Duality of estimation and control an its application to rare events simulation

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

Computational Statistics and Molecular Dynamics
CMO-BIRS, Oaxaca, 11–16 November 2018

Take-home message

- **Donsker & Varadhan:** Monte Carlo sampling of nonnegative random variables has an equivalent variational formulation.
- For path-dependent random variables the variational formulation boils down to an **optimal control** problem.
- The **numerical toolbox** for solving optimal control problems is different from the Monte Carlo toolbox.

Motivation: conformation dynamics of biomolecules

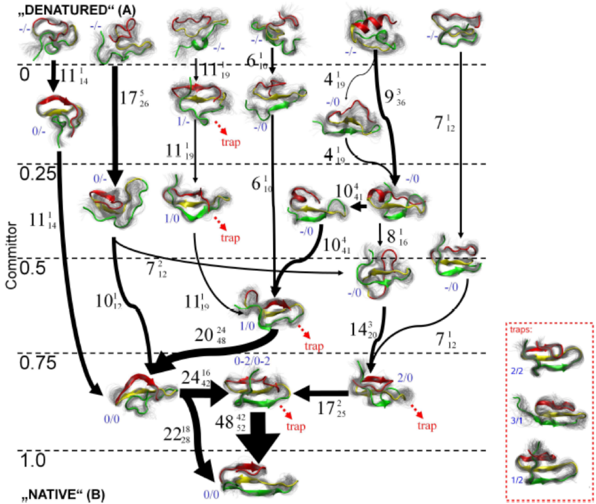

[Noé et al, PNAS, 2009]

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**, such as

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

with τ some random stopping time, or **free energies**

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

where W is some functional of X.

Outline

Importance sampling of rare events

Duality of estimation and control

Strategies for high-dimensional problems

Importance sampling of rare events

Duality of estimation and control

Strategies for high-dimensional problems

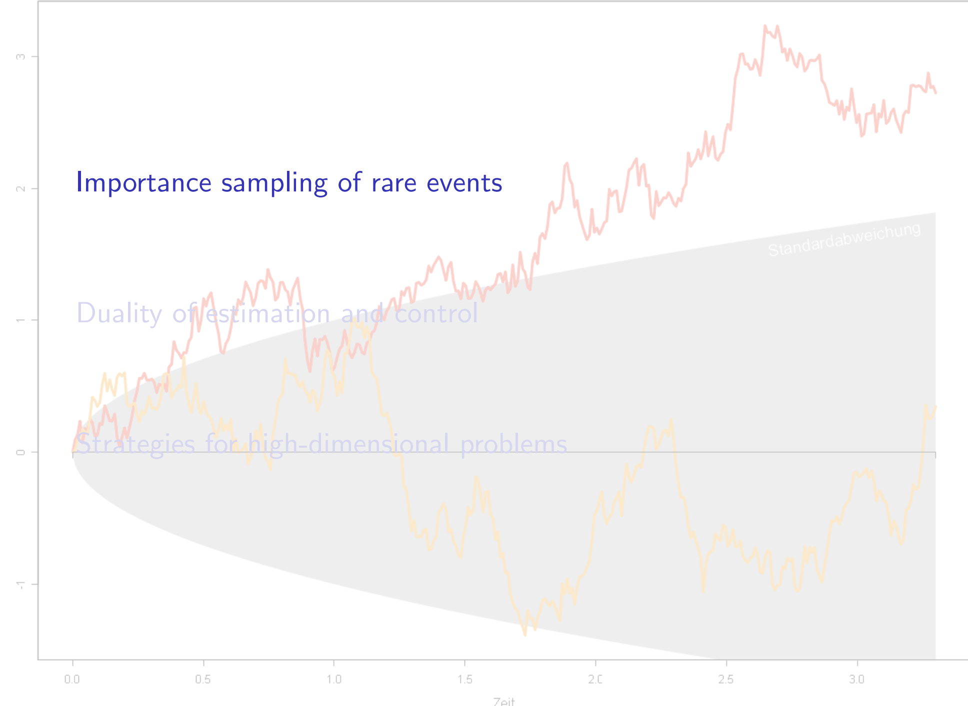

Illustrative example: bistable system

- Overdamped Langevin equation

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

- MC estimator of $p_\epsilon = P(\tau < T)$

$$\hat{p}_\epsilon^n = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}_{\{\tau_i < T\}}$$

- Small noise asymptotics (Kramers)

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

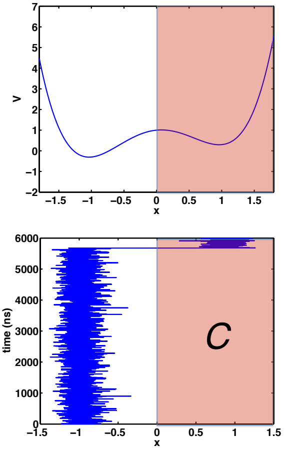

[Freidlin & Wentzell, 1984], [Berglund, Markov Processes Relat Fields 2013]

Illustrative example, cont'd

- **Relative error** of the MC estimator

$$\delta_\epsilon = \frac{\sqrt{\mathrm{Var}[\hat{p}_\epsilon^n]}}{\mathbb{E}\left[\hat{p}_\epsilon^n\right]}$$

- Varadhan's large deviations principle

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

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

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

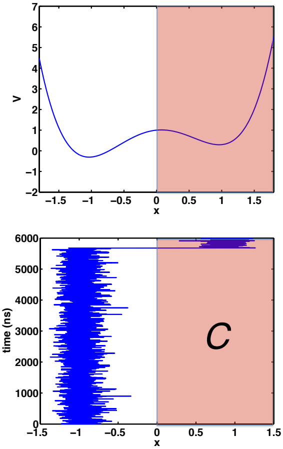

[Freidlin & Wentzell, 1984], [Berglund, Markov Processes Relat Fields 2013]

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

$$P(\tau < T) = \mathbb{E}\left[\mathbf{1}_{\{\tau<T\}}\right] \approx \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}_{\{\tau_i<T\}}\varphi^{-1}(\tau_i).$$

with τ_i now being independent draws from Q.

Optimal (zero-variance) change of measure is **infeasible**:

$$\varphi^* = \frac{dQ^*}{dP} = \frac{\mathbf{1}_{\{\tau<T\}}}{\mathbb{E}\left[\mathbf{1}_{\{\tau<T\}}\right]}\,.$$

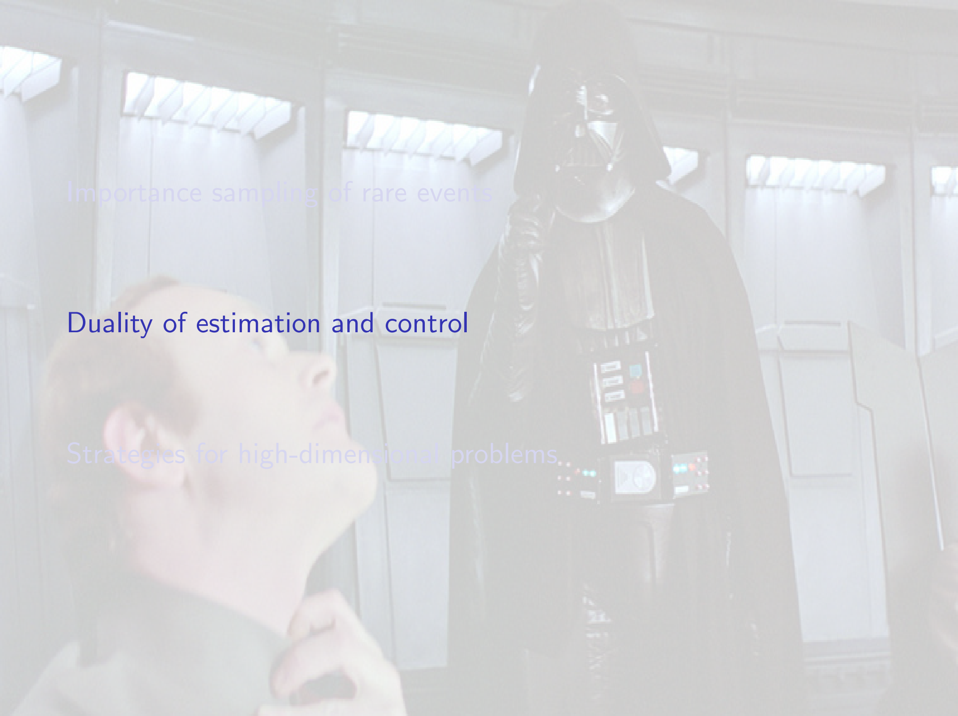

Importance sampling of rare events

Duality of estimation and control

Strategies for high-dimensional problems

Change of measure from nonequilibrium forcing

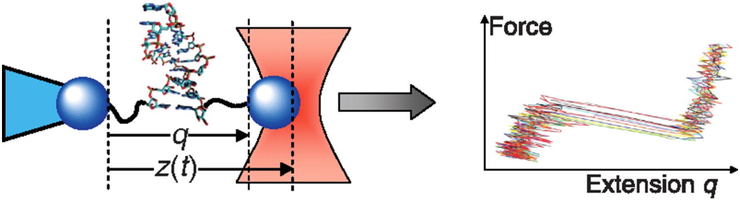

Single molecule pulling experiments, figure courtesy of G. Hummer, MPI Frankfurt

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], ...

Variational characterization of free energy

Theorem (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 $\mathbb{R}^n$,

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

and

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

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

Aim: Estimate the path functional

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

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

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

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

and a probability $Q \ll P$ on $C([0,\infty))$ with explicitly computable likelihood ratio $\varphi = dQ/dP$ (via Girsanov's Theorem).

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

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

under the **controlled dynamics**

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

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

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

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

Illustrative example, cont'd

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

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

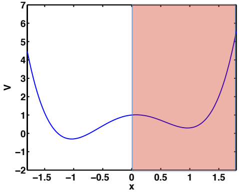

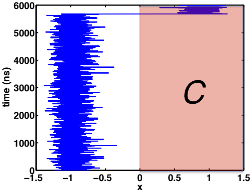

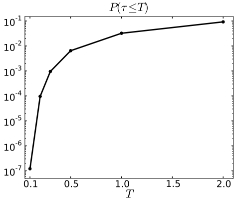

[Zhang et al, SISC, 2014], [Richter, MSc thesis, 2016], [H et al, Nonlinearity, 2016]

A few remarks

- The Theorem is a variant of the **Donsker–Varadhan** principle can be proved by both **probabilistic and PDE** arguments.
- If $\sigma\sigma^T > 0$ the optimal control has **gradient form**, i.e.

$$u_t^* = -2\sigma(X_t^{u^*})^T \nabla F(X_t^{u^*}, t)\,,$$

  with $F(x,t) = \min\{J(u)\colon X_t^u = x\}$ being the value function.
- **NFL Theorem:** $F = -\log\psi$ solves a nonlinear HJB equation,

$$-\frac{\partial F}{\partial t} + H\left(x, F, \nabla F, \nabla^2 F\right) = 0\,.$$

  (Remark: In some cases $F = F(x)$ will be stationary.)
- **Generalizations** include degenerate diffusions, Markov chains, infinite time-horizon, non-exponential functionals . . . .

[H et al, Entropy, 2014]; cf. [Fleming, SIAM J Control, 1978], [Dupuis & Wang, Stoch, 2004]

Related work (non-exhaustive)

- **Risk-sensitive control and dynamic games**: [Whittle, Eur J Oper Res, 1994], [James et al, IEEE TAC, 1994], [Dai Pra et al, Math Control Signals Systems, 1996], ...
- **Large deviations and control**: [Fleming, Appl Math Optim, 1977], [Fleming & Sheu, Ann Probab, 1997], [Pavon, Appl Math Optim, 1989], ...
- **Importance sampling of small noise diffusions**: [Dupuis & Wang, Stochastics, 2004], [Dupuis & Wang, Math Oper Res, 2007], [Vanden-Eijnden & Weare, CPAM, 2012], ...
- **Extension to multiscale systems**: [Spiliopoulos et al., SIAM MMS, 2012], [Hartmann et al, JCD, 2014], [Hartmann et al, Probab Theory Rel F, 2018], ...

Importance sampling of rare events

Duality of estimation and control

Strategies for high-dimensional problems

Two key facts about our control problem

Fact #1

Assuming that $\sigma\sigma^T > 0$ has a uniformly bounded inverse, the optimal control is a **feedback law** that can be represented as

$$u_t^* = \sigma(X_t^u) \sum_{i=1}^{\infty} c_i \nabla \phi_i(X_t^u, t)\,,$$

with coefficients $c_i \in \mathbb{R}$ and basis functions $\phi_i \in C^{1,0}(\mathbb{R}^n, [0, \infty))$.

Proof: HJB equation and Itô's formula.

Fact #2

Letting Q denote the probability (path) measure on $C([0,\infty))$ associated with the **tilted dynamics** X^u, it holds that

$$J(u) - J(u^*) = KL(Q, Q^*)$$

with $Q^* = Q(u^*)$ and

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

denoting the **relative entropy** (or: Kullback-Leibler divergence) between Q and Q^*.

Proof: Zero-variance property of $Q^* = Q(u^*)$.

Cross-entropy method for diffusions

Idea: seek a minimiser of J among all controls of the form

$$\hat{u}_t = \sigma(X_i^u) \sum_{i=1}^{M} c_i \nabla \phi_i(X_t^u, t)\,, \quad \phi_i \in (\mathbb{R}^n, [0, \infty))\,.$$

and minimise the Kullback-Leibler divergence

$$S(\mu) = KL(\mu, Q^*)$$

over all candidate probability measures of the form $\mu = \mu(\hat{u})$.

Remark: unique minimiser is $dQ^* = e^{F-W} dP = \psi^{-1} e^{-W} dP$.

[Zhang et al, SISC, 2014]; cf. [Oberhofer & Dellago, CPC, 2008]

Unfortunately, . . .

Cross-entropy method for diffusions, cont'd

...this is a nasty, non-convex minimisation problem.

Feasible cross-entropy minimisation

Minimisation of the relaxed functional $KL(Q^*, \cdot)$ is equivalent to **cross-entropy minimisation**: minimise

$$CE(\mu) = -\int \left(\log \frac{d\mu}{dP}\right) \frac{dQ^*}{dP} dP$$

over all admissible $\mu = \mu(\hat{u})$, with $dQ^* \propto e^{-W} dP$.

Note: $KL(\mu, Q^*)$=0 iff $KL(Q^*, \mu) = 0$, which holds iff $\mu = Q^*$.

Some remarks: algorithmic issues

- The cross-entropy minimisation can be recast as

$$\max_{c \in \mathbb{R}^M} \mathbb{E}\left[\log \varphi(\hat{u}) e^{-W(X^{\hat{u}})}\right]$$

  where the log likelihood ratio $\log \varphi(\hat{u})$ is **quadratic in the unknowns** $c = (c_1, \ldots, c_M)$ and can be explicitly computed.
- The **necessary optimality conditions** are of the form

$$Ac = \zeta$$

  with coefficients $A = (A_{ij})$, $\zeta = (\zeta_1, \ldots, \zeta_M)$ that are computable by Monte Carlo.
- In practice, annealing and clever choice of basis functions ϕ_i (e.g. global or local) greatly **enhances convergence**.

[Rubinstein & Kroese, Springer, 2004], [Zhang et al, SISC, 2014], [Badowski, PhD thesis, 2016]

Example I

Computing the mean first passage time ($n = 1$)

Minimise

$$J(u;\alpha) = \mathbb{E}\left[\alpha\tau^u + \frac{1}{4}\int_0^{\tau^u} |u_t|^2\, dt\right]$$

with $\tau^u = \inf\{t > 0 \colon X_t^u \in [-1.1, -1]\}$ and the dynamics

$$dX_t^u = (u_t - \nabla V(X_t^u))\, dt + 2^{-1/2}\, dB_t$$

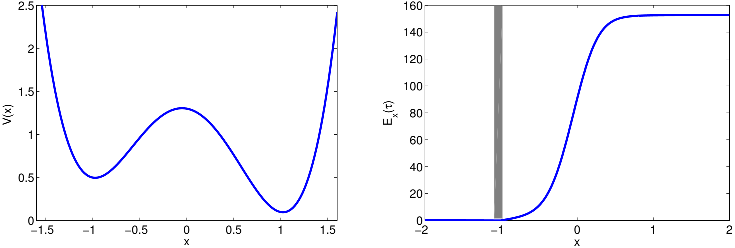

Skew double-well potential V and MFPT of the set $S = [-1.1, -1]$ from FEM reference solution).

Computing the mean first passage time, cont'd

Gradient descent approach using a parametric ansatz

$$c(x) = \sum_{i=1}^{10} c_i \nabla \phi_i(x)\,, \quad \phi_i : \text{ equispaced Gaussians}$$

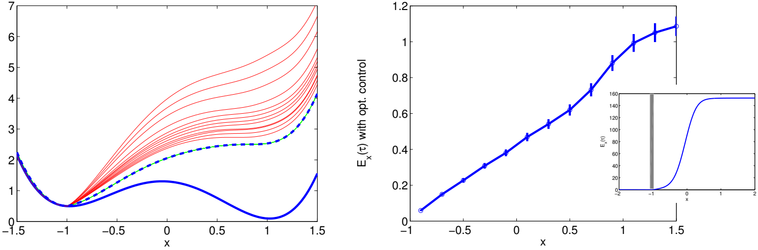

Biasing potential $V + 2F$ and unbiased estimate of the limiting MFPT.

cf. [H & Schütte, JSTAT, 2012], [Richter, MSc thesis, 2016]

Example II (suboptimal control)

Conformational transition of butane in water ($n = 16224$)

Probability of making a **gauche-trans transition** before time T:

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

with $\tau = \min\{\tau_C, T\}$ and τ_C denoting the first exit time from the gauche conformation "C" with smooth boundary ∂C

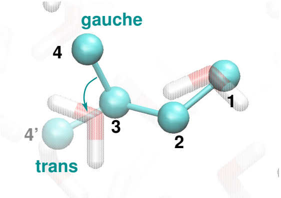

T [ps]	$\mathbf{P}(\tau \leq T)$	Error	Var	Accel. $\mathcal{I}$
0.1	4.30×10^{-5}	0.77×10^{-5}	3.53×10^{-6}	42.5
0.2	1.21×10^{-3}	0.11×10^{-3}	2.50×10^{-4}	26.0
0.5	6.85×10^{-3}	0.38×10^{-3}	2.88×10^{-3}	13.0
1.0	1.74×10^{-2}	0.08×10^{-2}	1.21×10^{-2}	7.0

IS of butane in a box of 900 water molecules (SPC/E, GROMOS force field) using **cross-entropy minimisation**

[Zhang et al, SISC, 2014], [H et al, Nonlinearity, 2016], [H et al, PTRF, 2018]

Alternatives to cross-entropy minimisation

- Minimise cost functional $J(\hat{u}(c))$ by **gradient descent:**

$$c^{(n+1)} = c^{(n)} - h_n \nabla J\left(\hat{u}\left(c^{(n)}\right)\right)\,,$$

  with $h_n \searrow 0$ as $n \to \infty$.

- Semi-explicit discretisation of FBSDE by **least-squares MC**

$$\begin{aligned} dX_s &= b(X_s, s)ds + \sigma(X_s)dB_s\,,\; X_t = x \\ dY_s &= h(X_s, Y_s, Z_s)ds + Z_s \cdot dB_s\,,\; Y_T = g(X_T)\,, \end{aligned}$$

  where $t \le s \le T$ and

$$F(x,t) = Y_t \quad \text{(as a function of the initial value } x\text{)}$$

- Approximate **policy iteration**

[H & Schütte, JSTAT, 2012], [Lie, PhD thesis, 2016], [Kebiri, Neureither & H, Proc IHP, 2018]

Take-home message (reloaded)

- Adaptive importance sampling scheme based on **dual variational formulation**; resulting control problem features short trajectories with **minimum variance estimators**.
- **Variational problem:** find the optimal perturbation by cross-entropy minimisation, gradient descent or the alike.
- Approach can (or better: should) be combined with **dimension reduction prior to optimization**.

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)