Solving high-dimensional Hamilton-Jacobi PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space

Nikolas Nüsken, **Lorenz Richter**

Freie Universität Berlin, lorenz.richter@fu-berlin.de

Controlled diffusions

Let us look at SDEs of the form

$$\mathrm{d}X_s = b(X_s, s)\,\mathrm{d}s + \sigma(X_s, s)\,\mathrm{d}W_s, \qquad X_0 = x_0, \tag{1}$$

and their controlled counterparts

$$\mathrm{d}X_s^u = (b(X_s^u, s) + \sigma(X_s^u, s)u(X_s^u, s))\,\mathrm{d}s + \sigma(X_s^u, s)\,\mathrm{d}W_s, \qquad X_t^u = x_0. \tag{2}$$

The goal is to find $u \in \mathcal{U}$ that steers the dynamics in a "good" way.

Equivalent problems

The following problems are (more or less) equivalent:

Problem 1 (Optimal control) *Find $u^* \in \mathcal{U}$ such that*

$$J(u^*) = \inf_{u \in \mathcal{U}} J(u), \tag{3}$$

where

$$J(u) = \mathbb{E}\left[\int_0^T \left(f(X_s^u, s) + \frac{1}{2}|u(X_s^u, s)|^2\right)\mathrm{d}s + g(X_T^u)\right]. \tag{4}$$

Problem 2 (Hamilton-Jacobi-Bellman PDE) *Find a solution V to the PDE*

$$(L + \partial_t)V(x,t) - \frac{1}{2}|\sigma^\top \nabla V(x,t)|^2 + f(x,t) = 0, \qquad (x,t) \in \mathbb{R}^d \times [0,T), \tag{5a}$$

$$V(x,T) = g(x), \qquad x \in \mathbb{R}^d, \tag{5b}$$

where

$$L = \frac{1}{2}\sum_{i,j=1}^d (\sigma\sigma^T)_{ij}(x,t)\partial_{x_i}\partial_{x_j} + \sum_{i=1}^d b_i(x,t)\partial_{x_i}. \tag{6}$$

Problem 3 (Forward-backward SDE) *Find progressively measurable stochastic processes $Y : \Omega \times \mathbb{R}^d \times [0,T] \to \mathbb{R}$ and $Z : \Omega \times \mathbb{R}^d \times [0,T] \to \mathbb{R}^d$ such that*

$$\mathrm{d}X_s = b(X_s, s)\,\mathrm{d}s + \sigma(X_s, s)\,\mathrm{d}W_s, \qquad X_0 = x_0, \tag{7a}$$

$$\mathrm{d}Y_s = -f(X_s, s)\,\mathrm{d}s + \frac{1}{2}|Z_s|^2\,\mathrm{d}s + Z_s \cdot \mathrm{d}W_s, \qquad Y_T = g(X_T). \tag{7b}$$

Problem 4 (Conditioning) *Denote by $\mathbb{P}$ and $\mathbb{P}^u$ the path measures associated to the solutions of (1) and (2), respectively. Define $\mathbb{Q}$ via*

$$\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} = \frac{e^{-\mathcal{W}}}{\mathcal{Z}}, \qquad \mathcal{Z} = \mathbb{E}\left[\exp(-\mathcal{W}(X))\right], \tag{8}$$

where

$$\mathcal{W}(X) = \int_0^T f(X_s, s)\,\mathrm{d}s + g(X_T). \tag{9}$$

Find $u^* \in \mathcal{U}$ such that $\mathbb{P}^{u^*} = \mathbb{Q}$.

Problem 5 (Importance sampling) *It holds that*

$$\mathbb{E}\left[\exp(-\mathcal{W}(X))\right] = \mathbb{E}\left[\exp(-\mathcal{W}(X^u))\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{P}^u}\right], \tag{10}$$

for all $u \in \mathcal{U}$. Find $u^* \in \mathcal{U}$ such that

$$\mathrm{Var}\left(\exp(-\mathcal{W}(X^{u^*}))\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{P}^{u^*}}\right) = \inf_{u \in \mathcal{U}} \mathrm{Var}\left(\exp(-\mathcal{W}(X^u))\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{P}^u}\right). \tag{11}$$

Connections

The solutions to the problems above coincide,

$$Y_s = V(X_s, s), \quad Z_s = -u^*(X_s, s) = \sigma^\top \nabla V(X_s, s). \tag{12}$$

The *target measure* $\mathbb{Q}$ characterises u^* uniquely.

A unifying perspective

We aim to find a functional $\mathcal{L}$ ("loss") that admits u^* as its unique global minimum. Then we can implement gradient-descent-like algorithms. We construct losses using divergences between path measures, leveraging the map $u \mapsto \mathbb{P}^u$ provided by (2):

$$\mathcal{L}_D(u) = D(\mathbb{P}^u|\mathbb{Q}), \qquad u \in \mathcal{U}, \tag{13}$$

where $D : \mathcal{P}(C([0,T];\mathbb{R}^d)) \times \mathcal{P}(C([0,T];\mathbb{R}^d)) \to \mathbb{R}_{\geq 0} \cup \{+\infty\}$ is a divergence between path measures. The perspective of approximating a measure on path space connects to variational inference.

We define the KL-based *relative entropy* and *cross-entropy* losses,

$$\mathcal{L}_{\mathrm{RE}}(u) = \mathbb{E}_{\mathbb{P}^u}\left[\log\frac{\mathrm{d}\mathbb{P}^u}{\mathrm{d}\mathbb{Q}}\right], \qquad \mathcal{L}_{\mathrm{CE}}(u) = \mathbb{E}_{\mathbb{Q}}\left[\log\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}^u}\right], \tag{14}$$

and, for $v \in \mathcal{U}$, the variance-based losses

$$\mathcal{L}_{\mathrm{Var}_v}(u) = \mathrm{Var}_{\mathbb{P}^v}\left(\frac{\mathrm{d}\mathbb{P}^u}{\mathrm{d}\mathbb{Q}}\right), \qquad \mathcal{L}^{\log}_{\mathrm{Var}_v}(u) = \mathrm{Var}_{\mathbb{P}^v}\left(\log\frac{\mathrm{d}\mathbb{P}^u}{\mathrm{d}\mathbb{Q}}\right). \tag{15}$$

We recover the formulation from the first column from (14) and (15):

Proposition 1 (Relative entropy vs. optimal control → Problem 1)

$$\mathcal{L}_{\mathrm{RE}}(u) = J(u) - \log\mathcal{Z}. \tag{16}$$

For arbitrary $v \in \mathcal{U}$ and $u^* = -Z$ we obtain the generalised forward-backward SDE system

$$\mathrm{d}X_s^v = (b(X_s^v, s) + \sigma(X_s^v, s)v(X_s^v, s))\,\mathrm{d}s + \sigma(X_s^v, s)\,\mathrm{d}W_s, \qquad X_0^v = x_0, \tag{17a}$$

$$\mathrm{d}Y_s^{u^*,v} = -f(X_s^v, s)\,\mathrm{d}s - (v \cdot u^*)(X_s^v, s)\,\mathrm{d}s + \frac{1}{2}|u_s^*|^2\,\mathrm{d}s - u_s^* \cdot \mathrm{d}W_s, \quad Y_T = g(X_T^v), \tag{17b}$$

Proposition 2 (Log-variance vs. forward-backward SDE → Problem 3)

$$\mathcal{L}^{\log}_{\mathrm{Var}_v}(u) = \mathrm{Var}(Y_T^{u,v} - g(X_T^v)). \tag{18}$$

This can be compared to a loss commonly used:

$$\mathcal{L}_{\mathrm{moment},v}(u, y_0) = \mathbb{E}\left[\left((Y_T^{u,v}(y_0) - g(X_T^v)\right)^2\right]. \tag{19}$$

Proposition 3 (Variance vs. importance sampling → Problem 5)

$$\mathcal{L}_{\mathrm{Var}_0}(u) = \mathrm{Var}\left(\exp(-\mathcal{W}(X^u))\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{P}^u}\right). \tag{20}$$

Infinite batch size properties

The loss $\mathcal{L}^{\log}_{\mathrm{Var}_v}$ is our favourite. Here are some properties:

Proposition 4 (Equivalence of log-variance and relative entropy loss) *The Gateaux-derivatives in direction $\phi \in C_b^1([0,T] \times \mathbb{R}^d;\mathbb{R}^d)$ satisfy*

$$\frac{1}{2}\left(\frac{\delta}{\delta u}\mathcal{L}^{\log}_{\mathrm{Var}_v}(u;\phi)\right)\Big|_{v=u} = \frac{\delta}{\delta u}\mathcal{L}_{\mathrm{RE}}(u;\phi). \tag{21}$$

This means that if the expectations required for the losses could be computed without sampling error, then $\mathcal{L}_{\mathrm{RE}}$ and $\mathcal{L}^{\log}_{\mathrm{Var}_v}$ would lead to equivalent algorithms.

Proposition 5 (Equivalence of log-variance and moment loss) *It holds that*

$$\left(\frac{\delta}{\delta u}\mathcal{L}_{\mathrm{moment},v}(u, y_0;\phi)\right)\Big|_{v=u} = \left(\frac{\delta}{\delta u}\mathcal{L}^{\log}_{\mathrm{Var}_v}(u;\phi)\right)\Big|_{v=u} \tag{22}$$

for all $\phi \in C_b^1(\mathbb{R}^d \times [0,T];\mathbb{R}^d)$ independently of $y_0 \in \mathbb{R}$.

To be precise, with $\widetilde{Y}_T^{u,v} := Y_T^{u,v} - y_0$ we have

$$\left(\frac{\delta}{\delta u}\mathcal{L}_{\mathrm{moment},v}(u, y_0;\phi)\right)\Big|_{v=u} = 2\,\mathbb{E}\left[\left(g(X_T^u) - \widetilde{Y}_T^{u,u}\right)\int_0^T \phi_s \cdot \mathrm{d}W_s\right] - 2y_0\,\mathbb{E}\left[\int_0^T \phi_s \cdot \mathrm{d}W_s\right] \tag{23}$$

and in Monte Carlo simulations y_0 has an impact on the variance of the estimator. In fact, the log-variance loss can be interpreted as a control variate version of the moment loss.

Proposition 6 (Absence of additional local minima) *Let $u \in \mathcal{U}$ and assume that*

$$\frac{\delta}{\delta u}\mathcal{L}_{\mathrm{RE}}(u;\phi) = 0, \tag{24}$$

for all $\phi \in C_b^1(\mathbb{R}^d \times [0,T];\mathbb{R}^d)$. Then $u = u^*$.

Robustness properties

However, $\mathcal{L}^{\log}_{\mathrm{Var}_v}$ is more stable, at least close to the solution u^*:

Proposition 7 (Stability of $\mathcal{L}^{\log}_{\mathrm{Var}_v}$ close to u^*) *At the solution u^*, the variance of the gradient estimator for the log-variance loss vanishes, i.e.*

$$\mathrm{Var}\left(\frac{\delta}{\delta u}\Big|_{u=u^*}\widehat{\mathcal{L}}^{\log}_{\mathrm{Var}_v}(u;\phi)\Big|_{v=u}\right) = 0. \tag{25}$$

This not true for the cross-entropy and relative entropy losses and holds for the moment loss if and only if $y_0 = -\log\mathcal{Z}$.

The log-variance estimator is stable in high dimensions:

Proposition 8 (Stability under tensorisation) *The relative error*

$$\frac{\sqrt{\mathrm{Var}\widehat{\mathcal{L}}^{\log}_{\mathrm{Var}}\left(\bigotimes_{i=1}^M \mathbb{P}_i|\bigotimes_{i=1}^M \mathbb{Q}_i\right)}}{\mathcal{L}^{\log}_{\mathrm{Var}}\left(\bigotimes_{i=1}^M \mathbb{P}_i|\bigotimes_{i=1}^M \mathbb{Q}_i\right)} \quad \textit{can be bounded uniformly in } M. \tag{26}$$

This is also true for the relative entropy and moment losses, but not for the variance and the cross-entropy losses.

Numerical examples

We approximate the optimal control u^* with a feed-forward neural network.

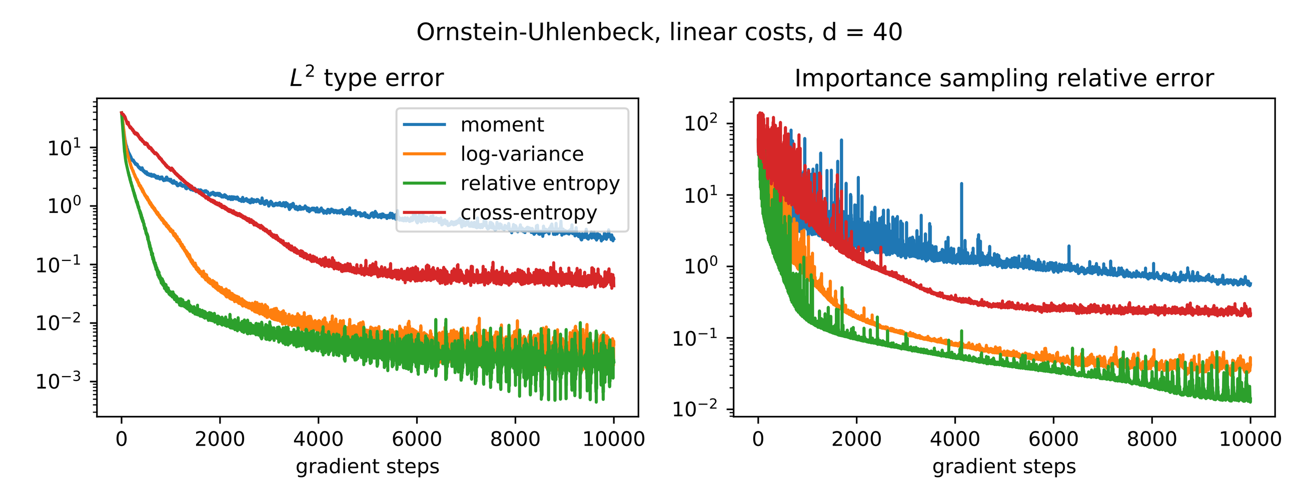

Fig. 1: Different convergence speeds of the losses, the independence of y_0 seems to be beneficial.

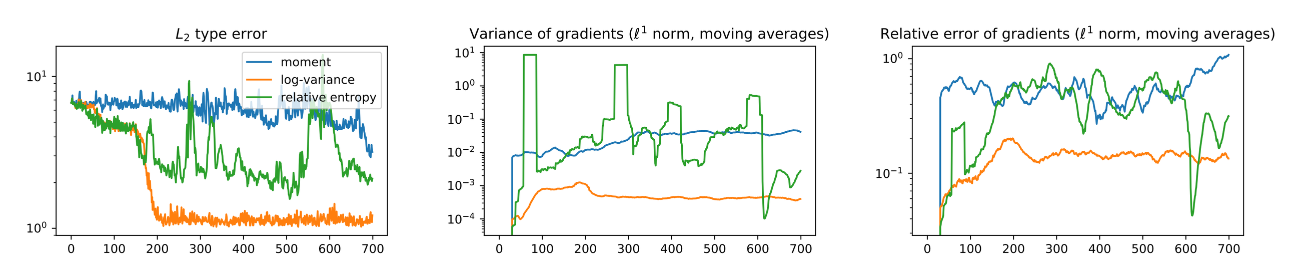

Fig. 2: One dimensional double well. This shows the instability of $\mathcal{L}_{\mathrm{RE}}$ close to u^*, see Proposition 7.

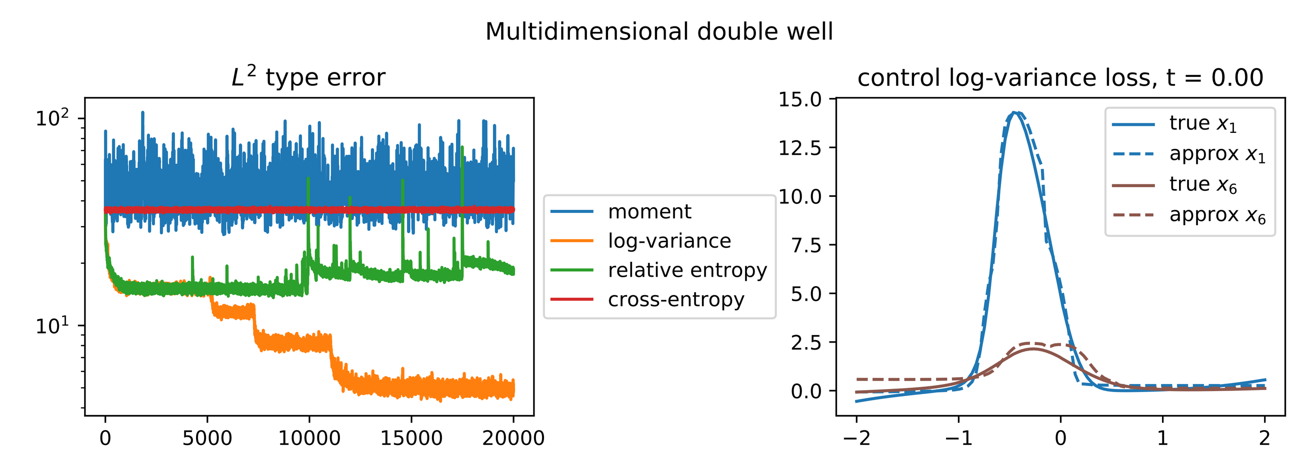

Fig. 3: Only the log-variance loss can cope with a 10-dimensional metastable double well setting, see Proposition 8.

Our paper will be on arXiv very soon. We appreciate any kind of question or comment on Slack or via e-mail.

Acknowledgement: This is joint work between projects A02 and A05 in CRC 1114.

SFB 1114