Séminaire de Mathématiques Appliquées du CERMICS

École des Ponts
ParisTech

A least-squares Monte-Carlo approach to rare events simulation

Carsten Hartmann (BTU Cottbus-Senftenberg)

27 novembre 2018

A least-squares Monte-Carlo approach to rare events simulation

Carsten Hartmann (BTU Cottbus-Senftenberg),
with Omar Kebiri, Lara Neureither and Lorenz Richter

CERMICS Applied Mathematics Seminar
École des Ponts Paris Tech, 27 November 2018

Motivation: conformation dynamics of biomolecules

Protein folding

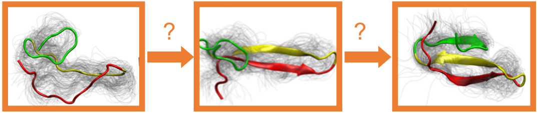

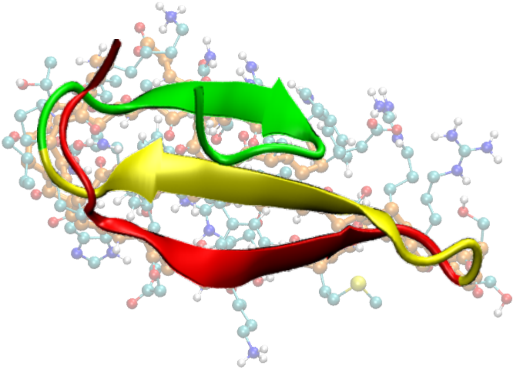

[Noé et al, PNAS, 2009]

Motivation: conformation dynamics of biomolecules

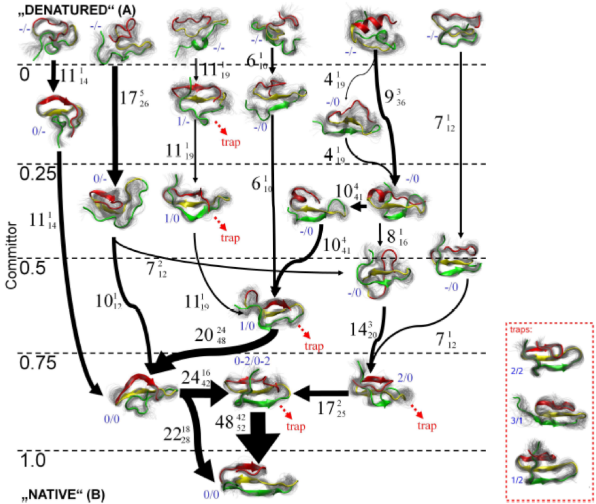

[Noé et al, PNAS, 2009]

Motivation: conformation dynamics of biomolecules

Given a **Markov process** $X = (X_t)_{t \geq 0}$, discrete or continuous in time, we want to **estimate probabilities** $p \ll 1$, such as

$$p = P\left(\tau < T\right),$$

or **rates**

$$k = \left(\mathbb{E}[\tau]\right)^{-1}$$

with τ some random first passage time and $\mathbb{E}[\cdot]$ the expectation with respect to the probability P.

Motivation: conformation dynamics of biomolecules

More specifically, we want to estimate **free energies**

$$F = -\log \mathbb{E}\left[e^{-W}\right],$$

where W is some functional of X.

For example, with $W = \alpha\tau$ and sufficiently small $\alpha > 0$, we have

$$-\alpha^{-1} F = \mathbb{E}[\tau] + \mathcal{O}(\alpha)$$

Illustrative example: bistable system

- Overdamped Langevin equation

$$dX_t = -\nabla V(X_t)dt + \sqrt{2\epsilon}dB_t\,.$$

- MC estimator of $\psi = \mathbb{E}[e^{-\alpha\tau_C}]$

$$\hat{\psi}_\epsilon^N = \frac{1}{N}\sum_{i=1}^{N} e^{-\alpha\tau_C^i}\,.$$

- Small noise asymptotics (Kramers)

$$\lim_{\epsilon\to 0} \epsilon \log \mathbb{E}[\tau_C] = \Delta V\,.$$

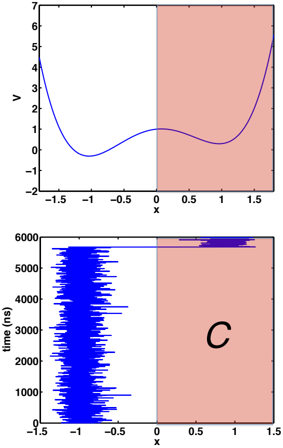

[Freidlin & Wentzell, 1984], [Berglund, Markov Processes Relat Fields 2013]

Illustrative example, cont'd

- **Relative error** of the MC estimator

$$\delta_\epsilon = \frac{\sqrt{\mathrm{Var}[\hat{\psi}_N^\epsilon]}}{\mathbb{E}\left[\hat{\psi}_N^\epsilon\right]}$$

- Varadhan's large deviations principle

$$\mathbb{E}\left[(\hat{\psi}_\epsilon^N)^2\right] \gg (\mathbb{E}\left[\hat{\psi}_\epsilon^N\right])^2\,,\ \epsilon \text{ small.}$$

- Unbounded relative error as $\epsilon \to 0$

$$\limsup_{\epsilon \to 0} \delta_\epsilon = \infty$$

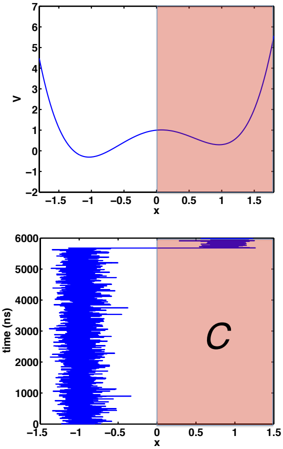

[Dupuis & Ellis, 1997]

Outline

Importance sampling of rare events

Duality of estimation and control

Least-squares regression

Numerical examples

- Importance sampling of rare events
- Duality of estimation and control
- Least-squares regression
- Numerical examples

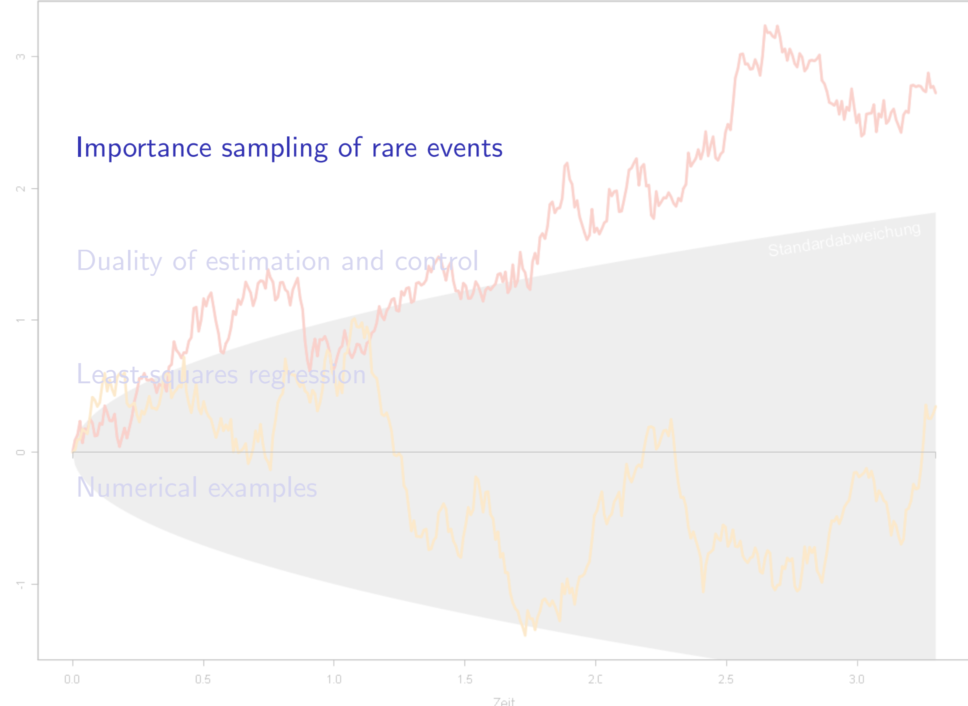

Optimal change of measure: zero variance

Pick another probability measure Q with likelihood ratio

$$\varphi = \frac{dQ}{dP} > 0\,,$$

under which the **rare event is no longer rare**, such that

$$\mathbb{E}[\exp(-\alpha\tau_C)] = \mathbb{E}_Q\left[\exp(-\alpha\tau_C)\varphi^{-1}\right].$$

Zero-variance change of measure exists and is given by

$$\varphi^* = \frac{dQ^*}{dP} = \frac{\exp(-\alpha\tau_C)}{\mathbb{E}[\exp(-\alpha\tau_C)]}\,,$$

but it depends on the quantity of interest, $\mathbb{E}[\exp(-\alpha\tau_C)]$.

Approaching zero variance (non-exhaustive list)

- Exponential tilting based on large deviations statistics:

$$dQ^* \approx \exp(\gamma - \alpha\tau_C)dP \quad \text{as } \epsilon \to 0,$$

where γ is related to the large deviations rate function.

Siegmund, Glasserman & Kou, Dupuis & Wang, Vanden-Eijnden & Weare, Spiliopoulos, ...

- Kullback-Leibler or cross-entropy minimisation:

$$Q^* \approx \operatorname*{argmin}_{Q \in \mathcal{M}} KL(Q, Q^*)\,,$$

with Q from some suitable ansatz space $\mathcal{M}$.

Rubinstein & Kroese, Zhang & H, Kappen & Ruiz, Opper, Quer, ...

- Mean square and work-normalised variance minimisation

Glynn & Whitt, Jourdain & Lelong, Su & Fu, Vázquez-Abad & Dufresne, ...

Another idea . . .

Exponential tilting from nonequilibrium forcing

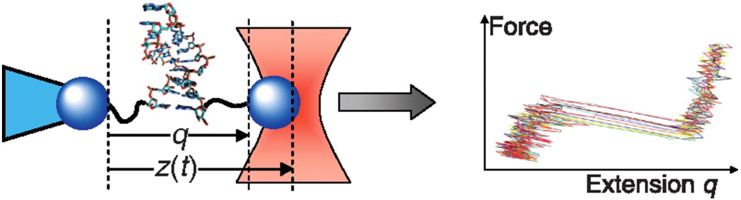

Single molecule pulling experiments, figure courtesy of G. Hummer, MPI Frankfurt

In vitro/in silico **free energy calculation** from forcing:

$$F = -\log \mathbb{E}\left[e^{-W}\right].$$

Forcing generates a "nonequilibrium" path space measure Q with typically **suboptimal likelihood quotient** $\varphi = dQ/dP$.

[Schlitter, J Mol Graph, 1994], [Hummer & Szabo, PNAS, 2001], Schulten & Park, JCP, 2004], ...

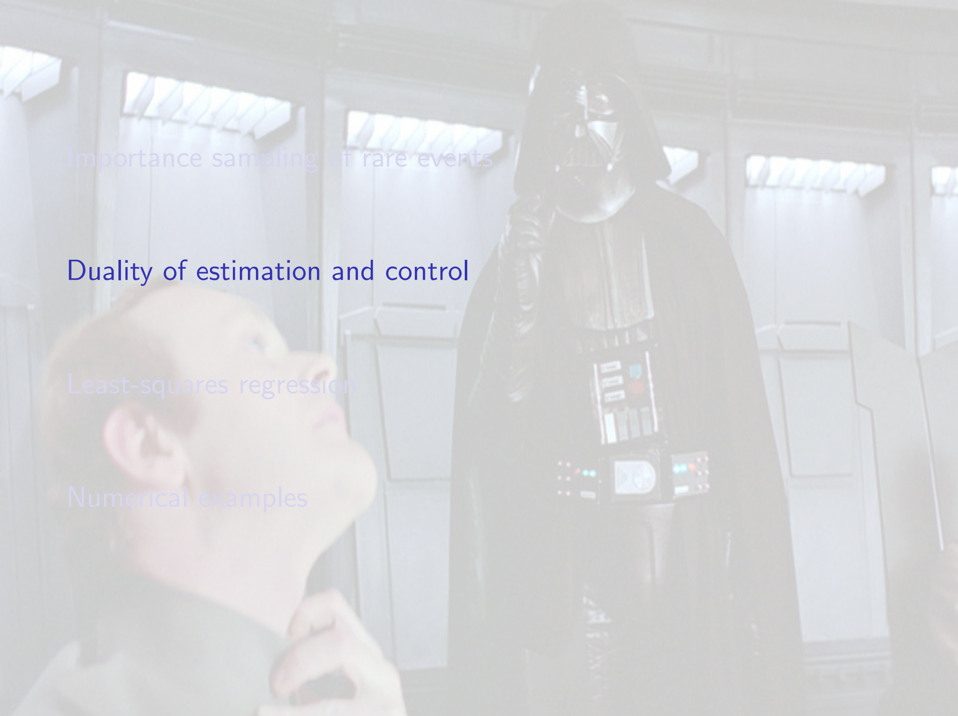

Importance sampling of rare events

Duality of estimation and control

Least-squares regression

Numerical examples

Variational characterization of free energy

Theorem (Gibbs, Donsker & Varadhan, ...)

For any bounded and measurable function W it holds

$$-\log \mathbb{E}\left[e^{-W}\right] = \min_{Q \ll P} \{\mathbb{E}_Q[W] + KL(Q,P)\}$$

where $KL(Q,P) \geq 0$ is the **relative entropy** between Q and P:

$$KL(Q,P) = \begin{cases} \int \log\left(\frac{dQ}{dP}\right) dQ & \text{if } Q \ll P \\ \infty & \text{otherwise} \end{cases}$$

Sketch of proof: Let $\varphi = dQ/dP$. Then

$$-\log \int e^{-W} dP = -\log \int e^{-W-\log\varphi} dQ \leq \int (W + \log\varphi)\, dQ$$

[Boué & Dupuis, LCDS Report #95-7, 1995], [Dai Pra et al, Math Control Signals Systems, 1996]

Same same, but different. . .

Set-up: uncontrolled ("equilibrium") diffusion process

Let $X = (X_s)_{s \geq 0}$ be a **diffusion process** on $\mathbb{R}^n$,

$$dX_s = b(X_s)ds + \sigma(X_s)dB_s \,, \quad X_0 = x \,,$$

and

$$W(X) = \int_0^\tau f(X_s)\,ds + g(X_\tau)\,,$$

for suitable functions f, g and a **a.s. finite stopping time** $\tau < \infty$.

Aim: Estimate the path functional

$$\psi(x) = \mathbb{E}\big[e^{-W(X)}\big]$$

Set-up: controlled ("nonequilibrium") diffusion process

Now given a **controlled diffusion process** $X^u = (X^u_s)_{s\geq 0}$,

$$dX^u_s = (b(X^u_s) + \sigma(X^u_s)u_s)ds + \sigma(X^u_s)dB_s\,, \quad X^u_0 = x\,,$$

and a probability $Q \ll P$ on $C([0,\infty))$ with **likelihood ratio**

$$\varphi(X^u) = \frac{dQ}{dP}\bigg|_{\mathcal{F}_\tau} = \exp\left(-\int_0^\tau u_s \cdot dB_s - \frac{1}{2}\int_0^\tau |u_s|^2\,ds\right)\,.$$

Now: Estimate the reweigthed path functional

$$\mathbb{E}\big[e^{-W(X)}\big] = \mathbb{E}\big[e^{-W(X^u)}(\varphi(X^u))^{-1}\big]$$

Variational characterization of free energies, cont'd

Theorem (H, 2012/2017)

Technical details aside, let u^* be a minimiser of the cost functional

$$J(u) = \mathbb{E}\left[W(X^u) + \frac{1}{2}\int_0^\tau |u_s|^2\, ds\right]$$

under the **controlled dynamics**

$$dX_s^u = (b(X_s^u) + \sigma(X_s^u)u_s)ds + \sigma(X_s^u)dB_s\,, \quad X_0^u = x\,.$$

The **minimiser is unique** with $J(u^*) = -\log\psi$. Moreover,

$$\psi(x) = e^{-W(X^{u^*})}(\varphi(X^{u^*}))^{-1} \quad \text{(a.s.)}\,.$$

[H & Schütte, JSTAT, 2012], [H et al, Entropy, 2017]

Illustrative example, cont'd

- Exit problem: $f = \alpha$, $g = 0$, $\tau = \tau_C$:

$$J(u^*) = \min_u \mathbb{E}\left[\alpha \tau_C^u + \frac{1}{4\epsilon} \int_0^{\tau_C^u} |u_s|^2 ds\right],$$

under the **tilted dynamics**

$$dX_t^u = (u_t - \nabla V(X_t^u))\, dt + \sqrt{2\epsilon} dB_t$$

- Optimally tilted potential

$$U^*(x, t) = V(x) - u_t^* \, x$$

with **stationary** feedback $u_t^* = c(X_t^{u^*})$.

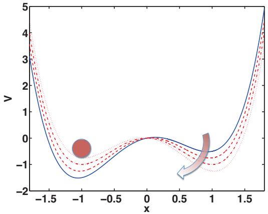

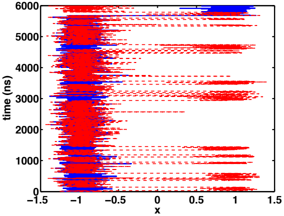

Sketch of proof: Fleming's log transformation

By the **Feynman-Kac theorem**,

$$\psi(x) = \mathbb{E}\left[\exp\left(-\int_0^\tau f(X_t)dt - g(X_\tau)\right)\,\middle|\, X_0 = x\right]$$

solves a linear BVP on an open and bounded set $O \subset \mathbb{R}^n$:

$$(\mathcal{L} - f)\psi = 0\,, \quad \psi|_{\partial\Omega_+} = \exp(-g) \qquad (\mathcal{L}\text{ generator})$$

The corresponding **semilinear BVP** for $F = -\log\psi$ reads

$$\mathcal{L}F - \frac{1}{2}|\nabla F|_a^2 + f = 0\,, \quad F|_{\partial\Omega_+} = g \qquad \text{with } \; a = \sigma\sigma^T$$

[H et al, JSTAT, 2012]; cf. [Fleming, SIAM J Control, 1978], [Boué & Dupuis, Ann Probab., 1998]

Sketch of proof, cont'd

The semilinear Hamilton-Jacobi-Bellman PDE

$$\mathcal{L}F - \frac{1}{2}|\nabla F|_a^2 + f = 0\,, \quad F|_{\partial\Omega_+} = g \qquad (a = \sigma\sigma^T)$$

is the **dynamic programming equation** for our stochastic control problem; it solution is the value function

$$F(x) = \min\{J(u) \colon X_0^u = x\}$$

If $F \in C^2(O) \cap C(\overline{O})$, the optimal control has **gradient form**, i.e.

$$u_t^* = -\sigma(X_t^{u^*})^T \nabla F(X_t^{u^*})\,,$$

Generalizations: degenerate diffusions, Markov chains, ...

[Schütte et al, Math Prog, 2012], [Banisch & Hartmann, MCRF, 2016], [H et al, Entropy, 2017]

- Importance sampling of rare events
- Duality of estimation and control
- **Least-squares regression**
- Numerical examples

From dynamic programming to a pair of SDE

The **semilinear HJB equation**

$$\mathcal{L}F + h(x, F, \sigma^T \nabla F) = 0\,,\ F|_{\partial\Omega_+} = g$$

is equivalent to the uncoupled **forward-backward SDE**

$$\begin{aligned} dX_s &= b(X_s)ds + \sigma(X_s)dB_s\,,\ X_t = x \\ dY_s &= -h(X_s, Y_s, Z_s)ds + Z_s \cdot dB_s\,,\ Y_\tau = g(X_\tau)\,, \end{aligned}$$

on a random time horizon $[0, \tau] \subset [0, \infty)$ where

$$Y_s = F(X_s)\,, \quad Z_s = \sigma(X_s)^T \nabla F(X_s).$$

Formal derivation: Itô's Lemma

[Pardoux & Peng, LNCIS 176, 1992], [Kobylanski, Ann Probab, 2000]

Some remarks

- The solution of the **forward-backward SDE** (FBSDE)

$$dX_s = b(X_s)ds + \sigma(X_s)dB_s\,,\ X_0 = x$$
$$dY_s = -h(X_s, Y_s, Z_s)ds + Z_s \cdot dB_s\,,\ Y_\tau = g(X_\tau)\,,$$

is a **triplet** (X, Y, Z) where (Y_s, Z_s) is adapted to $(X_u)_{u\leq s}$.

- Hence $Y_0 = F(x)$ is a **deterministic** function of the initial data $X_0 = x$, and Z controls this property.

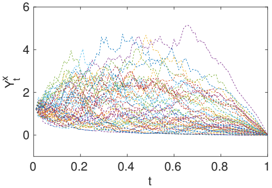

- **Existence and uniqueness** of (Y, Z) is guaranteed for terminal conditions g that are bounded & Lipschitz.

[Pardoux & Peng, LNCIS 176, 1992], [Kobylanski, Ann Probab, 2000]

Some remarks, cont'd

- A BSDE is **not a time-reversed SDE**; for example, the SDE

$$dY_s = Z_s \cdot dB_s\,, \quad Y_T = X_T$$

has two possible formal solutions

$$(Y_s, Z_s) = (X_s, 1) \quad \text{and} \quad (\tilde{Y}_s, \tilde{Z}_s) = (X_T, 0)\,,$$

only **one of which is adapted**, namely (Y_s, Z_s).

- **A fix:** project $\tilde{Y}_s$ onto the filtration generated by $(X_t)_{t\geq 0}$,

$$Y_s = \mathbb{E}[X_T|\mathcal{F}_s]\,.$$

- Then, by the **martingale representation theorem**, there is a unique, predictable process $Z \in L^2$ such that

$$Y_t = \mathbb{E}[X_T] + \int_0^t Z_s \cdot dB_s \quad \Longrightarrow \quad Y_t = X_T - \int_t^t Z_s \cdot dB_s\,.$$

Numerical discretisation of FBSDE

The **FBSDE is decoupled** and an explicit scheme can be based on

$$\hat{X}_{n+1} = \hat{X}_n + \Delta t\, b(\hat{X}_n) + \sqrt{\Delta t}\, \sigma(\hat{X}_n)\xi_{n+1}$$
$$\hat{Y}_{n+1} = \hat{Y}_n - \Delta t\, h(\hat{X}_n, \hat{Y}_n, \hat{Z}_n) + \sqrt{\Delta t}\, \hat{Z}_n \cdot \xi_{n+1}$$

with boundary values

$$\hat{X}_0 = x\,, \quad \hat{Y}_{n_\tau} = g(\hat{X}_{n_\tau})$$

Solution to stochastic two-point boundary value problem:

- least-squares Monte Carlo Gobet & Turkedjev, Bender et al.
- deep neural network approach E, Han & Jentzen

Solution by least-squares Monte-Carlo

Numerical discretisation of FBSDE, con't

Since $\hat{Y}_n$ **is adapted** we have $\hat{Y}_n = \mathbb{E}\big[\hat{Y}_n|\mathcal{F}_n\big]$ and thus

$$\begin{aligned}\hat{Y}_n &= \mathbb{E}\big[\hat{Y}_{n+1} + \Delta t\, h(\hat{X}_n, \hat{Y}_n, \hat{Z}_n)|\mathcal{F}_n\big] \\ &\approx \mathbb{E}\big[\hat{Y}_{n+1} + \Delta t\, h(\hat{X}_n, \hat{Y}_{n+1}, \hat{Z}_{n+1})|\mathcal{F}_n\big]\end{aligned}$$

where $\mathcal{F}_n = \sigma(\hat{X}_0, \ldots, \hat{X}_n)$ using that $\hat{Z}_n$ is independent of ξ_{n+1}.

The conditional expectation

$$\hat{Y}_n := \mathbb{E}\big[\hat{Y}_{n+1} + \Delta t\, h(\hat{X}_n, \hat{Y}_{n+1}, \hat{Z}_{n+1})|\mathcal{F}_n\big]$$

can be computed by **least-squares**:

$$\mathbb{E}\big[S|\mathcal{F}_n\big] = \underset{Y \in L^2,\, \mathcal{F}_n\text{-measurable}}{\operatorname{argmin}} \mathbb{E}[|Y - S|^2]\,.$$

[Gobet et al, AAP, 2005], [Bender & Steiner, Num Meth F, 2012], [Kebiri et al, Proc IHP, 2018]

Deep learning based approximation

Numerical discretisation of FBSDE, con't

Now consider the **forward iteration**

$$\mathcal{Y}_{n+1} = \mathcal{Y}_n - \Delta t\, h(\hat{X}_n, \mathcal{Y}_n, \mathcal{Z}_n) + \sqrt{\Delta t}\, \mathcal{Z}_n \cdot \xi_{n+1}\,,$$

with $\mathcal{Y}_n = \mathcal{Y}_n(x;\theta)$ and $\mathcal{Z}_n = \mathcal{Z}_n(\hat{X}_n;\theta)$ being the (non-adapted?) **deep neural net approximation** of $(\hat{Y}_n, \hat{Z}_n)$, so that

$$\mathcal{Y}_0 \approx F(x)\,, \quad \mathcal{Z}_n \approx (\sigma^T \nabla F)(\hat{X}_n)$$

The corresponding **loss function** is given by

$$\ell(\theta) = \mathbb{E}\left[|\mathcal{Y}_{n_\tau} - g(\hat{X}_{n_\tau})|^2\right]$$

(Note that $\mathbb{E}\left[|\mathcal{Y}_\tau - g(X_\tau)|^2\right] = 0$ for the exact solution.)

[E et al, Commun Math Stat, 2017], [H. et al, Preprint, 2018]

More remarks

- The LSMC scheme is **strongly convergent** of order $1/2$ in $\Delta t \to 0$ as $M, K \to \infty$ (M: sample size, K: # basis fcts.).
- A **zero-variance change of measure** is given by

$$\left.\frac{dQ}{dP}\right|_{\mathcal{F}_\tau} = \exp\left(\int_0^\tau Z_s \cdot dB_s - \frac{1}{2}\int_0^\tau |Z_s|^2\, ds\right)\,,$$

  for $\tau < \infty$ (a.s.) and the discretisation bias can be further reduced by using **importance sampling**.
- Under mild assumptions, the variance of the importance sampling estimator is **no worse than for crude MC**.
- **Generalisations include** bounded & deterministic τ, singular terminal condition, least-squares w/ change of drift.

[Turkedjiev, PhD thesis, 2013], [Kruse & Popier, SPA, 2016], [Kebiri & H, Preprint, 2018]

Importance sampling of rare events

Duality of estimation and control

Least-squares regression

Numerical examples

Example I: hitting probabilities

Probability of **hitting the set** $C \subset \mathbb{R}$ before time T:

$$-\log \mathbb{P}(\tau \leq T) = \min_u \mathbb{E}\left[\frac{1}{4}\int_0^{\tau\wedge T} |u_t|^2\, dt - \log \mathbf{1}_{\partial C}(X^u_{\tau\wedge T})\right],$$

with τ denoting the first hitting time of C under the dynamics

$$dX^u_t = (u_t - \nabla V(X^u_t))\, dt + \sqrt{2\epsilon}\, dB_t$$

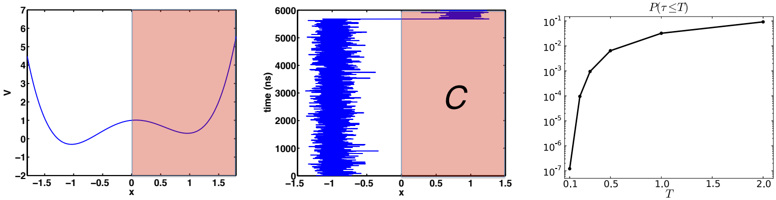

[Zhang et al, SISC, 2014], [Richter, MSc thesis, 2016], [H et al, Nonlinearity, 2016]

Example I, cont'd

Probability of **hitting** $C \subset \mathbb{R}$ before time T, starting from $x = -1$:

$$-\log \mathbb{P}(\tau \leq T) = \min_u \mathbb{E}\left[\frac{1}{4}\int_0^{\tau \wedge T} |u_t|^2 \, dt - \log \mathbf{1}_{\partial C}(X^u_{\tau \wedge T})\right],$$

(BSDE with singular terminal condition and random stopping time)

Simulation parameters	$F^{\epsilon}_{ref}(0,x)$	$\bar{F}^{\epsilon}(0,x)$	Var
$K=8,\ M=300,\ T=5,\ \Delta t=10^{-3},\ \epsilon=1$	0.3949	0.3748	10^{-3}
$K=5,\ M=300,\ T=1,\ \Delta t=10^{-3},\ \epsilon=1$	1.7450	1.6446	0.0248
$K=5, M=400, T=1, \Delta t=10^{-4}, \epsilon=0.6$	4.3030	4.5779	10^{-3}
$K=6, M=450, T=1, \Delta t=10^{-4}, \epsilon=0.5$	4.5793	4.6044	$5 \cdot 10^{-4}$

with K the number of Gaussians and M the number of realisations of the forward SDE.

[Ankirchner et al, SICON, 2014], [Kruse & Popier, SPA, 2016], [Kebiri et al, Proc IHP, 2018]

Example II: High-dimensional PDE

First exit time of a **Brownian motion** from an *n*-sphere of radius *r*:

$$\tau = \inf\{t > 0 \colon x + B_t \notin S^n_r\}$$

Cumulant generating function of first exit time satisfies

$$-\log \mathbb{E}[\exp(-\alpha\tau)] = \min_u \mathbb{E}\left[\alpha\tau^u + \frac{1}{2}\int_0^{\tau^u} |u_t|\, dt\right]$$

- BSDE on random time horizon with homogeneous terminal condition
- mean first exit time $\mathbb{E}[\tau] = \frac{r^2 - |x|^2}{n}$
- Least-squares MC w/ $K = 3, M \sim 10^2$

	$n = 3$	$n = 10$	$n = 100$	$n = 1000$
exact	1.00	1.00	1.00	1.00
CMC	0.98	0.99	1.08	1.04
LSMC	0.99	1.01	0.96	0.98

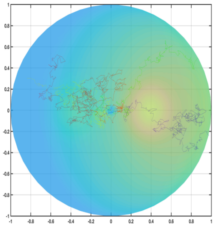

[Kebiri & H, Preprint, 2018]

Discussion: Markovian approximations

- The fact that the FBSDE is uncoupled implies that every approximation of X gives rise to an approximation of (Y, Z).
- **Slow-fast systems:** Strong error bound for limit BSDE

$$\sup\{|Y_t^\delta - \bar{Y}_t| \colon 0 \leq t \leq T\} \leq C\sqrt{\delta} \quad \delta = \frac{\tau_{\text{fast}}}{\tau_{\text{slow}}}$$

  as $\delta \to 0$, analogously for Z_t^δ (implies IS error bound).
- **Reversible metastable systems:** Approximation of X by s-state Markov jump process $(j_t)_{t\geq 0}$ implies that

$$\|\psi - \bar{\psi}\|_2 \leq C \inf_{v \in H} \|\psi - v\|_2$$

  with $C \approx 1$ and basis dependent best-approximation error.

[Banisch & H, MCRF, 2016], [H et al, PTRF, 2018], [Kebiri & H, Proc IHP, 2018]

Discussion: approximation by DNN

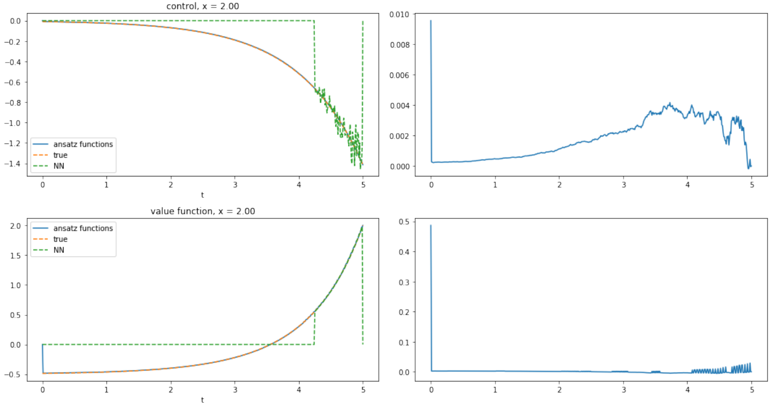

[H et al, Preprint, 2018]

Conclusions & outlook

- Adaptive importance sampling scheme based on **dual variational formulation**; resulting control problem features short trajectories with **minimum variance estimators**.
- **Variational problem** boils down to an uncoupled FBSDE with only one additional spatial dimension.
- **Error analysis** for unbounded stopping time & singular terminal condition is open, **least-squares algorithm** requires some fine-tuning (ansatz functions, change of drift, . . .).
- Clever choice of ansatz functions should be combined with **dimension reduction** (cf. results for slow-fast systems).

Thank you for your attention!

Acknowledgement:

Ralf Banisch
Omar Kebiri
Lara Neurither
Jannes Quer
Lorenz Richter
Christof Schütte
Wei Zhang

German Science Foundation (DFG)
Einstein Center for Mathematics Berlin (ECMath)