A robust deep FBSDE method for stochastic control

Kristoffer Andersson - Centrum Wiskunde & Informatica

Joint work with Adam Andersson and Cornelis W. Oosterlee

A Coruña,
April 19, 2022

Table of contents

Example: Control of pendulums on carts

Figure: Single and double pendulums on carts.

Example: Control of stochastic process

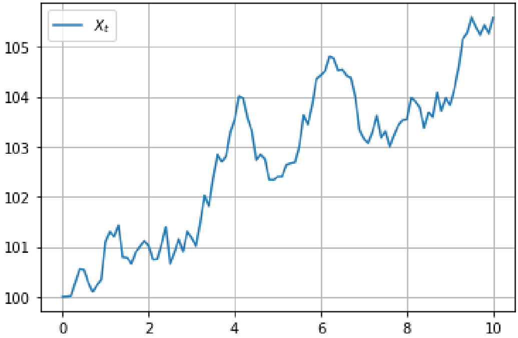

Figure: A trajectory of a stochastic process described by a SDE.

$$\textbf{SDE:} \quad X_t = x_0 + \int_0^t b(t, X_t)\mathrm{d}t + \int_0^t \sigma(t, X_t)\mathrm{d}W_t.$$

Example: Control of stochastic process

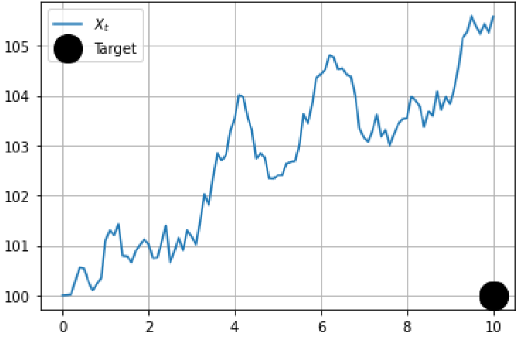

Figure: A trajectory of a stochastic process described by a SDE.

$$\textbf{SDE:} \quad X_t = x_0 + \int_0^t b(t, X_t)\mathrm{d}t + \int_0^t \sigma(t, X_t)\mathrm{d}W_t.$$

Example: Control of stochastic process

Figure: A trajectory of a stochastic process described by a SDE.

SDE: $X_t^u = x_0 + \int_0^t b(t, X_t^u, u_t)\mathrm{d}t + \int_0^t \sigma(t, X_t^u)\mathrm{d}W_t,$

Distance to target: $|X_T^u - 100|^2,$ **Expected distance to target:** $\mathbb{E}\left[|X_T^u - 100|^2\right].$

Stochastic control problem

Dynamical system described by an SDE referred to as the **State equation**

$$X_t^u = x_0 + \int_0^t \bar{b}(t, X_t^u, u_t)\mathrm{d}t + \int_0^t \bar{\sigma}(t, X_t^u)\mathrm{d}W_t.$$

$X^u = (X_t^u)_{t\in[0,T]}$ state of the system, $u = (u_t)_{t\in[0,T]}$ control of the system, taking on values in $\mathbb{R}^d$ and $U \subset \mathbb{R}^\ell$, respectively.

To measure performance of the control, a **cost functional** is used

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

The **control problem** is to find a control $u \in \mathcal{U}_{[0,T]}$($:=$ set of admissible controls) such that the cost functional is minimized.

FBSDE

Assuming $\bar{\sigma}$ is uniformly invertible, Itô's lemma gives the **FBSDE**

$$\begin{cases} X_t = x_0 + \int_0^t b(s, X_s, Z_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\mathrm{d}W_s, \\ Y_t = g(X_T) + \int_t^T f(s, X_s, Z_s)\mathrm{d}s - \int_t^T Z_s \mathrm{d}W_s, \end{cases} \tag{1}$$

where

$$\begin{aligned} Z_t &= \mathrm{D}_x V(t, X_t)^\top \sigma(t, X_t), \\ Y_t &= V(t, X_t), \end{aligned}$$

where $b(t, X_t, Z_t) := \bar{b}(t, X_t, u^*)$, $f(t, X_t, Z_t) := \bar{f}(t, X_t, u^*)$ and $\sigma := \bar{\sigma}$.

The solution to (1) is the triple (X, Y, Z) of adapted, square integrable processes.

Reformulation for deep BSDE solver

The **Deep BSDE solver**[1] uses the follwing reformulation of a FBSDE

$$\begin{cases} \inf\limits_{y_0,(z_t)_{t\in[0,T]}} \mathbb{E}|Y_T - g(X_T)|^2, \quad \text{subject to} \\ X_t = x_0 + \int_0^t b(s, X_s, z_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\mathrm{d}W_s, \\ Y_t = y_0 - \int_0^t f(s, X_s, z_s)\mathrm{d}s + \int_0^t z_s \mathrm{d}W_s, \end{cases} \tag{2}$$

compared to

$$\begin{cases} X_t = x_0 + \int_0^t b(s, X_s, Z_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\mathrm{d}W_s, \\ Y_t = g(X_T) + \int_t^T f(s, X_s, Z_s)\mathrm{d}s - \int_t^T Z_s \mathrm{d}W_s, \end{cases} \tag{3}$$

Motivation:

1. A solution to (3) solves (2);
2. By wellposedness of the FBSDE, the solution is unique.

Problem: Time discrete version does not in general converge to continuous problem. This is shown in our paper, and we elaborate on why and when this occurs.

---

[1] A. Jentzen et. al. *Solving high-dimensional partial differential equations using deep learning.* Proceedings of the National Academy of Sciences 115.34 (2018): 8505-8510.

Our reformulation - using properties from stochastic control problem

Our solution: Impose known structure from the stochastic control formulation of the problem. In our paper we show theoretically and empirically that our method converges.

$$\begin{cases} \inf\limits_{(z_t)_{t\in[0,T]}} \mathbb{E}[\mathcal{Y}_0(z)] + \lambda \,\mathrm{Var}[\mathcal{Y}_0(z)], \quad \text{subject to} \\ \mathcal{Y}_0(z) = g(X_T) + \int_0^T f(t, X_t, z_t)\mathrm{d}t - \int_0^T z_t \mathrm{d}W_t, \quad Y_0 = \mathbb{E}[\mathcal{Y}_0(z)] \\ X_t = x_0 + \int_0^t b(s, X_s, z_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\mathrm{d}W_s, \\ Y_t = Y_0 - \int_0^t f(s, X_s, z_s)\mathrm{d}s + \int_0^t z_s \mathrm{d}W_s. \end{cases} \tag{4}$$

1. Y_0 coincides with the value function of the control problem (property from the control problem),
2. Y_0 is $\mathcal{F}_0-$measurable and therefore has zero variance (property from the FBSDE).

Discussion

Why should our algorithms work better?

- Using mathematical structure from the specific problem leads to fewer entities to approximate;
- Loss surface of objective function seems to be nice and monotonic $\rightarrow$ easy to optimize;

Disdvantage of our algorithms:

- While deep BSDE solver is (at least conceptually) applicable for all FBSDEs, our algorithm is applicable only for FBSDEs steming from stochastic control problems.

Setup of numerical experiments

- All optimization problems are approximated with the help of ANNs, but any function approximator efficient enough could be used;
- In the following examples, we have analytical solutions available to compare with;
- For each solution component, X, Y and Z we compare to analytical counterpart in strong and weak sense;
- One problem with control in each spatial dimension $(\mathrm{d} = \ell) \rightarrow$ Algorithm 1;
- One problem with control in each spatial dimension $(\mathrm{d} > \ell) \rightarrow$ Algorithm 2.

Strong and weak approximations of X for $\ell = d = 2$

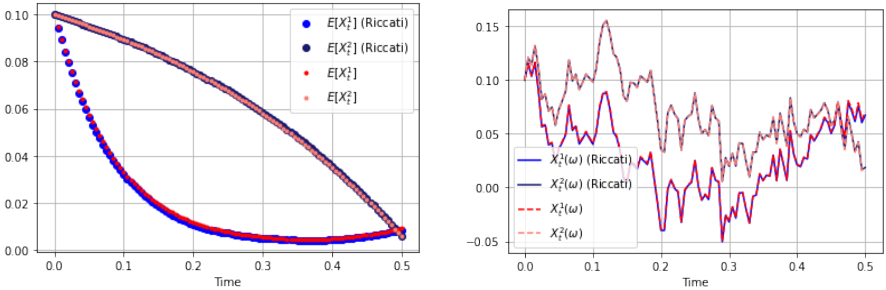

Figure: Average of solutions and a single solution path compared to their analytical counterparts.

Strong and weak approximations of Y for $\ell = d = 2$

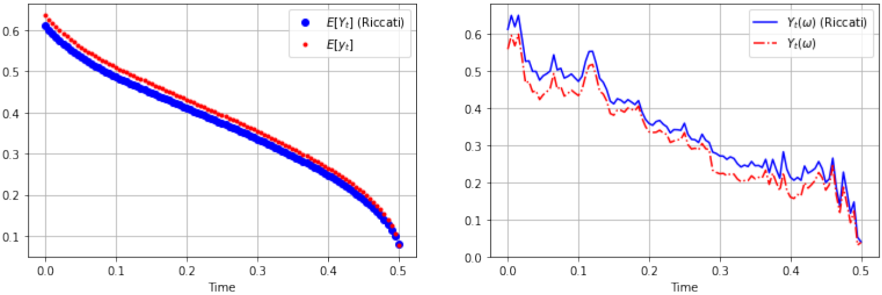

Figure: Average of solutions and a single solution path compared to their analytical counterparts.

Strong and weak approximations of Z for $\ell = d = 2$

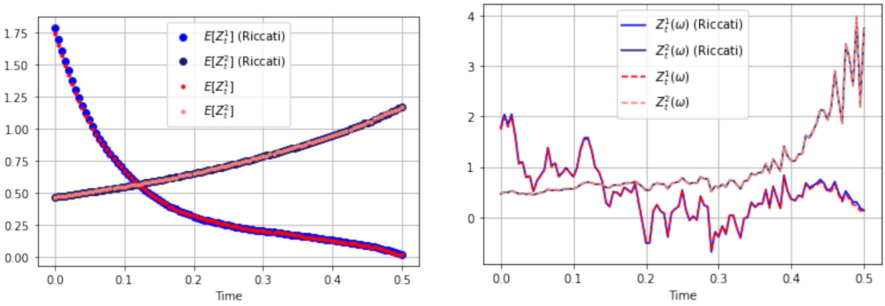

Figure: Average of solutions and a single solution path compared to their analytical counterparts.

Strong and weak approximations of X for $\ell = 2$ and $d = 6$

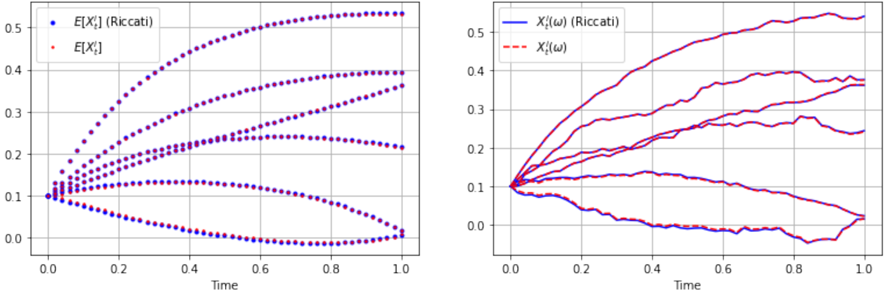

Figure: Average of solutions and a single solution path compared to their analytical counterparts.

Strong and weak approximations of Y for $\ell = 2$ and $d = 6$

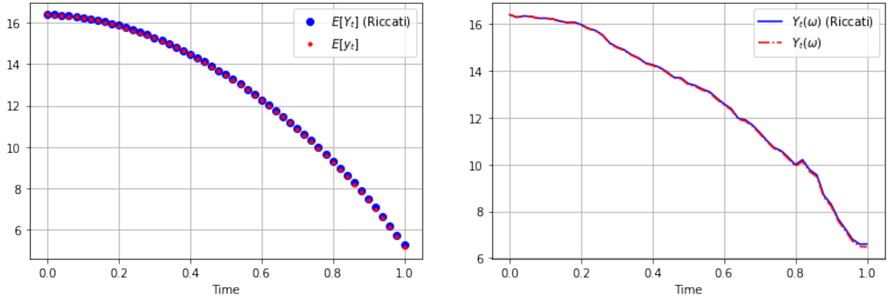

Figure: Average of solutions and a single solution path compared to their analytical counterparts.

Strong and weak approximations of Z for $\ell = 2$ and $d = 6$

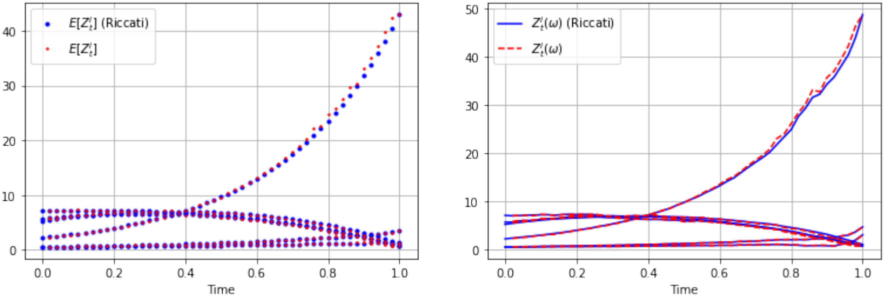

Figure: Average of solutions and a single solution path compared to their analytical counterparts.

Convergence analysis

Bound for strong error of our numerical scheme:

$$\left\|X-\hat{X}^{h,\lambda}\right\|_{\mathcal{S}^2(\mathbb{R}^d)}+\left\|Y-\hat{Y}^{h,\lambda}\right\|_{\mathcal{S}^2(\mathbb{R}^d)}+\left\|Z-\hat{Z}^{h,\lambda}\right\|_{\mathcal{H}^2(\mathbb{R}^k)}$$
$$\leq C\left(h^{\frac{\alpha}{2}}+\max_{0,\dots,N_h}\left(\mathbb{E}\left[\left\|\zeta^*(t_n,X_n^{\lambda,h})-\hat{\zeta}^*_{h,\lambda}(t_n,X_n^{\lambda,h})\right\|^2\right]\right)^{\frac{1}{2}}\right).$$

For the initial condition of BSDE:

$$\left|Y_0-Y_0^{h,\lambda}\right|+\mathrm{Var}\left(\mathcal{Y}_0^{h,\lambda}\right)\leq Ch^{\alpha}.$$

For the terminal condition of BSDE:

$$\left(\mathbb{E}\left[\left(g\left(X_N^{h,\lambda}\right)-Y_N^{h,\lambda}\right)^2\right]\right)^{\frac{1}{2}}\leq Ch^{\frac{\alpha}{2}}.$$

Convergence plots

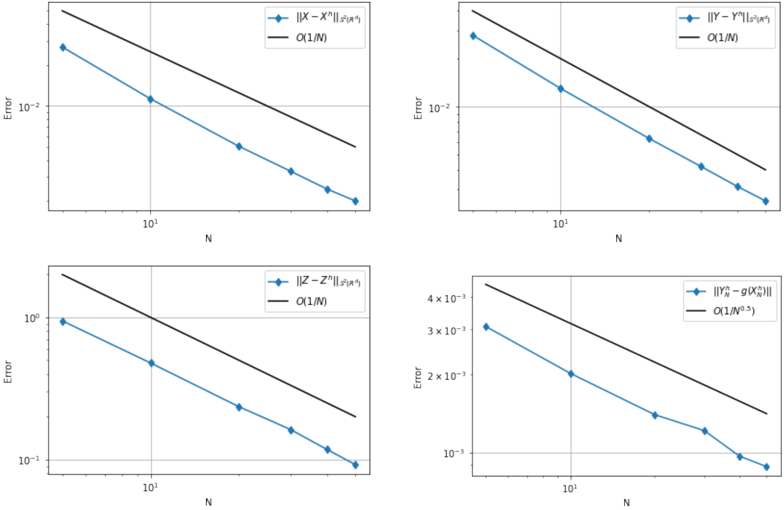

Figure: Convergence plots.

Thanks for your attention!

Value function and HJB equation

If infimum is attainable, the **value function**, V, satisfies

$$V(t,x) \in \inf_{u \in \mathcal{U}_{[t,T]}} J^u(t,x).$$

Under mild conditions, the value function satisfies a **Hamilton–Jacobi–Bellman** equation, which is a quasi-linear parabolic PDE given by

$$\begin{cases} \frac{\partial V}{\partial t}(t,x) + \mathcal{H}(t,x,\mathrm{D}_x V(t,x), \mathrm{D}_x^2 V(t,x)) = 0, & (t,x) \in [0,T) \times \mathbb{R}^n, \\ V(t,x) = g(x), & (t,x) \in \{T\} \times \mathbb{R}^n, \end{cases} \tag{5}$$

where the **Hamiltonian**, $\mathcal{H}$, is given by

$$\mathcal{H}(t,x,p,q) = \inf_{v \in U} \left[\bar{b}(t,x,v)^\top p + \bar{f}(t,x,v) \right],$$

for all $p \in \mathbb{R}^d$.

By inspection, **feedback control** on the form $u_t^* = u^*\left(t, X_t^{u^*}, D_x V(t, X_t^{u^*})\right)$.

Assumption: From now on, assume (5) has a solution V, with sufficiently many bounded derivatives.

Equidistant time grid $\pi := \{0 = t_0, t_1, \dots, t_N = T\}$, with $h = t_{n+1} - t_n$ and Brownian increments $\Delta W_n = W_{n+1} - W_n$.

Time discrete formulation:

$$\begin{cases} \inf_{\bar{y}_0, \{\bar{z}_k\})_{k \in \{0,1,\dots,n-1\}}} \mathbb{E}|\bar{Y}_N - g(\bar{X}_N)|^2, \quad \text{subject to} \\ \bar{X}_n = x_0 + \sum_{k=0}^{n-1} b(t_k, \bar{X}_k, \bar{z}_k)h + \sum_{k=0}^{n-1} \sigma(t_k, \bar{X}_k)\Delta W_k, \\ \bar{Y}_n = \bar{y}_0 - \sum_{k=0}^{n-1} f(t_k, \bar{X}_k, \bar{z}_k)h + \sum_{k=0}^{n-1} \bar{z}_k \Delta W_k. \end{cases} \tag{6}$$

To investigate convergence:

1. Fix $\bar{y}_0 \in \mathbb{R}$;
2. Minimize the objective in (6) (only over $\bar{z}_0, \bar{z}_1, \dots, \bar{z}_{N-1}$, since $\bar{y}_0$ fixed);
3. Explore the values of the objective function for different $\bar{y}_0$ (one optimization per $\bar{y}_0$).

If convergence, then for small h, $\bar{y}_0 \approx Y_0$ should yield the smallest value of the objective function.

Is this what we observe empirically? For nice and linear problems, yes. Otherwise, **No**!

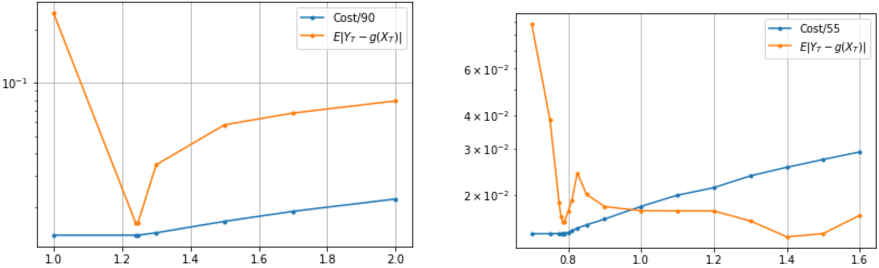

Figure: **Left:** Simple linear problem, $Y_0 = 1.27$. **Right:** Non-linear problem, $Y_0 = 0.79$.

Left: Y_0 is close to the global minimum of the objective function $\rightarrow$ convergence!

Right: Y_0 is close to a local minimum of the objective function. But **not** close to global minimum $\rightarrow$ **No** convergence!

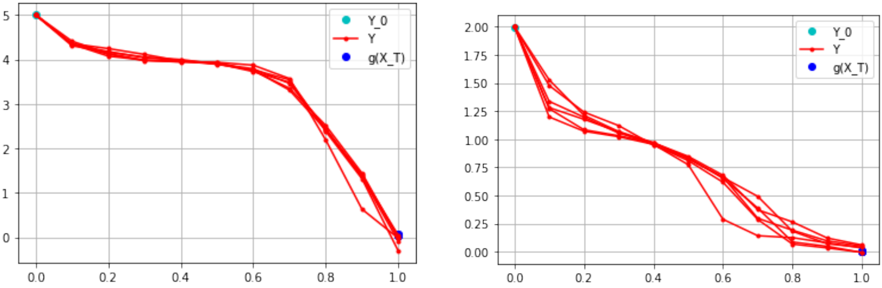

Figure: True $Y_0 = 1.27$ **Left:** $\bar{y}_0 = 5$. **Right:** $\bar{y}_0 = 2$.

Conclusion:

Possible to control system to "almost satisfy" terminal condition for many different $\bar{y}_0$
Therefore, no convergence.

From continuous to discrete formulations

Want to approximate FBSDE with a time discrete counterpart. Have seen equivalence between FBSDE and continuous variational problems.

Question: What happens to the reformulations in a time discrete setting?

Answer:

- (Discrete counterpart of) variational formulation used for deep BSDE solver converges for weakly coupled FBSDEs. Not easy to show convergence in the strongly coupled case;
- Some empirical evidence indicates no convergence;
- (Discrete counterparts of) our reformulations converges with mild assumptions also for strongly coupled FBSDEs in Y_0.
- Under additional assumptions proof that entire FBSDE converges;
- Empirical convergence for all problems investigated.