Séminaire de Calcul Scientifique du CERMICS

École des Ponts ParisTech

Optimal importance sampling using stochastic control

Lorenz Richter (BTU)

18 juin 2018

Optimal importance sampling using stochastic control

Lorenz Richter

Scientific computing seminar, Paris

Joint work with Carsten Hartmann

June 18, 2018

Freie Universität Berlin

SFB 1114

Brandenburgische Technische Universität Cottbus - Senftenberg

Outline

- The duality between sampling and control
  - applications: expectation values, hitting times
  - Doob's h-transform and conditioned processes
- Numerical approaches
  - gradient descent
  - controlled SMC
  - others
- Connections to large deviations
  - optimal control/variance reduction for infinite times?
  - sampling of rate functions, adaptive methods

The duality between sampling and control

- we consider the SDE

$$\mathrm{d}X_s = b(X_s)\mathrm{d}s + \sigma(X_s)\mathrm{d}W_s$$

  and want to estimate quantities of the form

$$\mathbb{E}\left[\exp(-W(X_{t:T}))\right] = \mathbb{E}\left[\exp\left(-\int_t^T f(X_s)\mathrm{d}s - g(X_T)\right)\right]$$

  with $f : \mathbb{R}^d \to \mathbb{R}, g : \mathbb{R}^d \to \mathbb{R}$
- applications:
  - sampling from stationary density π, namely $\mathbb{E}_\pi[h(X)]$
    - $b(x) = \nabla \log \pi(x), \sigma = \sqrt{2}, f = 0, g = -\log h$
  - computation of hitting times $\tau = \inf\{s \geq 0 : X_s \in B\}$
    - $T \leftrightarrow \tau, f = 1, g = 0$

The duality between sampling and control

- Monte Carlo:

$$\mathbb{E}\left[\exp(-W(X_{0:T}))\right] \approx \frac{1}{K}\sum_{k=1}^{K} \exp(-W(\hat{X}_{0:T}^{k})) =: \hat{W}_{\mathrm{MC}}$$

- possible high variance of estimators (e.g. for rare events)
- toy example: sampling from stationary density π, namely $\mathbb{E}_\pi[e^{-X}]$

$$b(x) = \nabla \log \pi(x), \sigma = \sqrt{2}, f = 0, g(x) = x$$

$$\mathrm{d}X_s = -X_s \mathrm{d}s + \sqrt{2}\mathrm{d}W_s$$

$$X_s \sim \mathcal{N}(X_0 e^{-s}, (1 - e^{-2s}))$$

  - compute $\mathbb{E}_{\mathbb{P}}[e^{-X_5}] \approx e^{0.5} \approx 1.649$
  - $\mathrm{Var}(\hat{W}_{\mathrm{MC}}) \approx 4.65$ with $K = 5000, \Delta t = 0.001$

The duality between sampling and control

- computing hitting times $\mathbb{E}[e^{\tau}]$ or $p = \mathbb{P}(\tau < T)$ in Langevin dynamics

$$\mathrm{d}X_s = -\nabla V(X_s)\mathrm{d}s + \sqrt{2\epsilon}\mathrm{d}W_s$$

- high variance since

$$\lim_{\epsilon \to 0} \epsilon \log \mathbb{E}[\tau] = \Delta V \qquad \delta_{\mathrm{rel}} = \frac{1}{p}\sqrt{\frac{p(1-p)}{K}} \xrightarrow{p \to 0} \infty$$

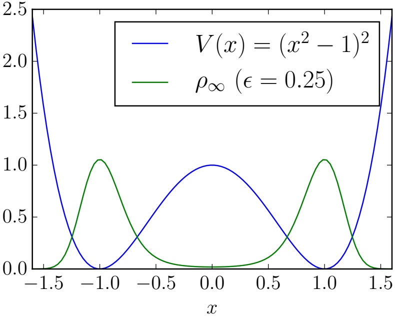

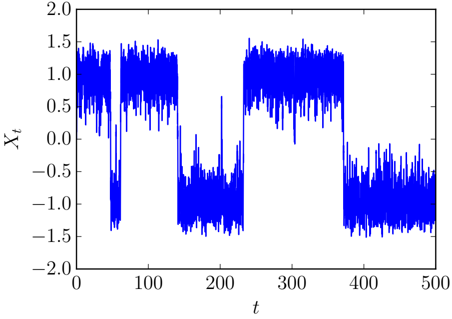

The duality between sampling and control

- strategy: importance sampling, sample from a different distribution to reduce variance

$$\mathbb{E}_{\mathbb{P}}[f(X)] = \mathbb{E}_{\mathbb{Q}}\left[f(X)\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right]$$

- $$\mathrm{d}\mathbb{Q} = \frac{f(X)}{\mathbb{E}_{\mathbb{P}}[f(X)]}\mathrm{d}\mathbb{P} \quad \Longrightarrow \quad \mathrm{Var}_{\mathbb{Q}}\left(f(X)\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right) = 0$$

The duality between sampling and control

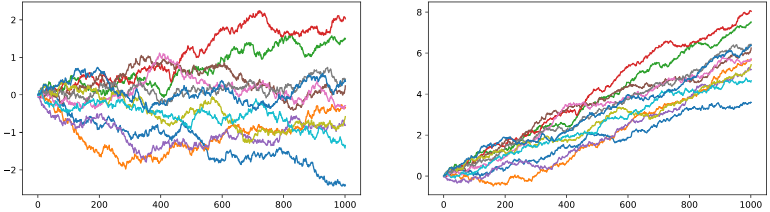

- importance sampling in path space corresponds to changing the drift of the stochastic process; the change of measure is given by Girsanov's theorem:

$$\mathrm{d}X_s = b(X_s)\mathrm{d}s + \sigma(X_s)\mathrm{d}W_s$$

$$\mathrm{d}X_s^u = (b(X_s^u) + \sigma(X_s^u)u(X_s^u, s))\mathrm{d}s + \sigma(X_s^u)\mathrm{d}W_s$$

$$\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}^u} = \exp\left(-\int_0^T u(X_s^u, s)\cdot \mathrm{d}W_s - \frac{1}{2}\int_0^T |u(X_s^u, s)|^2 \mathrm{d}s\right)$$

The duality between sampling and control

- zero-variance proposal density in path space

$$\mathrm{d}\mathbb{Q}^* = \frac{e^{-W(X_{0:T})}}{\mathbb{E}\left[e^{-W(X_{0:T})}\right]}\mathrm{d}\mathbb{P}$$

- of course this is circular since we do not know the expectation
- in path space, however, importance sampling corresponds to an optimal control problem
- we consider the SDE

$$\mathrm{d}X_s = b(X_s)\mathrm{d}s + \sigma(X_s)\mathrm{d}W_s$$

with generator

$$L = b \cdot \nabla + \frac{1}{2}\sigma\sigma^\top : \Delta$$

The duality between sampling and control

- Feynman-Kac: $(\partial_t + L - f)\psi(x,t) = 0, \psi(x, T) = e^{-g(x)}$

$$\psi(x,t) = \mathbb{E}\left[\exp\left(-\int_t^T f(X_s,s)\mathrm{d}s - g(X_T)\right)\middle| X_t = x\right]$$

- $\gamma(x,t) = \log\psi(x,t)$ brings Hamilton-Jacobi-Bellman equation (Fleming)

$$(\partial_t + L)\gamma(x,t) + \frac{1}{2}|\sigma^\top\nabla\gamma(x,t)|^2 - f = 0, \gamma(x, T) = -g(x)$$

  - $-\frac{1}{2}|\sigma^\top\nabla\gamma(x,t)|^2 = \min_{c\in\mathbb{R}^d}\left\{\sigma c\cdot\nabla\gamma(x,t) + \frac{1}{2}|c|^2\right\}$

- $\gamma(x,t)$ is the value function of a control problem with the cost functional

$$J(u) = \mathbb{E}\left[\int_t^T\left(f(X_s^u,s) + \frac{1}{2}|u(X_s^u,s)|^2\right)\mathrm{d}s + g(X_T^u)\middle| X_t = x\right]$$

The duality between sampling and control

- Donsker-Varadhan:

$$\gamma(x,t) = -\log \mathbb{E}_{\mathbb{P}}[\exp(-W)|X_t = x] = \inf_{\mathbb{Q}^u \ll \mathbb{P}} \{\mathbb{E}_{\mathbb{Q}^u}[W] + \mathrm{KL}(\mathbb{Q}^u \| \mathbb{P})\}$$

- one can show that indeed

$$\mathrm{Var}_{\mathbb{Q}^*}\left(\exp(-W)\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}^*}\right) = 0$$

- the optimal control is $u^*(x,t) = -\sigma^\top \nabla_x \gamma(x,t)$
- note that choosing a different $u(x,t)$ can increase the variance substantially

Three sides of the same problem

- free energy:

$$\gamma(x,t) = -\log \mathbb{E}[\exp(-W)|X_t = x]$$

- change of measure:

$$\gamma(x,t) = \inf_{\mathbb{Q}^u \ll \mathbb{P}} \{\mathbb{E}_{\mathbb{Q}^u}[W] + \mathrm{KL}(\mathbb{Q}^u \| \mathbb{P})\}$$

- optimal control:

$$\gamma(x,t) = \inf_u \mathbb{E}\left[\int_t^T \left(f(X_s^u, s) + \frac{1}{2}|u(X_s^u, s)|^2\right) \mathrm{d}s + g(X_T^u) \middle| X_t = x\right]$$

$$\mathrm{d}X_s^u = (b(X_s^u) + \sigma(X_s^u)u(X_s^u, s))\mathrm{d}s + \sigma(X_s^u)\mathrm{d}W_s$$

Numerics: gradient descent

- parametrize the control in ansatz functions $\varphi_i : \mathbb{R}^d \to \mathbb{R}^d$ and time-dependent coefficients $\alpha_i \in \mathbb{R}$:

$$\hat{u}(x,t) = \sum_{i=1}^{m} \alpha_i(t)\varphi_i(x)$$

- discretize time, $0 = t_1 < \cdots < t_N = T$, and run the algorithm backwards (in analogy to the dynamic programming principle)
- for each t_j compute the minimization of the costs $J(u)$ with a gradient descent in α, i.e.

$$\alpha^{k+1}(t_j) = \alpha^k(t_j) - \eta_k \nabla_\alpha \hat{J}(\hat{u}(\alpha^k(t_j)))$$

$$\hat{J}(\hat{u}) = \frac{1}{K}\sum_{k=1}^{K}\left(\sum_{i=j}^{N}\left(f(\hat{X}_i^{\hat{u},k}, i\Delta t) + \frac{1}{2}|\hat{u}(\hat{X}_i^{\hat{u},k}, i\Delta t)|^2\right)\Delta t + g(\hat{X}_N^{\hat{u},k})\right)$$

Numerics: gradient descent

- cost functional is strongly convex if ansatz functions are non-overlapping (Lie, 2016)
- the variances of different gradient estimators scale differently with the time horizon T (R, 2016):
  - $\mathrm{Var}(G_{\text{finite differences}}) \propto T$
  - $\mathrm{Var}(G_{\text{centered likelihood ratio}}) \propto T^2$
  - $\mathrm{Var}(G_{\text{likelihood ratio}}) \propto T^3$
- outlook: compute drift on the fly

Numerics: gradient descent

- sample $\mathbb{E}[\exp(-\alpha X_T)]$ with $\mathrm{d}X_t = -X_t \mathrm{d}t + \sqrt{2}\mathrm{d}W_t$
- $u^*(x,t) = -\sqrt{2}\alpha e^{t-T}$, $\hat{u}$ determined by gradient descent

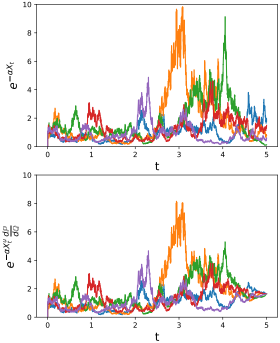

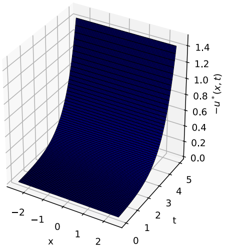

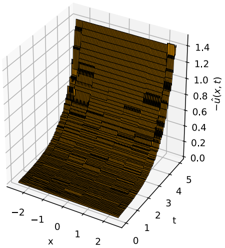

- variances for $T = 5, \alpha = 1, K = 5000$

Δt	vanilla	with u^*	with $\hat{u}$
10^{-2}	4.64	1.70×10^{-4}	5.69×10^{-2}
10^{-3}	4.02	1.69×10^{-7}	3.21×10^{-2}
10^{-4}	4.55	1.73×10^{-8}	6.45×10^{-2}

Numerical application in molecular dynamics

- let $f = 0, g(x) = -\log \mathbb{1}_B(x), \tau = \inf\{s \geq 0 : X_s \in B\}$

$$\mathbb{E}\left[e^{-g\left(X_{\min(\tau,T)}\right)}\right] = \mathbb{P}(\tau \leq T)$$

  i.e. the probability to reach B before time T (for analyzing transition mechanisms of molecules)
- conformational transitions of solvated butane, 900 water molecules, $d = 16224$ (Zhang, 2014)

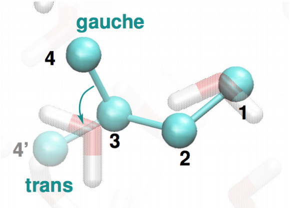

T [ps]	$\mathbf{P}(\tau \leq T)$	Error	Var	Accel. $\mathcal{I}$
0.1	4.30×10^{-5}	0.77×10^{-5}	3.53×10^{-6}	12.2
0.2	1.21×10^{-3}	0.11×10^{-3}	2.50×10^{-4}	4.8
0.5	6.85×10^{-3}	0.38×10^{-3}	2.88×10^{-3}	2.4
1.0	1.74×10^{-2}	0.08×10^{-2}	1.21×10^{-2}	1.4

Doob's *h*-transform and conditioned processes

- controlling a stochastic process also appears under the name *h*-transform
- idea: find an $h(x,t)$ s.t. for $0 \leq t < u \leq T$

$$(\partial_t + L)h(x,t) = 0 \quad \text{and} \quad \mathbb{E}[h(X_u,u)|X_t = x] = h(x,t)$$

- $h(x,t) = \mathbb{P}(X_T = \tilde{x}|X_t = x)$
  - this leads to the conditioned processes Y_s s.t. for all φ

$$\mathbb{E}[\varphi(Y_s)] = \mathbb{E}[\varphi(X_s)|X_T = \tilde{x}]$$

$$\mathrm{d}Y_t = (b(Y_s) + \sigma\sigma^\top \nabla \log h(Y_s,s))\mathrm{d}t + \sigma(Y_s)\mathrm{d}W_s$$

  with transition kernel $\mathbb{P}(Y_u = y|Y_t = x) = \frac{\mathbb{P}(X_u=y|X_t=x)h(u,y)}{h(t,x)}$

- $h(x,t) = \mathbb{E}[g(X_T)|X_t = x]$
  - this leads to the discussed zero-variance property
  - yet another application: Schrödinger bridge by taking $g(x) = \frac{\mathrm{d}\mu_T}{\mathrm{d}S_T\mu_0}(x)$, where μ_T is a desired target density (Dai Pra, 1991)

Doob's h-transform and conditioned processes

- in general the process is multidimensional and we sample $\mathbb{E}\left[\exp\left(-\int_0^T f(X_s)\mathrm{d}s - g(X_T)\right)\right], f, g : \mathbb{R}^d \to \mathbb{R}$
- only the one-dimensional quantity of interest is "conditioned", and we have some degree of freedom between the dimensions
- consider for instance $X_s \in \mathbb{R}^2, f = 0, g(x) = |x|^2$ then

$$\mathbb{E}_{\mathbb{P}}\left[\left(X_T^{(1)}\right)^2 + \left(X_T^{(2)}\right)^2\right] = \mathbb{E}_{\mathbb{Q}^u}\left[\left(X_T^{(1)}\right)^2 \frac{\mathrm{d}\,\mathbb{P}}{\mathrm{d}\,\mathbb{Q}^u}\right] + \mathbb{E}_{\mathbb{Q}^u}\left[\left(X_T^{(2)}\right)^2 \frac{\mathrm{d}\,\mathbb{P}}{\mathrm{d}\,\mathbb{Q}^u}\right]$$

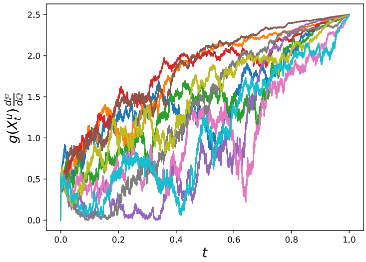

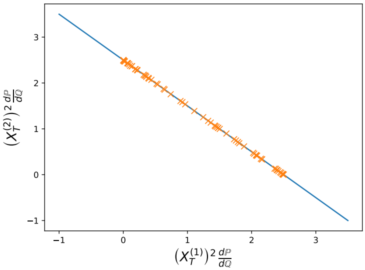

Numerics

- alternative numerical approaches are
  - cross-entropy method (Hartmann)
  - FBSDE (Hartmann, Kebiri)
  - approximate policy iteration (Bertsekas)
  - controlled SMC (Heng)

MALA and large deviations

- application: MALA

$$\mathrm{d}X_s = \nabla \log \pi(X_s)\mathrm{d}s + \sqrt{2}\mathrm{d}W_s$$

$$\sqrt{T}\bar{f}_T := \frac{1}{\sqrt{T}}\int_0^T f(X_s)\mathrm{d}s \xrightarrow{T\to\infty} \mathcal{N}(\mathbb{E}_\pi[f(X)], \sigma_f^2)$$

with asymptotic variance

$$\begin{aligned}\sigma_f^2 &= 2\langle f - \mathbb{E}_\pi[f], (-L)^{-1}(f - \mathbb{E}[f])\rangle_\pi \\ &= 2\int_0^\infty \mathbb{E}_\pi\left[f(X_s)f(X_0)\right]\mathrm{d}s\end{aligned}$$

- is there an optimal (or at least variance-reducing) control for infinite times? (ergodic control?)
- particularly relevant when considering rare events, e.g. $f(x) = \mathbb{1}_{x\in B}$
- compare to adding a nonreversible drift (Duncan, Rey-Bellet)

MALA and large deviations

- we can make statements about the large deviations of $\bar{f}_T$ from its typical value, namely

$$\lim_{T\to\infty} \frac{1}{T} \log \mathbb{P}\left(\bar{f}_T \in A\right) = -\inf_{a\in A} I(a)$$

with rate function

$$I(a) = \sup_{k\in\mathbb{R}}\{ka - \lambda_k\},$$

$$\lambda_k = \lim_{T\to\infty} \frac{1}{T} \log \mathbb{E}\left[\exp\left(kT\bar{f}_T\right)\right],$$

where λ_k is a principal eigenvalue of

$$(L + kf)r_k = \lambda_k r_k$$

MALA and large deviations

- we also have

$$\lambda_k = \lim_{T\to\infty} \sup_{\mathbb{Q}^u \ll \mathbb{P}} \left\{ \mathbb{E}_{\mathbb{Q}^u}\left[k\bar{f}\right] - \frac{1}{T}\,\mathrm{KL}(\mathbb{Q}^u|\mathbb{P}) \right\},$$

  which is the typical cost to be minimized in ergodic control problems as well as

$$\lim_{T\to\infty} \frac{\gamma(x,t)}{T} = \lambda_k$$

$$\lim_{T\to\infty} u^*(x,t) = \sigma\sigma^\top \nabla \log r_k(x)$$

MALA and large deviations

- taking the control

$$u(x) = \sigma\sigma^\top \nabla \log r_k(x)$$

  with $k = I'(a)$ makes the rare event $\bar{f}_T = a$ common in the long time limit (Touchette, 2015)
  - in analogy to Doob-conditioning define $L_k\varphi = \frac{(L+kf)r_k\varphi}{r_k} - \lambda_k\varphi$
  - the change of measure is $\frac{\mathrm{d}\mathbb{Q}_k}{\mathrm{d}\mathbb{P}} = \frac{r_k(X_T)}{r_k(x_0)} \exp(T(k\bar{f}_T - \lambda_k))$
- idea: importance sampling

$$\mathbb{P}(\bar{f}_T \in A) = \mathbb{E}_{\mathbb{Q}^u}\left[\mathbb{1}_{\bar{f}_T\in A}\frac{\mathrm{d}\,\mathbb{P}}{\mathrm{d}\,\mathbb{Q}^u}\right]$$

- but what about the numerics, how does the variance behave?

MALA and large deviations

- consider e.g. the one-dimensional Ornstein-Uhlenbeck process

$$\mathrm{d}X_s = -\gamma X_s \mathrm{d}s + \sqrt{2\epsilon}\mathrm{d}W_s$$

- we want to sample $\bar{f}_T = \frac{1}{T}\int_0^T X_s \mathrm{d}s$, then

$$r_k(x) = e^{\frac{kx}{\gamma}}, \qquad \ell_k(x) = e^{-\frac{\gamma}{2}(x-\frac{k}{\gamma^2})^2}, \qquad \lambda_k = \frac{\epsilon k^2}{\gamma^2}$$

- $u(x) = 2\epsilon \nabla \log r_k(x) = \frac{2\epsilon k}{\gamma}$
- want $\bar{f}_T = a \quad \Rightarrow \quad k = I'(a) = \frac{a\gamma^2}{2\epsilon} \quad \Rightarrow \quad u(x) = \gamma a$

$$X_s \sim \mathcal{N}\left(X_0 e^{-\gamma s} + a\left(1 - e^{-\gamma s}\right), \frac{\epsilon}{\gamma}\left(1 - e^{-2\gamma s}\right)\right) \xrightarrow{s\to\infty} \mathcal{N}\left(a, \frac{\epsilon}{\gamma}\right)$$

- but a constant drift brings a high variance of the weight
  - for a constant drift $u(x,t) = a$ we have $\mathrm{Var}_{\mathbb{Q}}\left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right) = e^{a^2 T} - 1$
  - conjecture: $\mathrm{Var}_{\mathbb{Q}}\left(g(X_T)\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right) \propto e^{a^2 T}$ (backed up by numerical simulations)
  - Girsanov weight $\exp\left(-\frac{\gamma}{\sqrt{2\epsilon}}\int_0^T a\mathrm{d}W_s - \frac{\gamma^2}{4\epsilon}\int_0^T a^2 \mathrm{d}s\right)$ degenerates (cut path into pieces?)

MALA and large deviations

- another idea: reduce the asymptotic variance by reweighting w.r.t. stationary density
- assume we want to sample $\mathbb{E}_\pi[e^{-g(X)}]$; the variance-zero (stationary) density is

$$\rho^*(x) = \frac{e^{-g(x)}\pi(x)}{\mathbb{E}_\pi[e^{-g(X)}]}$$

- we therefore consider the SDE

$$\mathrm{d}X_s = (\nabla \log \pi(X_s) - \nabla g(X_s))\mathrm{d}s + \sqrt{2}\mathrm{d}W_s$$

- we can reweight w.r.t. unnormalized stationary densities
  - (asymptotically unbiased) self-normalized IS:

$$\frac{\sum_{i=1}^N e^{-g(X_i)} \frac{\pi}{\rho^*}(X_i)}{\sum_{i=1}^N \frac{\pi}{\rho^*}(X_i)} = \frac{\sum_{i=1}^N e^{-g(X_i)} e^{g(X_i)}}{\sum_{i=1}^N e^{g(X_i)}} = \frac{N}{\sum_{i=1}^N e^{g(X_i)}}$$

  - no numerical advantage observed (shape of $e^{-g(x)}\pi(x)$ might be more complicated, possible decrease in convergence speed)

MALA and large deviations

- additional questions:
  - is there a control that optimally accelerates convergence speed (e.g. in the sense of $\mathrm{KL}(\rho_s \| \rho_\infty)$)?
  - how to combine Metropolis adjustment and Girsanov reweighting?

Numerics: controlled SMC (Heng, 2017)

- view process on a transition level with Markov kernels $M_i(x_{i-1}, \mathrm{d}x_i)$

$$\mathbb{P}(\mathrm{d}x_{0:N}) = \mu(\mathrm{d}x_0) \prod_{i=1}^{N} M_i(x_{i-1}, \mathrm{d}x_i)$$

- change of measure defined by "potential functions" $G_i(x_{i-1}, x_i)$

$$\mathbb{Q}^*(\mathrm{d}x_{0:N}) = Z^{-1} G_0(x_0) \prod_{i=1}^{N} G_i(x_{i-1}, x_i) \mathbb{P}(\mathrm{d}x_{0:N}),$$

- "twisting" by a policy ψ

$$\mathbb{Q}^\psi(\mathrm{d}x_{0:N}) = \mu^\psi(\mathrm{d}x_0) \prod_{i=1}^{N} M_i^\psi(x_{i-1}, \mathrm{d}x_i),$$

with $M_i^\psi(x_{i-1}, \mathrm{d}x_i) = \frac{\psi_i(x_{i-1}, x_i) M_i(x_{i-1}, \mathrm{d}x_i)}{\int_{\mathcal{X}} \psi_i(x_{i-1}, x_i) M_i(x_{i-1}, \mathrm{d}x_i)}$

Numerics: controlled SMC (Heng, 2017)

- for the optimal twisting ψ^* we have $G_t^{\psi^*}(x_{i-1}, x_i) = 1$ for $i \in [1 : N]$ and $G_0^{\psi^*}(x_0) = Z$ by construction, since we want $\mathbb{Q}^* = \mathbb{Q}^{\psi^*}$
- find optimal twisting ψ^* by backward recursion (in analogy to the dynamic programming principle):

$$\psi_N^*(x_{N-1}, x_N) = G_N(x_{N-1}, x_N)$$
$$\psi_i^*(x_{i-1}, x_i) = G_i(x_{i-1}, x_i) M_{i+1}(\psi_{i+1}^*)(x_i)$$
$$\psi_0^*(x_0) = G_0(x_0) M_1(\psi_1^*)(x_0)$$

- the potential functions $G_i^{\psi}(x_{i-1}, x_i)$ additionally get used as weights in a resampling scheme

References

Chetrite, R., Touchette, H.
Variational and optimal control representations of conditioned and driven processes
Journal of Statistical Mechanics: Theory and Experiment, 2015

Dai Pra, P.
A stochastic control approach to reciprocal diffusion processes
Applied Mathematics and Optimization, 23(1), 313-329, 1991

Fleming, W. H., Soner, H. M.
Controlled Markov processes and viscosity solutions.
Springer Science & Business Media, 2006

Hartmann, C., Richter, L., Schütte, C., Zhang, W.
Variational characterization of free energy: theory and algorithms
Entropy, 19(11), 626., 2017

Heng, J., Bishop, A. N., Deligiannidis, G., Doucet, A.
Controlled sequential Monte Carlo
arXiv preprint arXiv:1708.08396, 2017

Lie, H. C.
Convexity of a stochastic control functional related to importance sampling of Ito diffusions
arXiv preprint arXiv:1603.05900, 2016

Richter, L.
Efficient statistical estimation using stochastic control and optimization
Master's thesis, Freie Universität Berlin, 2016

Zhang, W., Wang, H., Hartmann, C., Weber, M., Schütte, C.
Applications of the cross-entropy method to importance sampling and optimal control of diffusions
SIAM Journal on Scientific Computing, 36(6), 2014

Merci pour votre attention!