed majorant regression

Let $(X_i, Y_i), i = 1, \ldots, n$ be a two-dimensional sample, stemming from a bivariate distribution, such that

$$\mathbb{E}(Y|X = x) = f_0(x).$$

Our goal is to find an estimate $\hat{f}_n(x)$ such that

$$P\{\hat{f}_n(x) \geq f_0(x) \quad \text{for all } x \} \geq 1 - \beta$$

Notice that we want this majorization to hold for all x and not only for the sampled X_i.
This topic is also treated by Alois Pichler in his talk on Monday. I am indebted to him for the idea to use the RKHS approach.

Reproducing Kernel Hilbert Spaces (RKHS)

Let $k(x,y)$ be a kernel function on $\mathbb{R}\times\mathbb{R}$ such that

- $k(x,y)=k(y,x)$
- $k(x,x)=1$; $0\le k(x,y)\le 1$
- $\sum_{i,j} v_i v_j k(x_i,x_j)\ge 0$ for all v_i, x_i (k is positive semidefinite)

(for instance $k(x,y)=\exp(-\alpha(x-y)^2)$).
Then the RKHS $\mathcal{H}$ is generated by the functions

$$x\mapsto \sum_i v_i k(x_i,x)$$

with inner product

$$<\sum_i v_i k(x_i,x),\sum_j w_i k(y_i,x)>_k=\sum_{i,j} v_i w_j k(x_i,y_j)$$

and is the closure of these functions under $\|f\|_k=\sqrt{<f,f>_k}$.

Basic properties

- Fundamental norm inequality: $\|f\|_\infty \leq \|f\|_k$
- Evalutation functional: $< k(x,\cdot), f(\cdot) >_k = f(x)$ for all x and all $f \in \mathcal{H}$.
- The operator K with $(Kf)(x) = \int k(x,y)\, f(y)\, dP(y)$ maps $\mathcal{H}$ to $\mathcal{H}$ and is self adjoint ($< Kf, g >_k = < f, Kg >_k$) and positive semidefinite ($< Kf, f >_k \geq 0$).
- Let $X = (X_1, \ldots, X_n)$ be an i.i.d. sample from P and let $\mathcal{H}_X$ be the Hilbert subspace generated by the functions $k(x, X_i), i = 1, \ldots, n$. If $\hat{P}$ is the empirical distribution ($\hat{P}_n = \frac{1}{n}\sum_i \delta_{X_i}$), then the operator $\hat{K}_n$ with

$$(\hat{K}_n f)(x) = \int k(x,y)\, f(y)\, d\hat{P}_n(y) = \frac{1}{n}\sum_i f(X_i) k(x, X_i)$$

maps $\mathcal{H}$ to $\mathcal{H}_X$ and is self adjoint and positive semidefinite.

The smoothing operator

Let $f \in \mathcal{H}$. Then the smoothed version $S_\lambda(f)$ is the argument minimum of

$$g \mapsto \|f-g\|_2^2+\lambda\|g\|_k^2 = \|f-g\|_2^2+\lambda \iint k(x,y)\, g(x)\, g(y)\, dP(x) dP(y).$$

The solution of this is

$$g = S_\lambda(f) = [K+\lambda I]^{-1} K f$$

where I is the identity. K has only nonnegative eigenvalues, say (μ_i), and hence $K+\lambda I$ is always invertible. Moreover, the eigenvalues of $[K+\lambda I]^{-1}K$ are $\mu_i/(\mu_i+\lambda)$ implying that $\|[K+\lambda I]^{-1}K\|_k \le 1$. Moreover,

$$\|S_\lambda(f) - f\|_k = \lambda \|[K+\lambda I]^{-1} f\|_k \le \lambda \|w\|_k$$

for $f = Kw$.

The emprirical smoothing operator

The empirical smoothing operator is the argument minimum of solution of

$$\begin{aligned} g &\mapsto \|f-g\|_2^2 + \iint k(x,y)\, g(x)\, g(y)\, d\hat{P}_n(x) d\hat{P}_n(y) \\ &= \|f-g\|_2^2 + \frac{1}{n^2}\sum_i k(X_i,X_j)\, g(X_i)\, g(X_j) \end{aligned}$$

The solution is

$$g(x) = \hat{S}_{\lambda,n}(f) = \sum_{i=1}^{n} v_i k(x, X_i)$$

where the vector $v = (v_1, \ldots, v_n)^\top$ satisfies

$$v = [\check{K}_n + \lambda I]^{-1} \cdot (f(X_1, \ldots, f(X_n))^\top$$

with $\check{K}_n$ being the $[n \times n]$ matrix with entries $\check{K}_n(i,j) = \frac{1}{n} k(X_i, X_j)$.
Notice that $\hat{S}_{\lambda,n}(f) \in \mathcal{H}_X =$ span $\{k(x, X_i); i = 1, \ldots, n\}$.

The regression estimate

Let $(X_i, Y_i)_{i=1,\dots,n}$ be the observed pairs. Then the regression estimate is

$$\hat{f}_n(x) = \sum_i v_i k(x, X_i)$$

where the vector $v = (v_1, \dots, v_n)^\top$ satisfies

$$v = [\check{K}_n + \lambda I]^{-1} \cdot (Y_1, \dots, Y_n)^\top.$$

If f_0 is the true regression function, then set

$$Y_i = f_0(X_i) + \epsilon_i$$

where the ϵ's are conditionally independent given the X_i's. Therefore

$$\hat{f}_n = \sum_i (w_i + z_i) k(x, X_i) = \hat{m}_n + \hat{R}_n$$

with

$$w = [\check{K}_n + \lambda I]^{-1} \cdot (f_0(X_1), \dots, f_0(X_n))^\top,$$
$$z = [\check{K}_n + \lambda I]^{-1} \cdot (\epsilon_1, \dots, \epsilon_n)^\top.$$

The main part $\hat{m}_n = \sum_i w_i k(x, X_i)$

Let $f_0(x) = \int k(x,y) w_0(y) dP(y)$. Then

$$\hat{m}_n(x) = \frac{1}{n} \sum_j ([\check{K}_n + \lambda I]^{-1} K w_0(X))_j k(x, X_j).$$

Let

$$\bar{f}_{0,n} = \frac{1}{n} \sum_i k(x, X_i) w_0(X_i) = \int k(x,y) w_0(y) d\hat{P}_n(y).$$

Then $\bar{f}_{0,n} \in \mathcal{H}_X$.

$$\begin{aligned} \|f_0 - \bar{f}_{0,n}\|_\infty &= \| \int k(x,y) w_0(y) dP(y) - \int k(x,y) w_0 d\hat{P}_n(y) \|_k \\ &= \leq [L(k)\|w_0\|_\infty + L(w_0)] d_W(\hat{P}_n, P) \end{aligned}$$

where $L(\cdot)$ are Lipschitz constants and $d_W(\hat{P}_n, P)$ is the Wasserstein distance between P and $\hat{P}_n$.

$$\begin{aligned}\|\hat{m}_n(x) - \bar{f}_{0,n}\|_\infty &\leq \lambda\|w_0\|_k + \mathbf{1}^\top[\hat{K}_n + \lambda I]^{-1}\check{K}_n[\check{K}_n + \lambda I]^{-1}]\mathbf{1} \times \\ &\quad \times[L(k)\|w_0\|_\infty + L(w_0)]d_W(\hat{P}_n, P).\end{aligned}$$

and hence

$$\begin{aligned}\|\hat{m}_n(x) - f_0(x)\|_\infty &\leq \lambda\|w_0\|_k \\ &+ (1 + \mathbf{1}^\top[\hat{K}_n + \lambda I]^{-1}\check{K}_n[\check{K}_n + \lambda I]^{-1}]\mathbf{1}) \\ &\times [L(k)\|w_0\|_\infty + L(w_0)]^2 d_W^2(\hat{P}_n, P) \\ &= \lambda\|w_0\|_k + c(\check{K}_n, \lambda, \hat{P}_n, P)\end{aligned}$$

The error part $\hat{R}_n$

For the error part $\hat{R}_n$, we have

$$\hat{R}_{\lambda,n}(x) = \frac{1}{n}\sum_j ([\hat{K}_n + \lambda I]^{-1}\epsilon)_j k(x, X_j)$$

with $\epsilon = (\epsilon_1, \ldots, \epsilon_n)^\top$. Conditioned on the sample X and assuming that $\mathbb{E}(\epsilon_i) \le \sigma^2$ one gets

$$\begin{aligned}
\mathbb{E}_X\left[\|\hat{R}_{\lambda,n}(.)\|_k^2\right] &\le \frac{1}{n^2}\sum_{i,j} k(X_i, X_j) e_i^\top [\check{K}_n + \lambda I]^{-1}\, \sigma^2\, [\check{K}_n + \lambda I]^{-1} e_j \\
&= \frac{\sigma^2}{n}\ \mathrm{tr}\ [\lambda\check{K}_n + \lambda I]^{-1}\check{K}_n[\check{K}_n + \lambda I]^{-1}) \\
&\le \frac{\sigma^2}{\lambda n}\kappa(\check{K}_n, \lambda)
\end{aligned}$$

where $\kappa(\hat{K}_n, \lambda) = \ \mathrm{tr}(\lambda[\hat{K}_n + \lambda I]^{-1}\check{K}_n[\check{K}_n + \lambda I]^{-1})$.

Suppose that $\hat{K}_n$ has eigenvalues $\hat{\mu}_i$. Then in the worst case

$$\kappa(\hat{K}_n, \lambda) = \sum_i \frac{\lambda \hat{\mu}_i}{(\lambda + \hat{\mu}_i)^2} \leq 1/\lambda \sum_i \hat{\mu}_i = 1/\lambda,$$

since $\sum_i \hat{\mu}_i = \text{ tr } \hat{K}_n = 1$. However, if the kernel functions are chosen in the right way, then $\kappa(\hat{K}_n, \lambda))$ is $O(1)$ as $\lambda \to 0$, since only a few eigenvalues are positive.
Consequently also integrating over all samples $X = (X_i)$,

$$\mathbb{E}\left[\|\hat{R}_{\lambda,n}(.)\|_k^2\right] \leq \frac{\sigma^2}{\lambda n} \kappa(\check{K}_n, \lambda).$$

and

$$P\left\{\|\hat{R}_{\lambda,n}(.)\|_k^2 \geq \eta\right\} \leq \mathbb{E}\left[\|\hat{R}_{\lambda,n}(.)\|_k^2\right]/\eta^2 \leq \frac{\sigma^2}{\eta^2 \lambda n} \kappa(\check{K}_n, \lambda).$$

From above we have for the main part

$$\|\hat{m}_n(x) - f_0(x)\|\infty \quad \leq \quad \lambda\|w_0\|_k + c(\check{K}_n, \lambda, \hat{P}_n, P).$$

Notice that

$$P\{d_W(\hat{P}_n, P) > \epsilon\} \leq \exp(-\frac{\lambda}{2} n\epsilon^2)$$

(Bolley et al. 2007) or we use a fixed design for the X_i.
Now choose β as the error for the significance level. Let

$$\eta = \frac{\sigma^2 \kappa(\check{K}_n, \lambda)}{n\lambda\beta}.$$

Then

$$P\{f_0(x) \leq \hat{f}_n(x) + \lambda\|w_0\|_k + c(\check{K}_n, \lambda, \hat{P}, P) + \eta \text{ for all } x \} \geq 1 - \beta.$$

Example

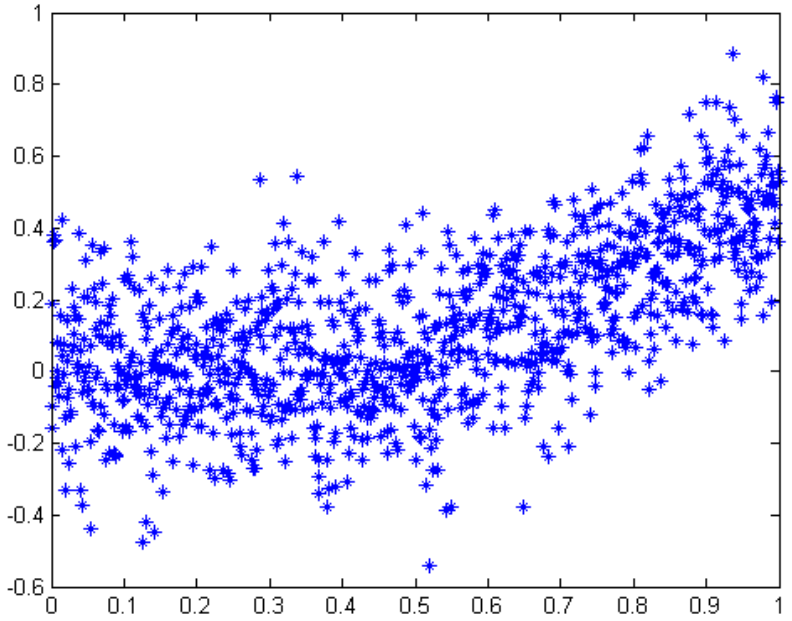

$$
\begin{aligned}
n &= 1000 \\
k(x,y) &= \exp(-20 \cdot (x-y)^2) \\
\lambda &= 0.01 \\
\sigma &= 0.2 \\
\beta &= 0,01
\end{aligned}
$$

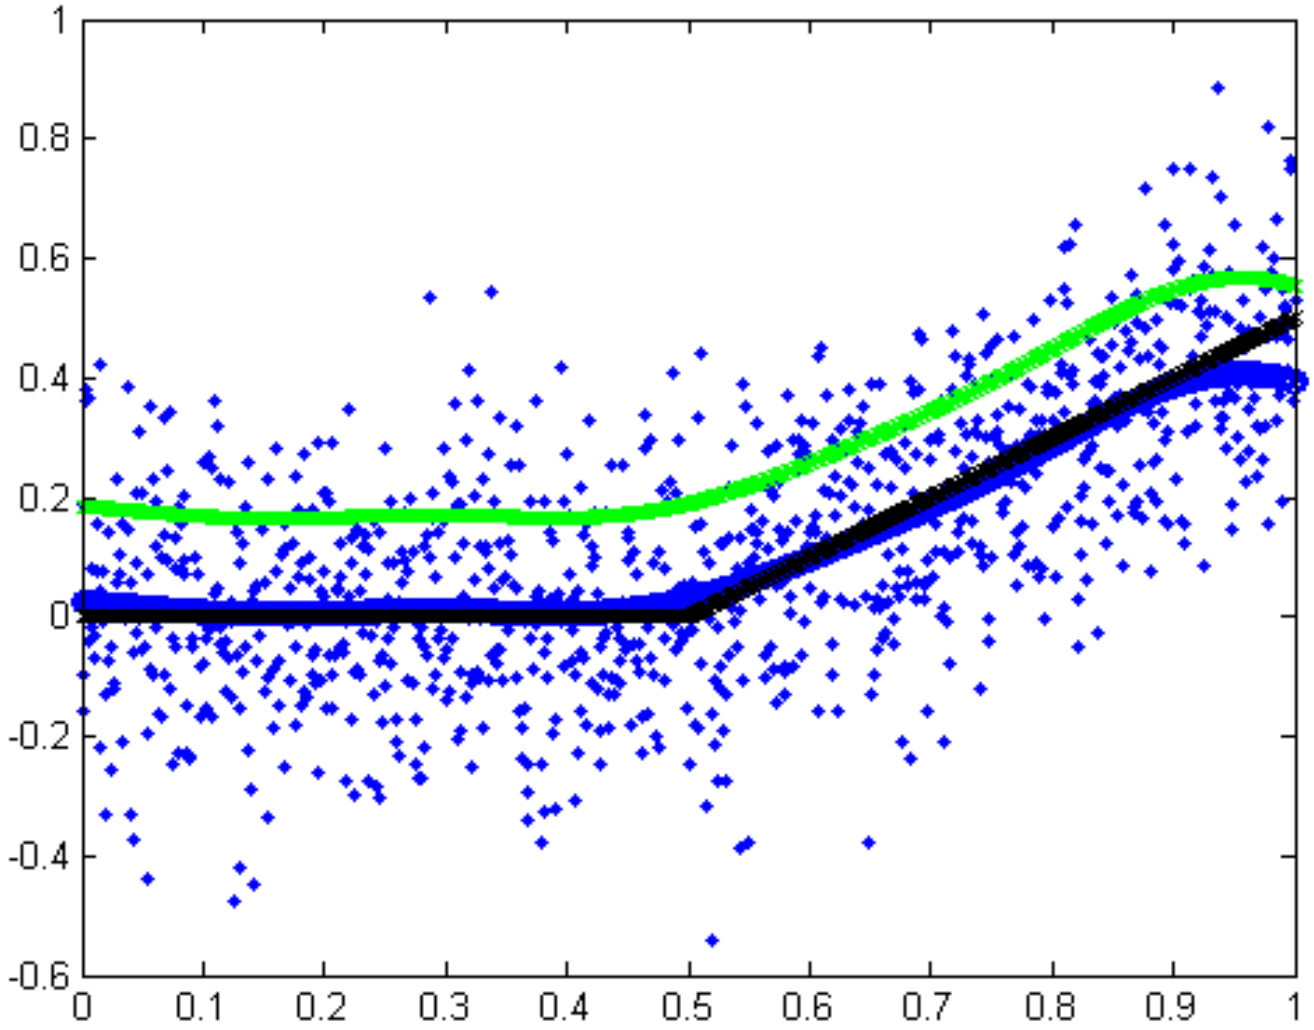

The matrix $\check{K}$ has 1000 eigenvalues, but only 17 of them are larger than 10^{-4}.

What needs to be done

- Using the guaranteed overestimation in Longstaff-Schwartz regression at each time step t.
- Relaxing the conditions about Lipschitz functions to Fortet-Mourier functions:

$$\mathcal{FM}_p = \{f : \exists L \text{ s.t. } |f(u)| \le L(\max(1,|u|)^p); \\ |f(u) - f(v)| \le L|u-v| \max(1,|u|,|v|)^{p-1}\} \qquad (2)$$

  The smallest L satifying (2) is called the Fortet-Mourier constant. We may replace Wasserstein distance by the Fortet-Mourier distance.
- Comparing this with the other methods for getting upper bounds.