

A robust deep learning algorithm for coupled FBSDEs

Kristoffer Andersson, Adam Andersson & C.W. Oosterlee

Centrum Wiskunde & Informatica

ICCF Wuppertal,
June 9, 2022

Table of contents

- 1 Motivation from stochastic control
- 2 The deep BSDE method and why it may not be suitable for coupled FBSDEs
- 3 Our proposed alternative algorithms
- 4 Numerical examples

Stochastic control problem

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

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

$X^u = (X_t^u)_{t \in [0, T]}$ 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

$$J^u(t, x) = \mathbb{E}_{t, x} \left[\int_t^T \bar{f}(t, X_t^u, u_t) dt + g(X_T^u) \right].$$

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.

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

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

FBSDE

Assumption: The value function V is sufficiently smooth and σ is of full rank.

Then by Itô's lemma we obtain the **FBSDE**

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

where

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

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

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

Reformulation of the FBSDE to a stochastic control problem

The **Deep BSDE method**¹ uses the following reformulation of a FBSDE

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

compared to

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

Motivation:

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

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

Time discretization of the control problem

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:

$$\left\{ \begin{array}{l} \inf_{\bar{y}_0, \{\bar{z}_k\}_{k \in \{0, 1, \dots, N-1\}}} \mathbb{E} |\bar{Y}_N - g(\bar{X}_N)|^2, \quad \text{subject to} \\ \bar{X}_n = x_0 + \sum_{k=0}^{n-1} b(t_k, \bar{X}_k, \bar{z}_k)h + \sum_{k=0}^{n-1} \sigma(t_k, \bar{X}_k) \Delta W_k, \\ \bar{Y}_n = \bar{y}_0 - \sum_{k=0}^{n-1} f(t_k, \bar{X}_k, \bar{z}_k)h + \sum_{k=0}^{n-1} \bar{z}_k \Delta W_k. \end{array} \right. \quad (4)$$

Aim: Approximate the FBSDE with discrete version of control problem.

Question: Does the discrete control problem converge to the continuous FBSDE?

- For weakly coupled FBSDEs an aposteriori error estimate exists ² but for general coupled FBSDE no convergence results available(?);
- We show empirically that for coupled FBSDEs convergence is highly non-trivial.

²Han, J and Long J. *Convergence of the deep BSDE method for coupled FBSDEs*. Probability, Uncertainty and Quantitative Risk 5.1 (2020): 1-33.

Investigate convergence of the deep FBSDE method

To investigate convergence:

- 1 Fix $\bar{y}_0 \in \mathbb{R}$;
- 2 Minimize the objective in (5) (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 $\bar{y}_0$).

$$\left\{ \begin{array}{l} \inf_{\{\bar{z}_k\}_{k \in \{0,1,\dots,N-1\}}} \mathbb{E} |\bar{Y}_N - g(\bar{X}_N)|^2, \quad \text{subject to} \\ \bar{X}_n = x_0 + \sum_{k=0}^{n-1} b(t_k, \bar{X}_k, \bar{z}_k) h + \sum_{k=0}^{n-1} \sigma(t_k, \bar{X}_k) \Delta W_k, \\ \bar{Y}_n = \bar{y}_0 - \sum_{k=0}^{n-1} f(t_k, \bar{X}_k, \bar{z}_k) h + \sum_{k=0}^{n-1} \bar{z}_k \Delta W_k. \end{array} \right. \quad (5)$$

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 very simple LQG control 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.

Alternative reformulation

When there is a one-to-one mapping between u_t^* and Z_t (usually when $\ell = d$) we use

$$\begin{cases} \inf_{(z_t)_{t \in [0, T]}} \mathbb{E}[\hat{y}_0(z)], & \text{subject to} \\ \hat{y}_0(z) = g(X_T) + \int_0^T f(t, X_t, z_t) dt - \int_0^T z_t dW_t, & \mathcal{Y}_0 = \mathbb{E}[\hat{y}_0(z)], \\ X_t = x_0 + \int_0^t b(s, X_s, z_s) ds + \int_0^t \sigma(s, X_s) dW_s, \\ Y_t = \mathcal{Y}_0 - \int_0^t f(s, X_s, z_s) ds + \int_0^t z_s dW_s. \end{cases} \quad (6)$$

compared to

$$\begin{cases} X_t = x_0 + \int_0^t b(s, X_s, Z_s) ds + \int_0^t \sigma(s, X_s) dW_s, \\ Y_t = g(X_T) + \int_t^T f(s, X_s, Z_s) ds - \int_t^T Z_s dW_s, \end{cases} \quad (7)$$

Motivation:

- 1 There is a one-to-one map between u_t^* and Z_t implying that we can optimize z instead of u . Therefore, a solution to (7) solves (6);
- 2 By wellposedness of the FBSDE, the solution is unique.

Extended reformulation

When there is **not** a one-to-one mapping between u_t^* and Z_t (usually when $\ell < d$).

Problem: No longer uniqueness of a minimizer in previous formulation.

Solution: Add an extra term to the objective function.

$$\left\{ \begin{array}{l} \inf_{(z_t)_{t \in [0, T]}} \mathbb{E}[\hat{y}_0(z)] + \lambda \text{Var}[\hat{y}_0(z)], \quad \text{subject to} \\ \hat{y}_0(z) = g(X_T) + \int_0^T f(t, X_t, z_t) dt - \int_0^T z_t dW_t, \quad \mathcal{Y}_0 = \mathbb{E}[\hat{y}_0(z)] \\ X_t = x_0 + \int_0^t b(s, X_s, z_s) ds + \int_0^t \sigma(s, X_s) dW_s, \\ Y_t = \mathcal{Y}_0 - \int_0^t f(s, X_s, z_s) ds + \int_0^t z_s dW_s. \end{array} \right. \quad (8)$$

Motivation:

There are infinitely many processes z that minimize the cost of the control problem but only one of them yields a deterministic initial condition of the FBSDE (by uniqueness of the FBSDE solution).

Discretization done by Euler-Maruyama of the two forward SDEs as in the deep BSDE method.

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 method is (at least conceptually) applicable for all FBSDEs, our algorithm is applicable only for FBSDEs stemming 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 ($d = \ell$) $\rightarrow$ Algorithm 1;
- One problem with control in each spatial dimension ($d > \ell$) $\rightarrow$ 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.

Summary and conclusion

- We have shown empirically that a straight forward extension of the deep BSDE method to strongly coupled FBSDEs fail to converge for many simple LQG control problems;
- We have introduced a method which does not have the initial condition of the BSDE as a free parameter to be optimized. Instead the initial condition is given by the minimized objective function;
- The most interesting extension would be to make method applicable also to FBSDEs not stemming from a stochastic control problem. For this to be successful, one either need to know the initial state of the BSDE a priori (which is almost never the case) or come up with a generalization of the objective function used in our work.

Thanks for your attention!