Rare event simulation of slow-fast systems

Carsten Hartmann (BTU Cottbus-Senftenberg),

with Omar Kebiri (Cottbus), Lara Neureither (Cottbus),
Lorenz Richter (Cottbus, Berlin) & Wei Zhang (Berlin)

Geneva, 29-30th January 2020

Motivation: conformation dynamics of biomolecules

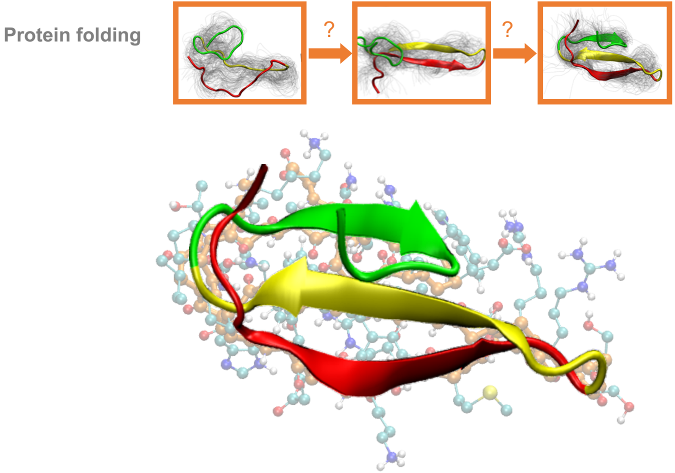

[Noé et al, PNAS, 2009]

Motivation: conformation dynamics of biomolecules

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

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

with τ some stopping time (e.g. a first passage time).

Given N **independent realizations** of X, the simplest way to estimate p is by

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

Motivation, cont'd: computational aspects

Although the naïve Monte-Carlo estimator is unbiased with bounded variance $p(1-p)/N \leq 1/(4N)$, the **relative error** is not:

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

blows up as $p \to 0$.

Remark (Varadhan's large deviations principle):

$$\frac{\mathbb{E}\left[(\hat{p}_N)^2\right]}{\left(\mathbb{E}[\hat{p}_N]\right)^2} \gg 1 \quad \text{for small } p.$$

Motivation, cont'd: importance sampling

We can improve the estimate of p by **sampling from an alternative distribution**, under which the variance becomes smaller (and the event is no longer rare):

$$P(\tau < T) = \int \mathbf{1}_{\{\tau<T\}} dP = \int \mathbf{1}_{\{\tau<T\}} \frac{dP}{dQ} \, dQ =: \mathbb{E}_Q \left[\mathbf{1}_{\{\tau<T\}} L^{-1} \right]$$

where L^{-1} in the inverse of the likelihood ratio $L = dQ/dP$ (assuming it exists).

An optimal (i.e. zero-variance) distribution Q^* **exists**, but it depends on p:

$$L^* = \frac{\mathbf{1}_{\{\tau<T\}}}{p}, \text{ i.e. } Q^* = P(\,\cdot\,|\,\tau < T\,).$$

Motivation, cont'd: importance sampling

We can improve the estimate of p by **sampling from an alternative distribution**, under which the variance becomes smaller (and the event is no longer rare):

$$P(\tau < T) = \int \mathbf{1}_{\{\tau<T\}} dP = \int \mathbf{1}_{\{\tau<T\}} \frac{dP}{dQ} dQ =: \mathbb{E}_Q\left[\mathbf{1}_{\{\tau<T\}} L^{-1}\right]$$

where L^{-1} in the inverse of the likelihood ratio $L = dQ/dP$ (assuming it exists).

An optimal (i.e. zero-variance) distribution Q^* **exists**, but it depends on p:

$$L^* = \frac{\mathbf{1}_{\{\tau<T\}}}{p}, \text{ i.e. } Q^* = P(\,\cdot\,|\,\tau < T\,).$$

Approaching minimum variance (non-exhaustive list)

- Exponential change of measures based on large deviations statistics:

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

  where γ is related to the large deviations rate function.

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

- Relative entropy (Kullback-Leibler divergence) or cross-entropy minimisation:

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

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

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

- Mean squared error and work-normalised variance minimisation

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

Outline

Duality of estimation and control

From dynamic programming to forward-backward SDE

Least squares regression

Importance sampling of multiscale systems

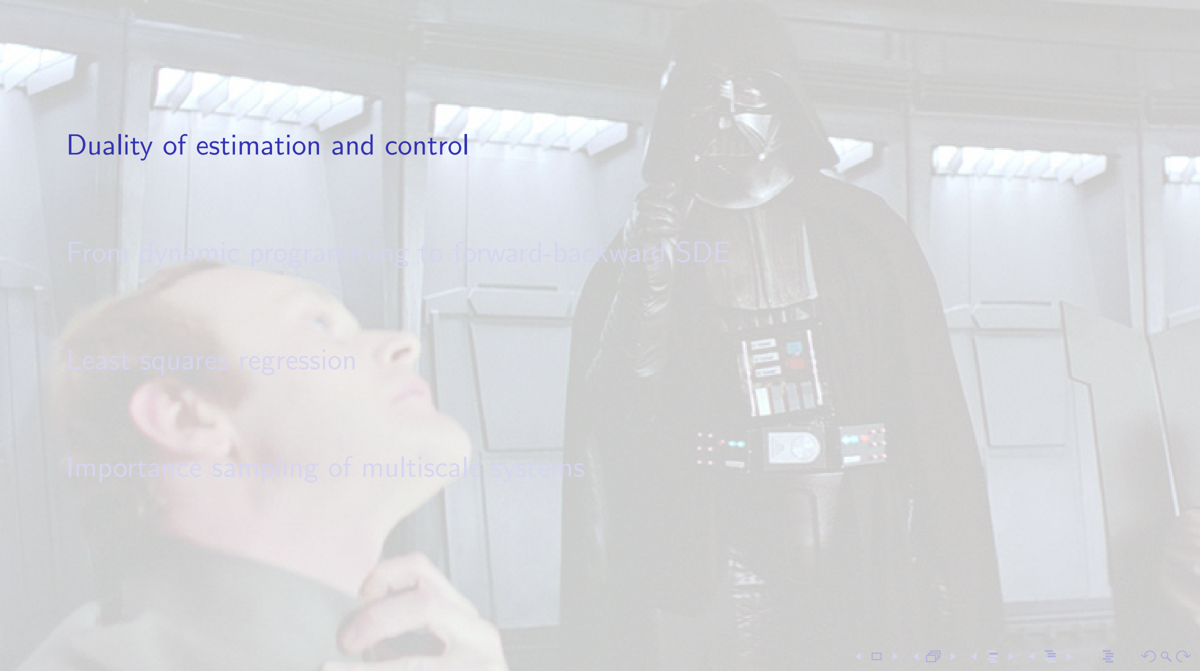

Duality of estimation and control

From dynamic programming to forward-backward SDE

Least squares regression

Importance sampling of multiscale systems

Change of measure and the Feynman-Kac theorem

Let us be specific and consider a d-**dimensional diffusion** on $[0, T]$ governed by

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

with generator $\mathcal{L}$. Then, for any function $g > 0$ **bounded away from zero**,

$$\mathbb{E}[g(X_T)] = \mathbb{E}\left[g(X_T)L^{-1}\right]\,, \quad L = \frac{g(X_T)}{\mathbb{E}[g(X_T)]}$$

Here $L = \exp(\log\psi(T, X_T) - \log\psi(0, x))$, with $\psi\colon [0, T] \times \mathbb{R}^d \to (0, \infty)$ solving

$$\left(\frac{\partial}{\partial t} + \mathcal{L}\right)\psi(t, x) = 0\,, \quad \psi(T, x) = g(x)\,.$$

Change of measure and the Feynman-Kac theorem

Let us be specific and consider a d-**dimensional diffusion** on $[0, T]$ governed by

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

with generator $\mathcal{L}$. Then, for any function $g > 0$ **bounded away from zero**,

$$\mathbb{E}[g(X_T)] = \mathbb{E}\left[g(X_T)L^{-1}\right]\,, \quad L = \frac{g(X_T)}{\mathbb{E}[g(X_T)]}$$

Here $L = \exp(\log\psi(T, X_T) - \log\psi(0, x))$, with $\psi\colon [0, T] \times \mathbb{R}^d \to (0, \infty)$ solving

$$\left(\frac{\partial}{\partial t} + \mathcal{L}\right)\psi(t, x) = 0\,, \quad \psi(T, x) = g(x)\,.$$

Change of measure and the Feynman-Kac theorem

Let us be specific and consider a d-**dimensional diffusion** on $[0, T]$ governed by

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

with generator $\mathcal{L}$. Then, for any function $g > 0$ **bounded away from zero**,

$$\mathbb{E}[g(X_T)] = \mathbb{E}\left[g(X_T)L^{-1}\right]\,, \quad L = \frac{g(X_T)}{\mathbb{E}[g(X_T)]}$$

Here $L = \exp(\log\psi(T, X_T) - \log\psi(0, x))$, with $\psi\colon [0, T] \times \mathbb{R}^d \to (0, \infty)$ solving

$$\left(\frac{\partial}{\partial t} + \mathcal{L}\right)\psi(t, x) = 0\,, \quad \psi(T, x) = g(x)\,.$$

Change of measure and the Feynman-Kac theorem, cont'd

By construction, $L > 0$ defines a **zero-variance change of measure** via $L = dQ^*/dP$.

Now, using Itô's formula, it follows that

$$L^{-1} = \exp\left(-\int_0^T u_t^* \cdot dW_t + \frac{1}{2}|u_t^*|^2 dt\right),$$

with $u_t^* = \sigma(X_t)^T \nabla \log \psi(t, X_t)$.

By Girsanov's Thm, Q^* is generated by the previous SDE with **new drift** $b^u = b + \sigma u^*$.

Change of measure and the Feynman-Kac theorem, cont'd

By construction, $L > 0$ defines a **zero-variance change of measure** via $L = dQ^*/dP$.

Now, using Itô's formula, it follows that

$$L^{-1} = \exp\left(-\int_0^T u_t^* \cdot dW_t + \frac{1}{2}|u_t^*|^2 dt\right),$$

with $u_t^* = \sigma(X_t)^T \nabla \log \psi(t, X_t)$.

By Girsanov's Thm, Q^* is generated by the previous SDE with **new drift** $b^u = b + \sigma u^*$.

Computational aspects: some observations

- We cannot draw directly from Q^*, because we **cannot simulate** X^* as it depends on the unknown quantity ψ via the extra drift term $u_t^* = \sigma(X_t)^T \nabla \log \psi(t, X_t)$.
- The extra drift u^* **minimises the second moment** of the importance sampling estimator (and hence the variance),

$$u^* = \operatorname*{argmin}_u \mathbb{E}\left[(g(X_T))^2 \exp\left(-\int_0^T u_t \cdot dW_t + \frac{1}{2}|u_t|^2 dt\right)\right],$$

but computing it e.g. by stochastic gradient descent is notoriously difficult.
- So we have replaced a difficult rare event estimation problem by a **potentially more difficult** PDE numerics or variational problem.

[Bardou, PhD Thesis, 2005], [H & Schütte, JSTAT, 2012]

Computational aspects: some observations

- We cannot draw directly from Q^*, because we **cannot simulate** X^* as it depends on the unknown quantity ψ via the extra drift term $u_t^* = \sigma(X_t)^T \nabla \log \psi(t, X_t)$.
- The extra drift u^* **minimises the second moment** of the importance sampling estimator (and hence the variance),

$$u^* = \operatorname*{argmin}_u \mathbb{E}\left[(g(X_T))^2 \exp\left(-\int_0^T u_t \cdot dW_t + \frac{1}{2}|u_t|^2 dt\right)\right],$$

but computing it e.g. by stochastic gradient descent is notoriously difficult.
- So we have replaced a difficult rare event estimation problem by a **potentially more difficult** PDE numerics or variational problem.

[Bardou, PhD Thesis, 2005], [H & Schütte, JSTAT, 2012]

Computational aspects: some observations

- We cannot draw directly from Q^*, because we **cannot simulate** X^* as it depends on the unknown quantity ψ via the extra drift term $u_t^* = \sigma(X_t)^T \nabla \log \psi(t, X_t)$.
- The extra drift u^* **minimises the second moment** of the importance sampling estimator (and hence the variance),

$$u^* = \underset{u}{\operatorname{argmin}} \, \mathbb{E}\left[(g(X_T))^2 \exp\left(- \int_0^T u_t \cdot dW_t + \frac{1}{2} |u_t|^2 dt \right) \right],$$

but computing it e.g. by stochastic gradient descent is notoriously difficult.
- So we have replaced a difficult rare event estimation problem by a **potentially more difficult** PDE numerics or variational problem.

[Bardou, PhD Thesis, 2005], [H & Schütte, JSTAT, 2012]

Computational aspects: more observations

- $-\log\psi$ is the value function of the **linear-quadratic optimal control problem**

$$-\log\psi(t,x) = \min_u \mathbb{E}\left[\frac{1}{2}\int_t^T |u_s|^2 ds - \log g(X_T^u) \,\middle|\, X_t^u = x\right]$$

under the controlled dynamics

$$dX_t^u = (b(X_t^u) + \sigma(X_t^u)u_t)\,dt + \sigma(X_t^u)dW_t\,.$$

- The necessary and sufficient condition for optimality is that $v = -\log\psi$ is a (sufficiently regular) solution of the semilinear **dynamic programming equation**

$$\left(\frac{\partial}{\partial t} + \mathcal{L}\right)v + h(x, v, \sigma^T\nabla v) = 0\,, \quad v(T,x) = g(x)$$

with the nonlinearity $h(x,y,z) = \min_\alpha\left\{\alpha\cdot z + \frac{1}{2}|\alpha|^2\right\} = -\frac{1}{2}|z|^2$.

[Fleming & Soner, 2006]

Computational aspects: more observations

- $-\log\psi$ is the value function of the **linear-quadratic optimal control problem**

$$-\log\psi(t,x) = \min_u \mathbb{E}\left[\frac{1}{2}\int_t^T |u_s|^2 ds - \log g(X_T^u) \,\middle|\, X_t^u = x\right]$$

  under the controlled dynamics

$$dX_t^u = (b(X_t^u) + \sigma(X_t^u)u_t)\,dt + \sigma(X_t^u)dW_t\,.$$

- The necessary and sufficient condition for optimality is that $v = -\log\psi$ is a (sufficiently regular) solution of the semilinear **dynamic programming equation**

$$\left(\frac{\partial}{\partial t} + \mathcal{L}\right) v + h(x, v, \sigma^T \nabla v) = 0\,, \quad v(T,x) = g(x)$$

  with the nonlinearity $h(x,y,z) = \min_\alpha \left\{\alpha\cdot z + \frac{1}{2}|\alpha|^2\right\} = -\frac{1}{2}|z|^2$.

[Fleming & Soner, 2006]

Illustrative example (with a slight abuse of the previous formalism)

- Hitting probability: $g = \mathbf{1}_{\tau < T}$:

$$-\log P(\tau < T) = \min_u \mathbb{E}\left[\frac{1}{2}\int_0^{\tau\wedge T} |u_s|^2 ds - \log \mathbf{1}_{\tau<T}\right]$$

under the **tilted dynamics**

$$dX_t^u = (u_t - \nabla V(X_t^u))\, dt + dW_t$$

- Optimally tilted potential

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

with **non-stationary** feedback $u_t^* = c(t, X_t^*)$.

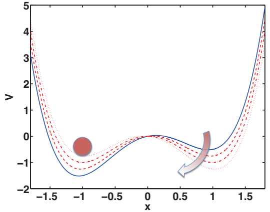

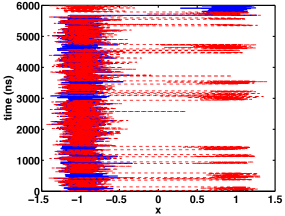

[H et al, JSTAT, 2012], [H. et al, Entropy, 2014]

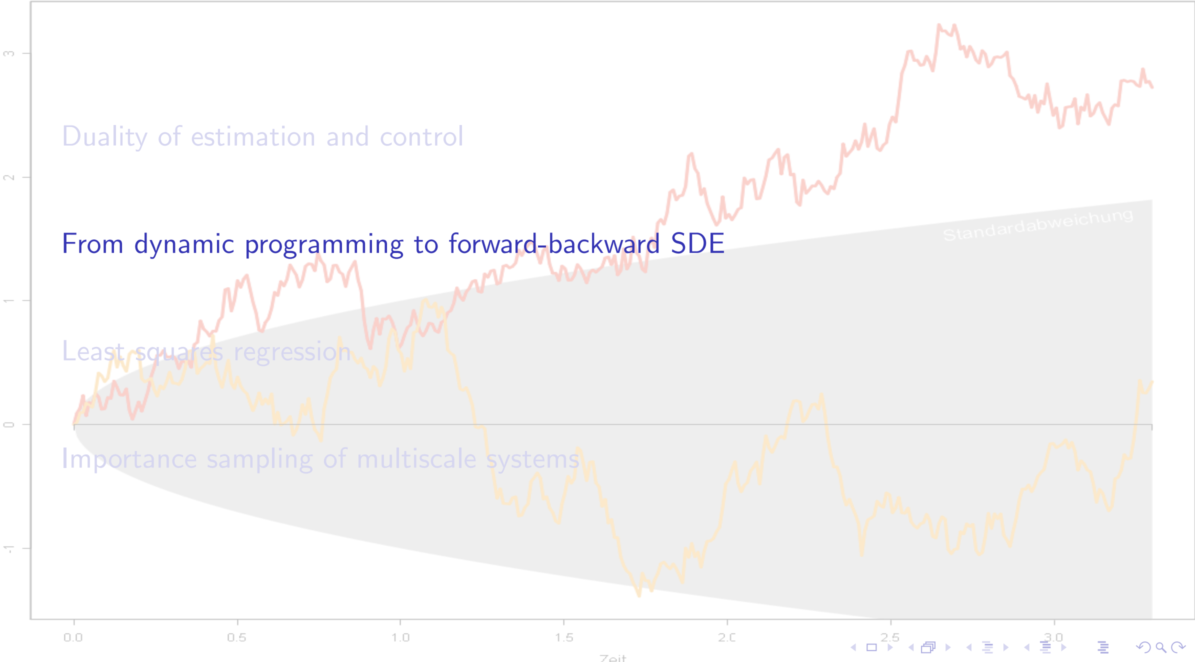

Duality of estimation and control

From dynamic programming to forward-backward SDE

Least squares regression

Importance sampling of multiscale systems

A wish list

Solving either the dynamic programming or the Feynman-Kac PDE is **not an option**! Hence we seek a **reformulation of the problem**, such that the problem

1. remains numerically tractable in high dimensions,
2. is amenable to model reduction when multiple scales are present,
3. can be solved iteratively by adaptively improving the control,
4. does not require too much expert-knowledge (e.g. specific basis functions).

From dynamic programming to a pair of SDE

The **semilinear dynamic programming equation**

$$\left(\frac{\partial}{\partial t} + \mathcal{L}\right) v + h(x, v, \sigma^T \nabla v) = 0\,, \quad v(T, x) = g(x)$$

is equivalent to the uncoupled **forward-backward SDE (FBSDE)**

$$dX_s = b(X_s)ds + \sigma(X_s)dW_s\,,\ X_t = x$$
$$dY_s = -h(X_s, Y_s, Z_s)ds + Z_s \cdot dW_s\,,\ Y_T = g(X_T)\,,$$

in dimension $d + 1$ and on a finite time horizon $[0, T]$ where

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

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

From dynamic programming to a pair of SDE: sketch of derivation

By Itô's Lemma and the dynamic programming PDE, we have for $s \in (t, T)$:

$$
\begin{aligned}
dY_s &= \left(\frac{\partial}{\partial t} + \mathcal{L}\right) v(s, X_s) + \nabla v(s, X_s) \cdot \sigma(X_s) dW_s \\
&= -h(X_s, Y_s, Z_s) + Z_s \cdot dW_s \,.
\end{aligned}
$$

where $\mathcal{L}$ is the **generator of the uncontrolled SDE**.

For $s = T$, the process Y satisfies the **terminal condition**

$$
Y_T = v(T, X_T) = g(X_T) \,.
$$

Remark

The solution to an FBSDE is a **triplet** (X, Y, Z) where (Y_s, Z_s) is adapted to the filtration generated by $(X_u)_{u \leq s}$. Consequently, $Y_t = v(t, x)$ is **deterministic**.

From dynamic programming to a pair of SDE: sketch of derivation

By Itô's Lemma and the dynamic programming PDE, we have for $s \in (t, T)$:

$$
\begin{aligned}
dY_s &= \left(\frac{\partial}{\partial t} + \mathcal{L}\right) v(s, X_s) + \nabla v(s, X_s) \cdot \sigma(X_s) dW_s \\
&= -h(X_s, Y_s, Z_s) + Z_s \cdot dW_s \, .
\end{aligned}
$$

where $\mathcal{L}$ is the **generator of the uncontrolled SDE**.

For $s = T$, the process Y satisfies the **terminal condition**

$$
Y_T = v(T, X_T) = g(X_T) \, .
$$

Remark

The solution to an FBSDE is a **triplet** (X, Y, Z) where (Y_s, Z_s) is adapted to the filtration generated by $(X_u)_{u \leq s}$. Consequently, $Y_t = v(t, x)$ is **deterministic**.

Remark

A BSDE is **not a time-reversed SDE** in the sense of $Y_t = f(X_{T-t})$: the FBSDE

$$dX_s = dW_s \,, \quad dY_s = Z_s \cdot dW_s \,,$$

with terminal condition $Y_T = X_T$ has **two possible formal solutions**

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

but only one of them is adapted.

This observation has **two important consequences**: (a) for the numerics and (b) for the multiscale analysis of our optimal control problem.

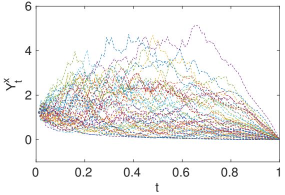

Numerical discretisation of FBSDE

The **FBSDE is decoupled** and an explicit time stepping 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

$$\hat{X}_0 = x\,, \quad \hat{Y}_N = g(\hat{X}_N)$$

Solution to stochastic two-point boundary value problem:

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

Solution by least-squares Monte-Carlo

Numerical discretisation of FBSDE

Euler-discretised FBSDE:

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

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

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

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

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

Numerical discretisation of FBSDE, cont'd

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**:

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

Specifically,

$$\hat{Y}_n \approx \underset{Y = Y_K(\hat{X}_n)}{argmin} \frac{1}{M} \sum_{m=1}^{M} \left| Y - \hat{Y}_{n+1}^{(m)} - \Delta t\, h\big(\hat{X}_n^{(m)}, \hat{Y}_{n+1}^{(m)}, \hat{Z}_{n+1}^{(m)}\big) \right|^2,$$

where $Y_K(x) = \alpha_1 \phi_1(x) + \ldots + \alpha_K \phi_K(x)$ is a parametric representation of Y.

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

Numerical discretisation of FBSDE, cont'd

The conditional expectation

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

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

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

Specifically,

$$\hat{Y}_n \approx \operatorname*{\mathit{argmin}}_{Y = Y_K(\hat{X}_n)} \frac{1}{M} \sum_{m=1}^{M} \left| Y - \hat{Y}_{n+1}^{(m)} - \Delta t\, h\left(\hat{X}_n^{(m)}, \hat{Y}_{n+1}^{(m)}, \hat{Z}_{n+1}^{(m)}\right) \right|^2 ,$$

where $Y_K(x) = \alpha_1 \phi_1(x) + \ldots + \alpha_K \phi_K(x)$ is a parametric representation of Y.

[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

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

The corresponding **loss function** is given by

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

(Note that $\mathbb{E}\left[|Y_T - g(X_T)|^2\right] = 0$ for the exact solution.)

[E et al, Commun Math Stat, 2017], [H. et al, Chaos, 2019], [Pham et al. Meth. Comput. Appl. Prob., 2020]

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

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

The corresponding **loss function** is given by

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

(Note that $\mathbb{E}\left[|Y_T - g(X_T)|^2\right] = 0$ for the exact solution.)

[E et al, Commun Math Stat, 2017], [H. et al, Chaos, 2019], [Pham et al. Meth. Comput. Appl. Prob., 2020]

Remark: iterative computation of the optimal control

Both LSMC and DL methods can be improved by **iterative learning of optimal control** based on the FBSDE

$$dX_s = (b(X_s) + \sigma(X_s)\xi_s)\,ds + \sigma(X_s)dW_s\,,\ X_0 = x$$
$$dY_s = (Z_s \cdot \xi_s - h(X_s, Y_s, Z_s))\,ds + Z_s \cdot dB_s\,,\ \ Y_T = g(X_T)$$

that, for any measurable ξ represents the same PDE.

Observations

- variance at most MC variance
- family of zero-variance estimators
- iteration may not converge

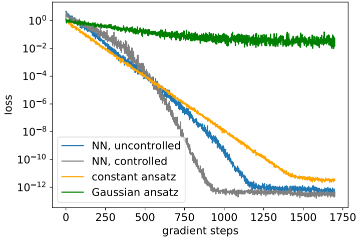

[Kebiri & H, Computation, 2018], [Kebiri et al, Proc. IHP, 2019], [H et al, Chaos, 2019]

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

$$\frac{dQ}{dP}\bigg|_{\mathcal{F}_T} = \exp\left(\int_0^T Z_s \cdot dW_s - \frac{1}{2}\int_0^T |Z_s|^2\, ds\right),$$

  for $T < \infty$ (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** random time horizon, singular terminal condition, ...

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

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

for $T < \infty$ (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** random time horizon, singular terminal condition, ...

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

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

  for $T < \infty$ (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** random time horizon, singular terminal condition, ...

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

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

  for $T < \infty$ (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** random time horizon, singular terminal condition, ...

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

Duality of estimation and control

From dynamic programming to forward-backward SDE

Least squares regression

Importance sampling of multiscale systems

Example I: hitting probabilities

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

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

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

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

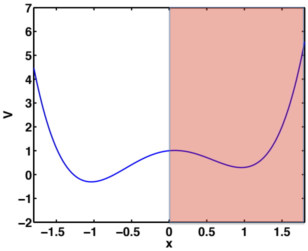

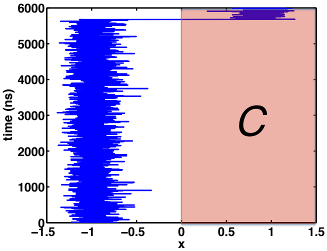

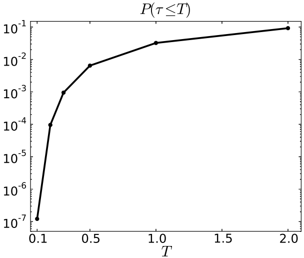

[Zhang et al, SISC, 2014], [Richter, MSc thesis, 2016], [H et al, Nonlinearity, 2016]

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 d-sphere of radius r:

$$\tau = \inf\{t > 0 \colon x + W_t \notin S_r^d\}$$

Cumulant generating function of first exit time satisfies

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

- Least-squares MC w/ $K = 3, M \sim 10^2$

	$d = 3$	$d = 10$	$d = 100$	$d = 1000$
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

- mean first exit time $\mathbb{E}[\tau] = \frac{r^2 - |x|^2}{d}$

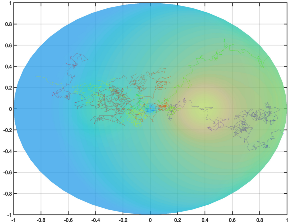

[H, et al., Chaos, 2019]

Suboptimal controls for multiscale problems

Suboptimal controls from averaging

The fact that the FBSDE is uncoupled implies that every strong approximation X gives rise to an approximation of (Y, Z).

Averaged control problem: minimize

$$J(\eta) = \mathbb{E}\left[\frac{1}{2}\int_0^T |\eta_s|^2\, ds + \bar{g}(x_T)\right]$$

subject to the averaged dynamics

$$dx_t^\eta = (\Sigma(x_t^\eta)\eta_t + B(x_t^\eta))dt + \Sigma(x_t^\eta)dW_t$$

Control approximation strategy when $x = \xi(X)$

$$u_t^* \approx \nabla\xi(X_t^*)\eta_t^*\,.$$

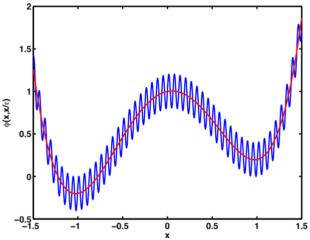

[H et al, Nonlinearity, 2016],[H. et al, PTRF, 2018]; cf. [Legoll & Lelièvre, Nonlinearity, 2010], [H. et al]

Slow-fast systems: some results

- Uniform bound of the relative error using averaged optimal controls

$$\delta_{\mathrm{rel}} \leq CN^{-1/2}\,\varepsilon^{1/8}\,, \quad \varepsilon = \frac{\tau_{\mathrm{fast}}}{\tau_{\mathrm{slow}}}$$

- Slightly stronger error bound for limit BSDE

$$\sup\{|Y_t^\delta - \bar{Y}_t| \colon 0 \leq t \leq T\} \leq C\sqrt{\varepsilon}$$

  as $\delta \to 0$, analogously for Z_t^δ (implies importance sampling $\mathcal{O}(\varepsilon^{1/4})$ error bound).
- Issues for **highly oscillatory controls** due to quadratic nonlinearity. Log efficiency in this case has been proved by Dupuis, Spiliopoulos and Wang.

[Spiliopoulos et al, MMS, 2012], [Banisch & H, MCRF, 2016], [H et al, PTRF, 2018], [Kebiri & H, Computation, 2018]

Conclusions & outlook

- Adaptive importance sampling scheme based on **dual stochastic control formulation** features short trajectories with **minimum variance estimators**.
- Optimal control problem boils down to an **uncoupled FBSDE** with only one additional spatial dimension.
- **Error analysis** of the FBSDE algorithms for unbounded stopping time & singular terminal condition is largely open. **LSMC algorithm** requires some fine-tuning.
- Numerics should be combined with **dimension reduction** (e.g. averaging).

Conclusions & outlook

- Adaptive importance sampling scheme based on **dual stochastic control formulation** features short trajectories with **minimum variance estimators**.
- Optimal control problem boils down to an **uncoupled FBSDE** with only one additional spatial dimension.
- **Error analysis** of the FBSDE algorithms for unbounded stopping time & singular terminal condition is largely open. **LSMC algorithm** requires some fine-tuning.
- Numerics should be combined with **dimension reduction** (e.g. averaging).

Conclusions & outlook

- Adaptive importance sampling scheme based on **dual stochastic control formulation** features short trajectories with **minimum variance estimators**.
- Optimal control problem boils down to an **uncoupled FBSDE** with only one additional spatial dimension.
- **Error analysis** of the FBSDE algorithms for unbounded stopping time & singular terminal condition is largely open. **LSMC algorithm** requires some fine-tuning.
- Numerics should be combined with **dimension reduction** (e.g. averaging).

Conclusions & outlook

- Adaptive importance sampling scheme based on **dual stochastic control formulation** features short trajectories with **minimum variance estimators**.
- Optimal control problem boils down to an **uncoupled FBSDE** with only one additional spatial dimension.
- **Error analysis** of the FBSDE algorithms for unbounded stopping time & singular terminal condition is largely open. **LSMC algorithm** requires some fine-tuning.
- Numerics should be combined with **dimension reduction** (e.g. averaging).

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)