Cross-entropy minimization, optimal control and importance sampling of rare events

Carsten Hartmann, FU Berlin
(jointly with Christof Schütte & Wei Zhang)

IPAM, 19–22 January 2016

Motivation: conformation dynamics of biomolecules

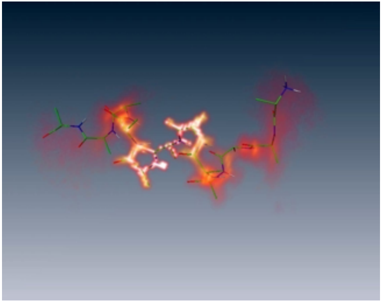

$1.3\mu s$ MD simulation of *dodeca-alanin* at $T = 300K$
(GROMOS96, visualization: Amira@ZIB)

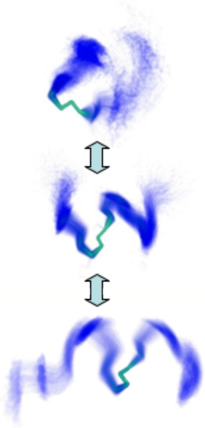

Motivation: conformation dynamics of biomolecules

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

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

with τ the time to reach the target conformation, or rates

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

where $\mathbb{E}[\cdot]$ is the expectation with respect to P.

Guiding example: bistable system

- Overdamped Langevin equation

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

- **Small noise** asymptotics for $\tau = \tau_\epsilon$

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

- Hence, for moderate values of T,

$$p_\epsilon = P(\tau < T)$$

is **exponentially small** in ϵ.

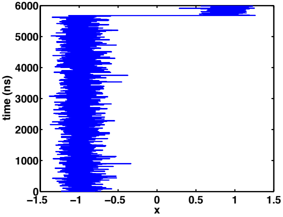

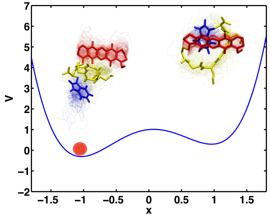

[Freidlin & Wentzell, 1984], [Berglund, Markov Processes Relat Fields 2013]

Motivation, cont'd: computational aspects

Given N **independent realizations** of $X_t = X_t^\epsilon(\omega)$, the simplest way to estimate probabilities, such as

$$p_\epsilon = P(\tau < T)$$

is by **Monte-Carlo**:

$$p_\epsilon \approx \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}_{\{\tau(\omega_i)<T\}}$$

Motivation, cont'd: computational aspects

Although the naïve MC estimator is unbiased with bounded variance $p_\epsilon(1 - p_\epsilon)/N$, the **relative error** is not:

$$\delta_{\text{rel}}(\epsilon) = \frac{\text{standard deviation}}{\text{mean}} = \frac{1}{p_\epsilon}\sqrt{\frac{p_\epsilon(1 - p_\epsilon)}{N}}$$

blows up as $p_\epsilon \to 0$.

This is a **common feature when estimating rare events**.

[Asmussen et al, Encyclopedia of Operations Research and Management Sciences, 2012]

Outline

Adaptive importance sampling of rare events

Optimal controls from cross-entropy minimization

Numerical strategies for high-dimensional problems

Adaptive importance sampling of rare events

Optimal controls from cross-entropy minimization

Numerical strategies for high-dimensional problems

Importance sampling of rare events

We can improve our estimate of p_ϵ by **sampling from an alternative distribution**, under which the event is no longer rare:

$$\begin{aligned} P(\tau < T) &= \int \mathbf{1}_{\{\tau(\omega)<T\}} dP(\omega) \\ &= \int \mathbf{1}_{\{\tau<T\}} \frac{dP}{d\mu}\, d\mu \\ &=: \mathbb{E}^\mu\left[\mathbf{1}_{\{\tau<T\}} \frac{dP}{d\mu}\right] \end{aligned}$$

Problem: Optimal (zero-variance) distribution depends on p_ϵ,

$$\mu = \frac{\mathbf{1}_{\{\tau<T\}}}{p_\epsilon} P = P(\,\cdot\,|\,\tau < T\,)\,.$$

Sampling of rare events based on large deviation asymptotics

Logarithmic asymptotic efficiency

Let Y^ϵ be any unbiased importance sampling (IS) estimator of p_ϵ. Ideally, we would like Y_ϵ to have a **bounded relative error**:

$$\limsup_{\epsilon \to 0} \frac{\mathrm{Var}^\mu[Y_\epsilon]}{p_\epsilon^2} < \infty$$

In practice, however, this is the exception.

Minimum requirement: logarithmic asymptotic efficiency

$$\lim_{n \to \infty} \frac{\log \mathbb{E}^\mu[Y_\epsilon^2]}{\log p_\epsilon} = 2\,.$$

Note that $\mathbb{E}^\mu[Y_\epsilon^2] \geq (\mathbb{E}^\mu[Y_\epsilon])^2 = p_\epsilon^2$, hence p_ϵ^2 does not decay much faster than the variance of Y_ϵ under this assumption.

Logarithmic asymptotic efficiency: large deviations

Often p_ϵ satisfies a large deviations principle of the form

$$\lim_{\epsilon \to 0} \epsilon \log p_\epsilon = -\gamma$$

for some constant $\gamma > 0$, and log efficient estimators can be based on an **exponentially tilted distribution** $\mu = \mu_\gamma$, such that

$$\lim_{\epsilon \to 0} \epsilon \log \mathbb{E}^\mu[Y_\epsilon^2] = -2\gamma$$

(Recall that $\lim_{\epsilon \to 0} \epsilon \log \mathbb{E}[\tau] = \Delta V$ in our guiding example.)

[Siegmund, Ann Stat, 1976], [Heidelberger, ACM TMCS, 1995], [Glasserman & Kou, AAP, 1997]

Logarithmic asymptotic efficiency: observations

- The large deviations principle says that

$$\mathbb{E}^{\mu}[Y_\epsilon^2] = e^{-2\gamma/\epsilon + o(1/\epsilon)}\,,$$

  so the quality of the estimator heavily depends on whether one can **control the $e^{o(1/\epsilon)}$ prefactor**.
- Extension due to Dupuis & Wang: **adaptive change of measure** based on underlying dynamic programming equation.
- But: estimators that are asymptotically log efficient may perform worse than the vanilla MC estimator when ϵ **is finite**.

[Glasserman & Wang, AAP, 1997], [Dupuis & Wang, Stochastics, 2004], [Dupuis & Wang, Math Oper Res, 2007], cf. [Vanden-Eijnden & Weare, CPAM, 2012]

Sampling of rare events based on optimal control

Single molecule experiments as paradigm

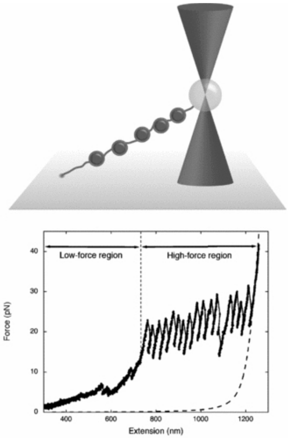

- Estimation of molecular properties in **thermodynamic equilibrium**, e.g.

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

(includes rates, statistical weights, etc.)

- Perturbation drives the system out of equilibrium with **likelihood quotient**

$$\varphi = \frac{dP}{d\mu}.$$

- Experimental and numerical realization: AFM, SMD, TMD, Metadynamics, ...

[Schlitter, J Mol Graph, 1994], [Schulten & Park, JCP, 2004], [H. et al, Proc Comput Sci, 2010]

Variational characterization of free energies

Theorem (Dai Pra et al, 1996)

For any bounded and measurable function W it holds

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

where $KL(\mu, P) \geq 0$ is the KL divergence between μ and P.

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

$$\begin{aligned} -\log \int e^{-W} dP &= -\log \int e^{-W+\log\varphi} d\mu \\ &\leq \int (W - \log\varphi)\, d\mu \end{aligned}$$

with **equality iff** $W - \log\varphi$ **is constant** (μ-a.s.).

[Fleming, SIAM J Control, 1978], [Dai Pra et al, Math Control Signals Systems, 1996]

Same same, but different. . .

Set-up: equilibrium diffusion process

Given an "equilibrium" diffusion process $X = (X_t)_{t\geq 0}$ on $\mathbb{R}^n$,

$$dX_t = b(X_t)dt + \sigma(X_t)dB_t\,, \quad X_0 = x\,,$$

we want to **estimate path functionals** of the form

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

Example: exit time statistics of a set $C \subset \mathbb{R}^n$

Let $W = \alpha\tau_C$. Then, for sufficiently small $\alpha > 0$,

$$-\alpha^{-1}\log\psi = \mathbb{E}[\tau_C] + \mathcal{O}(\alpha)$$

Set-up: nonequilibrium diffusion process

Given a "nonequilibrium" (tilted) diffusion process $X^u = (X_t^u)_{t \geq 0}$,

$$dX_t^u = (b(X_t^u) + \sigma(X_t^u)u_t)dt + \sigma(X_t^u)dB_t\,, \quad X_0^u = x\,,$$

estimate a **reweigthed version** of ψ:

$$\mathbb{E}\left[e^{-W(X)}\right] = \mathbb{E}^\mu\left[e^{-W(X^u)}\varphi(X^u)\right]$$

with equilibrium/nonequilibrium likelihood ratio $\varphi = \frac{dP}{d\mu}$.

Remark: We allow for W's of the general form

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

for suitable functions f, g and a bounded stopping time $\tau \leq T$.

Guiding example, cont'd

- **Mean first exit time** for small ϵ

$$\mathbb{E}[\tau] \asymp \exp(\Delta V/\epsilon)$$

- **Tilting the potential**

$$U(x,t) = V(x) - u_t x$$

decreases the energy barrier.

- Overdamped Langevin equation

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

with **time-dependent forcing** u_t.

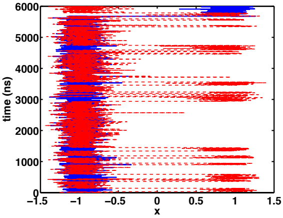

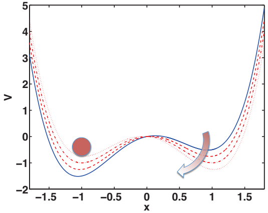

Can we systematically **speed up the sampling** while **controlling the variance** by tilting the energy landscape?

Variational characterization of free energies, cont'd

Theorem (H, 2012)

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

$$J(u) = \mathbb{E}\left[W(X^u) + \frac{1}{4}\int_0^{\tau^u} |u_s|^2\, ds\right]$$

under the nonequilibrium dynamics X_t^u, with $X_0^u = x$.

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

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

[H & Schütte, JSTAT, 2012], [H et al, Entropy, 2014]

Guiding example, cont'd

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

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

- Recovering **equilibrium statistics** by

$$\mathbb{E}[\tau_C] = \frac{d}{d\alpha}\bigg|_{\alpha=0} J(u^*)$$

- Optimally tilted potential

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

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

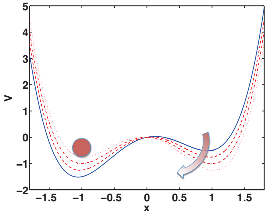

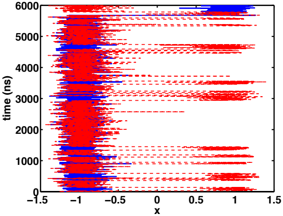

Some remarks . . .

Duality between estimation and control

The optimal control is a feedback control in gradient form ,

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

with the bias potential being the value function

$$F(x,t) = \min\{J(u) \colon X_t^u = x\}\,.$$

(Remark: In many interesting cases, $F = F(x)$ will be stationary.)

NFL Theorem: The bias potential is given by

$$F = -\log\psi,$$

i.e., u^* depends on the quantity we want to estimate.

[H & Schütte, JSTAT, 2012], [H et al, Entropy, 2014]; cf. [Fleming, SIAM J Control, 1978]

More on the duality between estimation and control

The Legendre-type variational principle for the free energy furnishes an equivalence between the **dynamic programming equation**

$$-\frac{\partial F}{\partial t} + \min_{c \in \mathbb{R}^k} \left\{ LF + (\sigma c) \cdot \nabla F + \frac{1}{2}|c|^2 + f \right\} = 0 \ + \ b.c.$$

for F and the **Feynman-Kac formula** for $e^{-F} = \mathbb{E}[e^{-W}]$:

$$\left(\frac{\partial}{\partial t} - L\right) e^{-F} = 0 \,,$$

with L being the infinitesimal generator of $X_t^{u=0}$.

Related work on risk sensitivity and large deviations: Fleming & Sheu, Whittle, James, Dupuis & Wang, Rubinstein & Kroese, Asmussen, Spiliopoulos, Vanden-Eijnden & Weare, ...

Adaptive importance sampling of rare events

Optimal controls from cross-entropy minimization

Numerical strategies for high-dimensional problems

Two key facts about our control problem

Fact #1

The optimal control is a **feedback law** of the form

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

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

Fact #2

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

$$J(u) - J(u^*) = KL(\mu, \mu^*)$$

with $\mu^* = \mu(u^*)$ and

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

the **Kullback-Leibler divergence** between μ and μ^*.

Cross-entropy method for diffusions

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

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

and minimize the Kullback-Leibler divergence

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

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

Remark: unique minimizer is given by $d\mu^* = \psi^{-1} e^{-W} dP$.

cf. [Oberhofer & Dellago, CPC, 2008], [Aurell et al, PRL, 2011]

Unfortunately, ...

Cross-entropy method for diffusions, cont'd

...that doesn't work without knowing the normalization factor ψ.

Feasible cross-entropy minimization

Minimization of the relaxed functional $KL(\mu^*, \cdot)$ is equivalent to **cross-entropy minimization**: minimize

$$CE(\mu) = -\int \log \mu \, d\mu^*$$

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

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

[Rubinstein & Kroese, Springer, 2004], [Zhang et al, SISC, 2014], [Badowski, PhD thesis, 2015]

Example I (guiding example)

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

Minimize

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

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

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

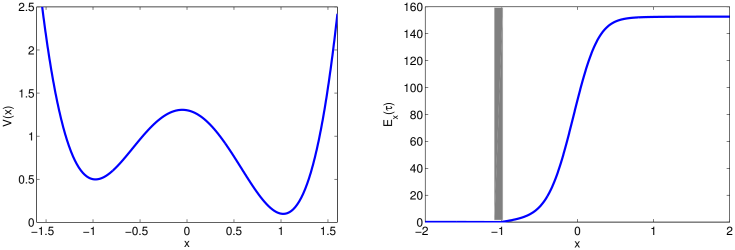

Skew double-well potential V and MFPT of the set $S = [-1.1, -1]$ (FEM reference solution).

Computing the mean first passage time, cont'd

Cross-entropy minimization using a parametric ansatz

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

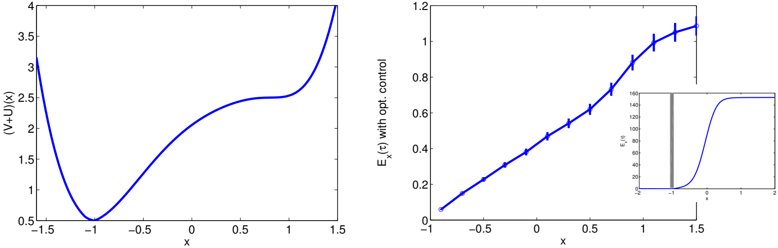

Biasing potential $V + 2F$ and unbiased estimate of the limiting MFPT.

cf. [H & Schütte, JSTAT, 2012]

Adaptive importance sampling of rare events

Optimal controls from cross-entropy minimization

Numerical strategies for high-dimensional problems

The bad news

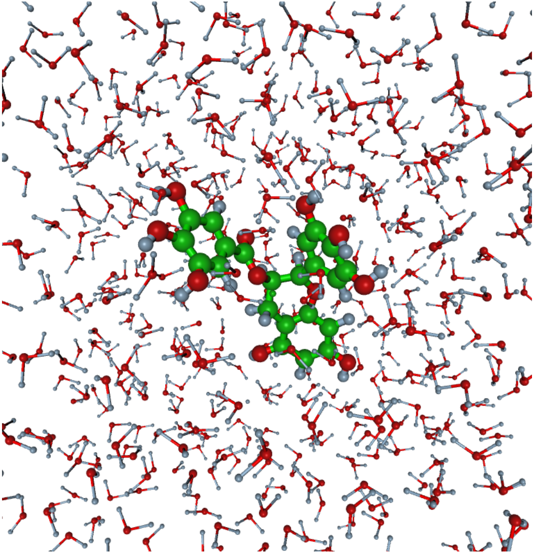

The good news

Often a “reaction coordinate” or a low-dimensional model is given.

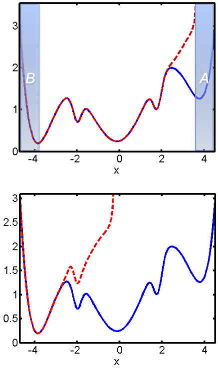

committor probabilities $P_x(\tau_B < \tau_A)$

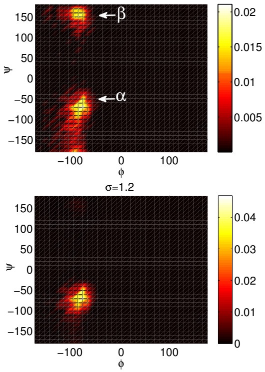

α-helical conformation of ADP in water

Cf. [Schütte et al, Math Prog, 2012], [Sarich et al, Entropy, 2014], [Banisch & H, MCRF, 2015]

Suboptimal controls from averaging

Averaged control problem: minimize

$$I(v) = \mathbb{E}\left[\bar{W}(\xi^v) + \frac{1}{4}\int_0^{\tau^v} |v_s|^2\, ds\right]$$

subject to the averaged dynamics

$$d\xi_t^u = (\Sigma(\xi_t^v)v_t - B(\xi_t^v))dt + \Sigma(\xi_t^v)dW_t$$

Control approximation strategy

$$u_t^* \approx c(\xi(X_t^{u^*}), t) = \nabla\xi(X_t^{u^*})v_t^*\,.$$

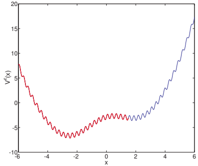

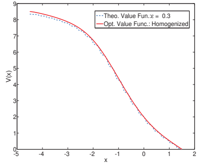

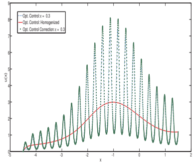

[H et al, Nonlinearity, submitted]; cf. [Legoll & Lelièvre, Nonlinearity, 2010]

Suboptimal controls, cont'd

Uniform bound of the relative error using "averaged" optimal controls

$$\delta_{\text{rel}} \leq C N^{-1/2} \eta^{1/8}, \quad \eta = \frac{\tau_{\text{fast}}}{\tau_{\text{slow}}}$$

Issues for **highly oscillatory controls**:

$$u^\eta \rightharpoonup u \quad \not\Rightarrow \quad J(u^\eta) \to J(u)$$

(relative error may be unbounded)

Log efficient estimators based on HJB subsolutions due to Spiliopoulos et al.

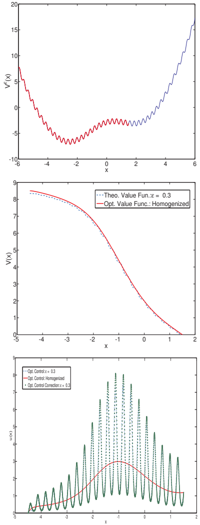

[H et al, J Comp Dyn, 2014],[H et al, submitted], cf. [Spiliopoulos et al, MMS, 2012]

Example II (suboptimal control)

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

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

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

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

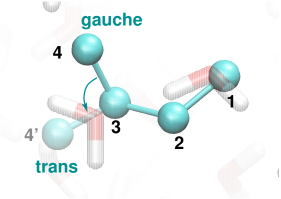

T [ps]	$\mathbf{P}(\tau \leq T)$	Error	Var	Accel. $\mathcal{I}$
0.1	4.30×10^{-5}	0.77×10^{-5}	3.53×10^{-6}	42.5
0.2	1.21×10^{-3}	0.11×10^{-3}	2.50×10^{-4}	26.0
0.5	6.85×10^{-3}	0.38×10^{-3}	2.88×10^{-3}	13.0
1.0	1.74×10^{-2}	0.08×10^{-2}	1.21×10^{-2}	7.0

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

[Zhang et al, SISC, 2014]

Take-home message

- Optimally designed nonequilibrium perturbations can **mimic thermodynamic equilibrium**.
- **Variational problem:** find the optimal perturbation by cross-entropy minimization.
- Method features short trajectories with **minimum variance estimators** of the rare event statistics.

- **To do:** adaptivity, error analysis, data-driven framework, ...

Thank you for your attention!

Acknowledgement:

Tomasz Badowski
Ralf Banisch
Han Cheng Lie
Christof Schütte
Wei Zhang

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