



FiNANCE FOR ENERGY MARKET RESEARCH CENTRE



Neural networks-based backward scheme for fully nonlinear PDEs *

Huyen Pham, Xavier Warin,

Working Paper
RR-FiME-19-07

August 2019

Neural networks-based backward scheme for fully nonlinear PDEs *

Huyên PHAM[†]

Xavier WARIN[‡]

August 22, 2019

Abstract

We propose a numerical method for solving high dimensional fully nonlinear partial differential equations (PDEs). Our algorithm estimates simultaneously by backward time induction the solution and its gradient by multi-layer neural networks, through a sequence of learning problems obtained from the minimization of suitable quadratic loss functions and training simulations. This methodology extends to the fully nonlinear case the approach recently proposed in [HPW19] for semi-linear PDEs. Numerical tests illustrate the performance and accuracy of our method on several examples in high dimension with nonlinearity on the Hessian term including a linear quadratic control problem with control on the diffusion coefficient.

Key words: Deep neural networks, fully nonlinear PDEs in high dimension.

MSC Classification: 60H35, 65C20, 65M12.

1 Introduction

This paper is devoted to the resolution in high dimension of fully nonlinear parabolic partial differential equations (PDEs) of the form

$$\begin{cases} \partial_t u + f(\cdot, \cdot, u, D_x u, D_x^2 u) = 0, & \text{on } [0, T] \times \mathbb{R}^d, \\ u(T, \cdot) = g, & \text{on } \mathbb{R}^d, \end{cases} \quad (1.1)$$

with a non-linearity in the solution, its gradient and its hessian via the function $f(t, x, y, z, \gamma)$ defined on $[0, T] \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{S}^d$ (where \mathbb{S}^d is the set of symmetric $d \times d$ matrices), and a terminal condition g .

The numerical resolution of this class of PDEs is far more difficult than the one of classical semi-linear PDEs where the nonlinear function f does not depend on γ . In fact, rather few methods are available to solve fully nonlinear equations even in moderate dimension.

- First based on the work of [Che+07], an effective scheme developed in [FTW11] using some regression techniques has been shown to be convergent under some ellipticity conditions later removed by [Tan+13]. Due to the use of basis functions, this scheme does not permit to solve PDE in dimension greater than 5.
- A scheme based on nesting Monte Carlo has been recently proposed in [War18]. It seems to be effective in very high dimension for maturities not too long and linearities not too important.
- A numerical algorithm to solve fully nonlinear equations has been proposed by [BEJ17] based on the representation of [Che+07] and global deep neural networks minimizing a terminal objective function, but no test on real fully nonlinear case is given.
- The Deep Galerkin method proposed in [SS18] based on some machine learning techniques and using some numerical differentiation of the solution seems to be effective on some cases. It has been tested in [AA+18] for example on the Merton problem.

In this article, we introduce a numerical method based on machine learning techniques and backward in time iterations, which extends the proposed scheme for semi-linear PDEs in the recent work [HPW19]. A first idea to extend this work consists in using the representation proposed in [Che+07] and applied numerically in [FTW11] and [BEJ17] to solve the problem: more precisely, at each time step t_n of an Euler scheme, $D_x^2 u$ at t_n is approximated by a neural network minimizing some local L_2 criterion associated to a BSDE involving $D_x u$

*This work is supported by FiME, Laboratoire de Finance des Marchés de l’Energie, and the ”Finance and Sustainable Development” EDF - CACIB Chair.

[†]LPSM, Paris-Diderot University, CREST-ENSAE & FiME pham at lpsm.paris

[‡]EDF R&D & FiME xavier.warin at edf.fr

at date t_{n+1} and $D_x^2 u$. Then, the pair $(u, D_x u)$ at date t_n is approximated/learned with a second minimization similarly as in the method described by [HPW19]. The first minimization can be implemented with different variations but numerical results show that the global scheme does not scale well with the dimension. Instability on the $D_x^2 u$ calculation rapidly propagates during the backward resolution. Besides, the methodology appears to be costly when using two optimizations at each time step. An alternative approach that we develop here, is to combine the ideas of [HPW19] and the splitting method in [Bec+19] in order to derive a new deep learning scheme that requires only one local optimization during the backward resolution for learning the pair $(u, D_x u)$ and approximating $D_x^2 u$.

The outline of the paper is organized as follows. In Section 2, we briefly recall the mathematical description of the classical feedforward approximation, and then derive the proposed neural networks-based backward scheme. We test our method in Section 3 on several examples: first we illustrate our results with a PDE involving a non linearity of type $uD_x^2 u$, then we consider a stochastic linear quadratic problem with controlled volatility where an analytic solution is available, and we test the performance and accuracy of our algorithm up to dimension 20. The convergence analysis of our scheme is postponed to a forthcoming work.

2 The proposed deep backward scheme

Our aim is to numerically approximate the function $u : [0, T] \times \mathbb{R}^d \mapsto \mathbb{R}$, assumed to be the unique smooth solution to the fully nonlinear PDE (1.1) under suitable conditions. This will be achieved by means of neural networks approximations for u and its gradient $D_x u$, relying on a backward scheme and training simulated data of some forward diffusion process.

2.1 Feedforward neural network to approximate functions

We denote by d_0 the dimension of the input variables, and d_1 the dimension of the output variable. A (deep) neural network is characterized by a number of layers $L + 1 \in \mathbb{N} \setminus \{1, 2\}$ with m_ℓ , $\ell = 0, \dots, L$, the number of neurons (units or nodes) on each layer: the first layer is the input layer with $m_0 = d$, the last layer is the output layer with $m_L = d_1$, and the $L - 1$ layers between are called hidden layers, where we choose for simplicity the same dimension $m_\ell = m$, $\ell = 1, \dots, L - 1$.

A feedforward neural network is a function from \mathbb{R}^{d_0} to \mathbb{R}^{d_1} defined as the composition

$$x \in \mathbb{R}^d \mapsto A_L \circ \varrho \circ A_{L-1} \circ \dots \circ \varrho \circ A_1(x) \in \mathbb{R}. \quad (2.1)$$

Here A_ℓ , $\ell = 1, \dots, L$ are affine transformations: A_1 maps from \mathbb{R}^d to \mathbb{R}^m , A_2, \dots, A_{L-1} map from \mathbb{R}^m to \mathbb{R}^m , and A_L maps from \mathbb{R}^m to \mathbb{R}^{d_1} , represented by

$$A_\ell(x) = \mathcal{W}_\ell x + \beta_\ell, \quad (2.2)$$

for a matrix \mathcal{W}_ℓ called weight, and a vector β_ℓ called bias term, $\varrho : \mathbb{R} \rightarrow \mathbb{R}$ is a nonlinear function, called activation function, and applied component-wise on the outputs of A_ℓ , i.e., $\varrho(x_1, \dots, x_m) = (\varrho(x_1), \dots, \varrho(x_m))$. Standard examples of activation functions are the sigmoid, the ReLu, the Elu, tanh.

All these matrices \mathcal{W}_ℓ and vectors β_ℓ , $\ell = 1, \dots, L$, are the parameters of the neural network, and can be identified with an element $\theta \in \mathbb{R}^{N_m}$, where $N_m = \sum_{\ell=0}^{L-1} m_\ell(1+m_{\ell+1}) = d_0(1+m) + m(1+m)(L-2) + m(1+d_1)$ is the number of parameters. We denote by $\mathcal{N}_{d_0, d_1, L, m}$ the set of all functions generated by (2.1) for $\theta \in \mathbb{R}^{N_m}$.

2.2 Forward-backward representation

Let us introduce a forward diffusion process

$$X_t = X_0 + \int_0^t \mu(s, X_s) ds + \int_0^t \sigma(s, X_s) dW_s, \quad 0 \leq t \leq T, \quad (2.3)$$

where μ is a function defined on $[0, T] \times \mathbb{R}^d$ with values in \mathbb{R}^d , σ is a function defined on $[0, T] \times \mathbb{R}^d$ with values in \mathbb{M}^d the set of $d \times d$ matrices, and W a d -dimensional Brownian motion on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$ satisfying the usual conditions. The process X will be used for the simulation of training data in our deep learning algorithm, and we shall discuss later the choice of the drift and diffusion coefficients μ and σ , see Remark 2.2.

Let us next denote by (Y, Z, Γ) the triple of \mathbb{F} -adapted processes valued in $\mathbb{R} \times \mathbb{R}^d \times \mathbb{S}^d$, defined by

$$Y_t = u(t, X_t), \quad Z_t = D_x u(t, X_t), \quad \Gamma_t = D_x^2 u(t, X_t), \quad 0 \leq t \leq T. \quad (2.4)$$

By Itô's formula applied to $u(t, X_t)$, and since u is solution to (1.1), we see that (Y, Z, Γ) satisfies the backward equation:

$$\begin{aligned} Y_t &= g(X_T) - \int_t^T [\mu(s, X_s) \cdot Z_s + \frac{1}{2} \text{tr}(\sigma \sigma^\top(s, X_s) \Gamma_s) - f(s, X_s, Y_s, Z_s, \Gamma_s)] ds \\ &\quad - \int_t^T \sigma^\top(s, X_s) Z_s \cdot dW_s, \quad 0 \leq t \leq T. \end{aligned} \quad (2.5)$$

2.3 Algorithm

We now provide a numerical approximation of the forward backward system (2.3)-(2.5), and consequently of the solution u (as well as its gradient $D_x u$) to the PDE (1.1).

We start from a time grid $\pi = \{t_i, i = 0, \dots, N\}$ of $[0, T]$, with $t_0 = 0 < t_1 < \dots < t_N = T$, and time steps $\Delta t_i := t_{i+1} - t_i, i = 0, \dots, N-1$. The time discretization of the forward process X on π is then equal (typically when μ and σ are constants) or approximated by an Euler scheme:

$$X_{t_{i+1}} = X_{t_i} + \mu(t_i, X_{t_i}) \Delta t_i + \sigma(t_i, X_{t_i}) \Delta W_{t_i}, \quad i = 0, \dots, N-1,$$

where we set $\Delta W_{t_i} := W_{t_{i+1}} - W_{t_i}$ (by misuse of notation, we keep the same notation X for the continuous time diffusion process and its Euler scheme). The backward SDE (2.5) is approximated by the time discretized scheme

$$Y_{t_i} \simeq Y_{t_{i+1}} - [\mu(t_i, X_{t_i}) \cdot Z_{t_i} + \frac{1}{2} \text{tr}(\sigma \sigma^\top(t_i, X_{t_i}) \Gamma_{t_i}) - f(t_i, X_{t_i}, Y_{t_i}, Z_{t_i}, \Gamma_{t_i})] \Delta t_i - \sigma^\top(t_i, X_{t_i}) Z_{t_i} \cdot \Delta W_{t_i},$$

that is written in forward form as

$$Y_{t_{i+1}} \simeq F(t_i, X_{t_i}, Y_{t_i}, Z_{t_i}, \Gamma_{t_i}, \Delta t_i, \Delta W_{t_i}), \quad i = 0, \dots, N-1, \quad (2.6)$$

with

$$\begin{aligned} F(t, x, y, z, \gamma, h, \Delta) &:= y - \tilde{f}(t, x, y, z, \gamma) h + z^\top \sigma(t, x) \Delta, \\ \tilde{f}(t, x, y, z, \gamma) &:= f(t, x, y, z, \gamma) - \mu(t, x) \cdot z - \frac{1}{2} \text{tr}(\sigma \sigma^\top(t, x) \gamma). \end{aligned} \quad (2.7)$$

The idea of the proposed scheme is the following. Similarly as in [HPW19], we approximate at each time t_i , $u(t_i, \cdot)$ and its gradient $D_x u(t_i, \cdot)$, by neural networks $x \in \mathbb{R}^d \mapsto (\mathcal{U}_i(x; \theta), \mathcal{Z}_i(x; \theta))$ with parameter θ that are learned optimally by backward induction: suppose that $\hat{\mathcal{U}}_{i+1} := \mathcal{U}_{i+1}(\cdot; \theta_{i+1}^*)$, $\hat{\mathcal{Z}}_{i+1} := \mathcal{Z}_{i+1}(\cdot; \theta_{i+1}^*)$ is an approximation of $u(t_{i+1}, \cdot)$ and $D_x u(t_{i+1}, \cdot)$ at time t_{i+1} , then θ_i^* is computed from the minimization of the quadratic loss function:

$$\hat{L}_i(\theta) = \mathbb{E} \left| \hat{\mathcal{U}}_{i+1} - F(t_i, X_{t_i}, \mathcal{U}_i(X_{t_i}; \theta), \mathcal{Z}_i(X_{t_i}; \theta), D \hat{\mathcal{Z}}_{i+1}(\mathcal{T}(X_{t_{i+1}})), \Delta t_i, \Delta W_{t_i}) \right|^2 \quad (2.8)$$

where \mathcal{T} is a truncation operator such that $\mathcal{T}(X)$ is bounded for example by a quantile of the diffusion process and $D \hat{\mathcal{Z}}_{i+1}$ stands for the numerical differentiation of $\hat{\mathcal{Z}}_{i+1}$. The truncation permits to avoid that the oscillations of the neural network fitted in zone where the simulations are scarce propagate to areas of importance.

The intuition for the relevance of this scheme to the approximation of the PDE (1.1) is the following. From (2.4) and (2.6), the solution u to (1.1) should approximately satisfy

$$u(t_{i+1}, X_{t_{i+1}}) \simeq F(t_i, X_{t_i}, u(t_i, X_{t_i}), D_x u(t_i, X_{t_i}), D_x^2 u(t_i, X_{t_i}), \Delta t_i, \Delta W_{t_i}).$$

Suppose that at time t_{i+1} , $\hat{\mathcal{U}}_{i+1}$ is an estimation of $u(t_{i+1}, \cdot)$. Recalling the expression of F in (2.7), the quadratic loss function at time t_i is then approximately equal to

$$\begin{aligned} \hat{L}_i(\theta) &\simeq \mathbb{E} \left| u(t_i, X_{t_i}) - \mathcal{U}_i(X_{t_i}; \theta) + (D_x u(t_i, X_{t_i}) - \mathcal{Z}_i(X_{t_i}; \theta))^\top \sigma(t_i, X_{t_i}) \Delta W_{t_i} \right. \\ &\quad \left. - \Delta t_i [\tilde{f}(t_i, X_{t_i}, u(t_i, X_{t_i}), D_x u(t_i, X_{t_i}), D_x^2 u(t_i, X_{t_i})) - \tilde{f}(t_i, X_{t_i}, \mathcal{U}_i(X_{t_i}; \theta), \mathcal{Z}_i(X_{t_i}; \theta), D \hat{\mathcal{Z}}_{i+1}(\mathcal{T}(X_{t_{i+1}})))] \right|^2. \end{aligned}$$

By assuming that \tilde{f} has small nonlinearities in its arguments (y, z, γ) , say Lipschitz, possibly with a suitable choice of μ, σ , the loss function is thus approximately equal to

$$\hat{L}_i(\theta) \simeq (1 + O(\Delta t_i)) \mathbb{E} |u(t_i, X_{t_i}) - \mathcal{U}_i(X_{t_i}; \theta)|^2 + O(\Delta t_i) \mathbb{E} |D_x u(t_i, X_{t_i}) - \mathcal{Z}_i(X_{t_i}; \theta)|^2 + O(|\Delta t_i|^2).$$

Therefore, by minimizing over θ this quadratic loss function, via stochastic gradient descent (SGD) based on simulations of $(X_{t_i}, X_{t_{i+1}}, \Delta W_{t_i})$ (called training data in the machine learning language), one expects the neural networks \mathcal{U}_i and \mathcal{Z}_i to learn/approximate better and better the functions $u(t_i, \cdot)$ and $D_x u(t_i, \cdot)$ in view of the universal approximation theorem for neural networks. The rigorous convergence of this algorithm is postponed to a future work.

To sum up, the global algorithm is given in Algo 1 in the case where g is Lipschitz and the derivative can be analytically calculated almost everywhere. If the derivative of g is not available, it can be calculated by numerical differentiation.

Algorithm 1 Algorithm for full non linear equations.

1: Use a single deep neural network $(\mathcal{U}_N(\cdot; \theta), \mathcal{Z}_N(\cdot; \theta)) \in \mathcal{N}_{d,1+d,L,m}$ and minimize (by SGD)

$$\begin{cases} \hat{L}_N(\theta) := \mathbb{E} \left| \mathcal{U}_N(X_{t_N}; \theta) - g(X_{t_N}) \right|^2 + \mathbb{E} \left| \mathcal{Z}_N(X_{t_N}; \theta) - Dg(X_{t_N}) \right|^2 \\ \theta_N^* \in \arg \min_{\theta \in \mathbb{R}^{N_m}} \hat{L}_N(\theta). \end{cases} \quad (2.9)$$

2: $\hat{\mathcal{U}}_N = \mathcal{U}_N(\cdot; \theta_N^*)$, and set $\hat{\mathcal{Z}}_N = \mathcal{Z}_N(\cdot; \theta_N^*)$

3: **for** $i = N - 1, \dots, 0$ **do**

4: Use a single deep neural network $(\mathcal{U}_i(\cdot; \theta), \mathcal{Z}_i(\cdot; \theta)) \in \mathcal{N}_{d,1+d,L,m}$ for the approximation of $(u(t_i, \cdot), D_x u(t_i, \cdot))$, and compute (by SGD) the minimizer of the expected quadratic loss function

$$\begin{cases} \hat{L}_i(\theta) := \mathbb{E} \left| \hat{\mathcal{U}}_{i+1}(X_{t_{i+1}}) - F(t_i, X_{t_i}, \mathcal{U}_i(X_{t_i}; \theta), \mathcal{Z}_i(X_{t_i}; \theta), D\hat{\mathcal{Z}}_{i+1}(\mathcal{T}(X_{t_{i+1}})), \Delta t_i, \Delta W_{t_i}) \right|^2 \\ \theta_i^* \in \arg \min_{\theta \in \mathbb{R}^{N_m}} \hat{L}_i(\theta). \end{cases} \quad (2.10)$$

5: Update: $\hat{\mathcal{U}}_i = \mathcal{U}_i(\cdot; \theta_i^*)$, and set $\hat{\mathcal{Z}}_i = \mathcal{Z}_i(\cdot; \theta_i^*)$.

Remark 2.1 A variation in the algorithm consists in using two neural networks for $\hat{\mathcal{U}}_i$ and $\hat{\mathcal{Z}}_i$ instead of one. \square

Remark 2.2 The diffusion process X is used for the training simulations in the stochastic gradient descent method for finding the minimizer of the quadratic loss function in (2.10). The choice of the drift coefficient is typically related to the underlying probabilistic problem associated to the PDE (for example a stochastic control problem), but does not really matter. The choice of the diffusion coefficient σ is more delicate: large σ induces a better exploration of the state space, but would require a lot of neurons. Moreover, for the applications in stochastic control, we might explore some region that are visited with very small probabilities by the optimal state process, hence representing few interest. On the other hand, small σ means a weak exploration, and we might lack information and precision on some region of the state space. In practice and for the numerical examples in the next section, we test the scheme for different σ and by varying the number of time steps, and if it converges to the same solution, one can consider that we have obtained the correct solution. \square

3 Numerical results

We first construct an example with different non linearities in the Hessian term and the solution. We graphically show that the solution is very well calculated in dimension one and then move to higher dimensions. We next use an example derived from a stochastic optimization problem with an analytic solution and show that we are able to accurately calculate the solution.

In the whole numerical part, we use a classical Feed Forward Network using layers with n neurons each and a tanh activation function, the output layer uses an identity activation function. At each time step the resolution of equation (2.10) is achieved using some mini-batch with 1000 trajectories. Every 50 inner iterations the convergence rate is checked using 10000 trajectories and an adaptation of the learning rate is achieved using an Adams gradient descent. Notice that the adaptation of the learning rate is not common with the Adams method but in our case it appears to be crucial to have a steady diminution of the loss of the objective function.

During time resolution, it is far more effective to initialize the solution of equations (2.10) with the solution $(\mathcal{U}, \mathcal{Z})$ at the next time step and to use an initial learning rate $\ell_2^i \leq \ell_1^i$. The number of outer iterations is fixed for each optimization. It is set to 500 for the first optimization at date N and then a value of 100 outer iteration is used at the dates $i < N$. All experiments have achieved using Tensorflow [Aba+15]. In the sequel, the PDE solutions on curves are calculated as the average of 10 runs. Each time we provide the curves giving

the standard deviation observed for the given results. We also show the impact of the choice of the diffusion coefficient σ , and the influence of the number of neurons on the accuracy of the results.

3.1 A non linearity in uD_x^2u

We consider the case where $T = 1$,

$$f(t, x, y, z, \gamma) = y\text{tr}(\gamma) + \frac{y}{2} + 2y^2 - 2y^2 - 2y^4 e^{-(T-t)},$$

and $g(x) = \tanh\left(\frac{\sum_{i=1}^d x_i}{\sqrt{d}}\right)$, so that an analytical solution is available:

$$u(t, x) = \tanh\left(\frac{\sum_{i=1}^d x_i}{\sqrt{d}}\right) e^{-\frac{T-t}{2}}.$$

We choose to evaluate the solution at $t = 0$ and $x = 0.5\frac{1_d}{\sqrt{d}}$, for which $u(t, x) = 0.761902$ while its derivative is equal to 1.2966. This initial value x is chosen such that independently of the dimension the solution is varying around this point and not in a region where the tanh function is close to -1 or 1 .

The coefficients of the forward process used to solve the equation are

$$\sigma = \frac{\hat{\sigma}}{\sqrt{d}} \mathbf{I}_d, \quad \mu = 0,$$

and the truncation operator indexed by a parameter p is chosen equal to

$$\mathcal{T}_p(X_t^{0,x}) = (x - \sigma\sqrt{t}\phi_p) \vee X_t^{0,x} \wedge (x + \sigma\sqrt{t}\phi_p),$$

where $\phi_p = \mathcal{N}^{-1}(p)$, \mathcal{N} is the CDF of a unit centered Gaussian random variable. In the numerical results we take $p = 0.999$ and $n = 20$ neurons.

We first begin in dimension one, and show in figure 1 how u , $D_x u$ and $D_x^2 u$ are well approximated by the resolution method.

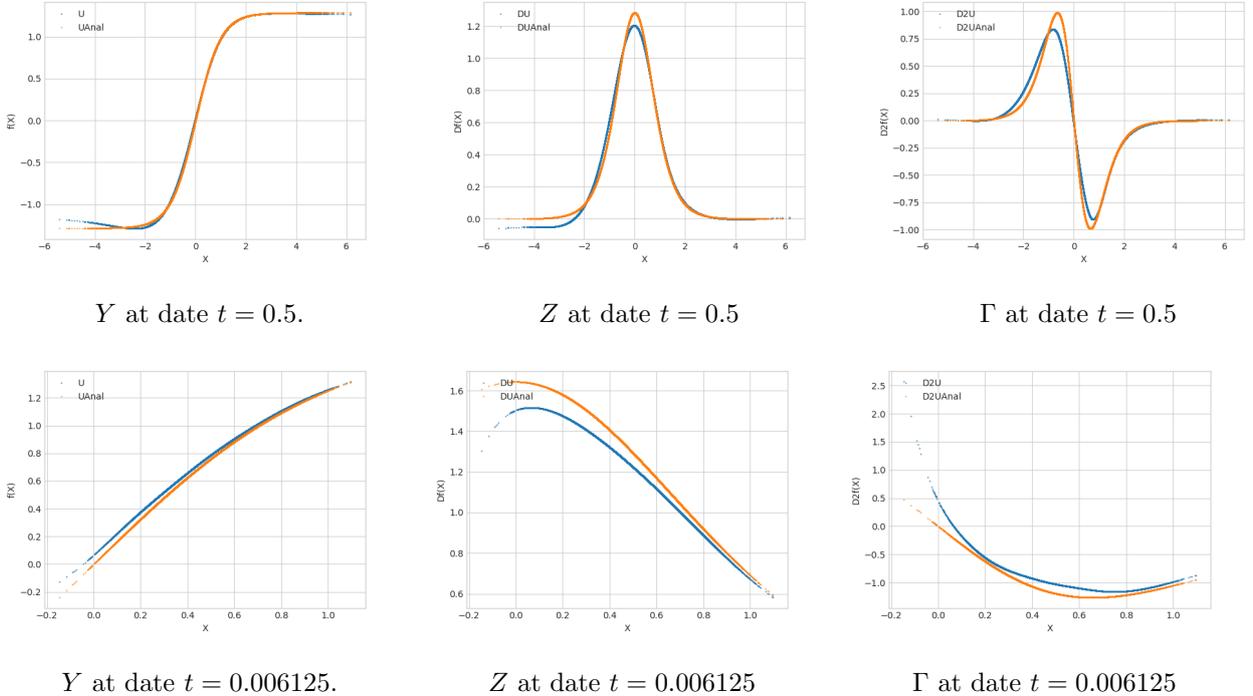
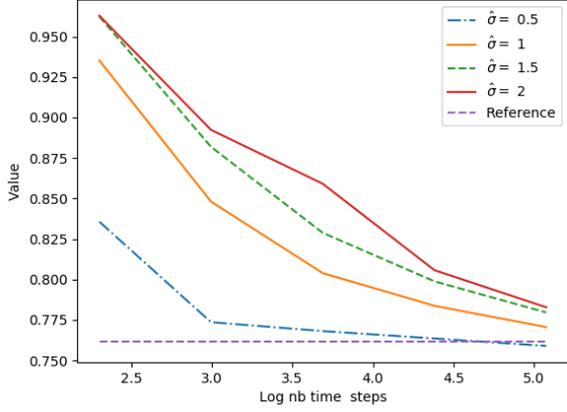
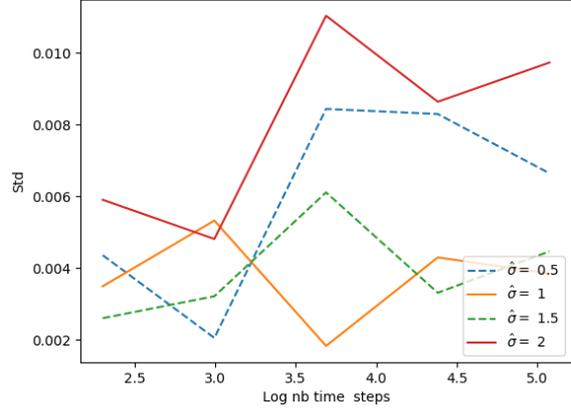


Figure 1: A single valuation run for test case one 1D using 160 time steps, $\hat{\sigma} = 2.$, $p = 0.999$, 20 neurons, 2 layers.

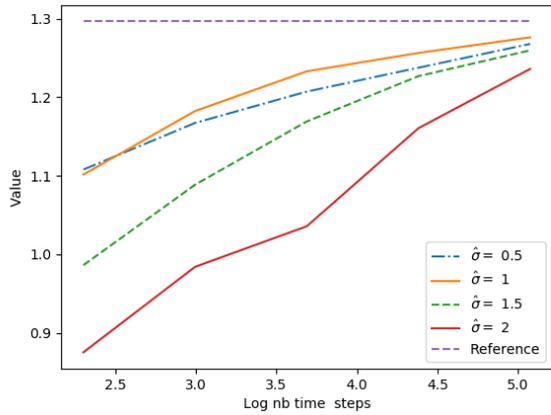
On figure 2, we check the convergence, for different values of $\hat{\sigma}$ of both the value function and its derivative at point x and date 0. Standard deviation of the function value is very low and the standard deviation of the derivative still being low.



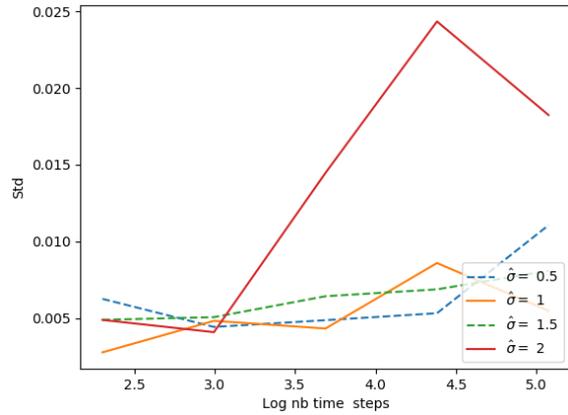
Convergence of u depending on $\hat{\sigma}$



Standard deviation of u



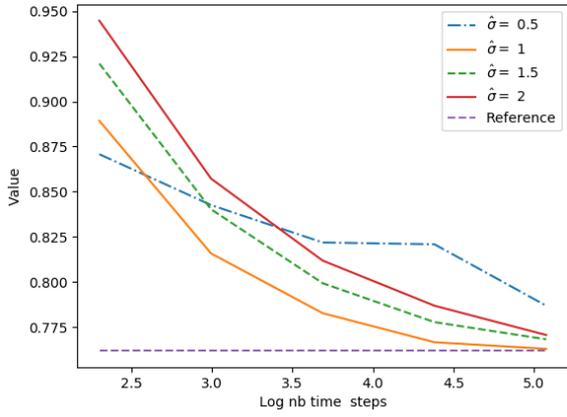
Convergence of $D_x u$ depending on $\hat{\sigma}$



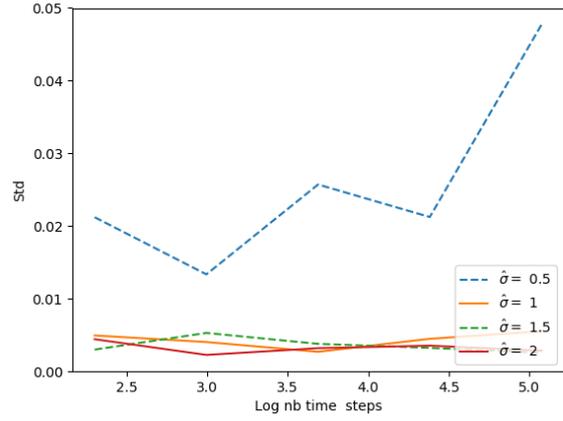
Standard deviation of $D_x u$

Figure 2: Convergence in 1D of the case one, number of neurons par layer equal to 20, 2 layers, $p = 0.999$.

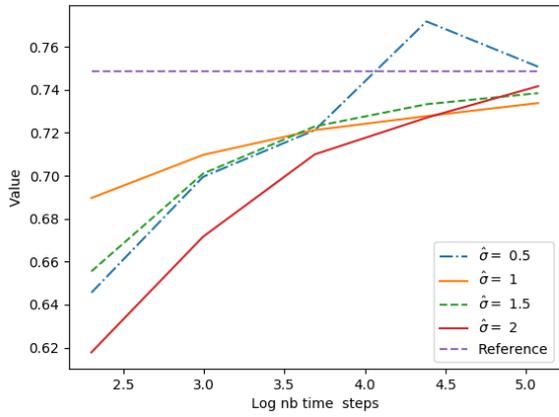
As the dimension increases, we have to increase the value of $\hat{\sigma}$ of the forward process. In dimension 3, the value $\hat{\sigma} = 0.5$ gives high standard deviation in the result obtained as shown on figure 3, while in dimension 10, see Figure 4, we see that the value $\hat{\sigma} = 1$ is too low to give good results. We also clearly notice that in 10D, a smaller time step should be used but in our test cases we decided to consider a maximum number of time steps equal to 160.



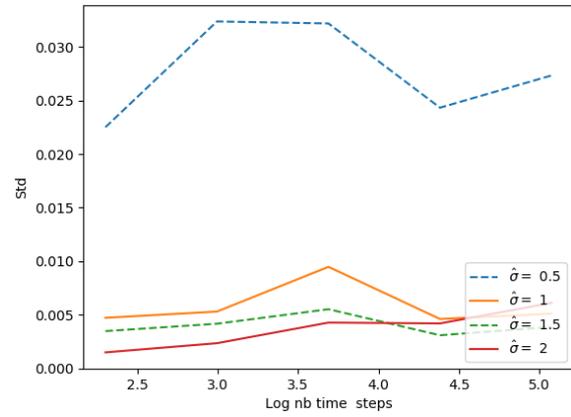
Convergence of u depending on $\hat{\sigma}$



Standard deviation of u

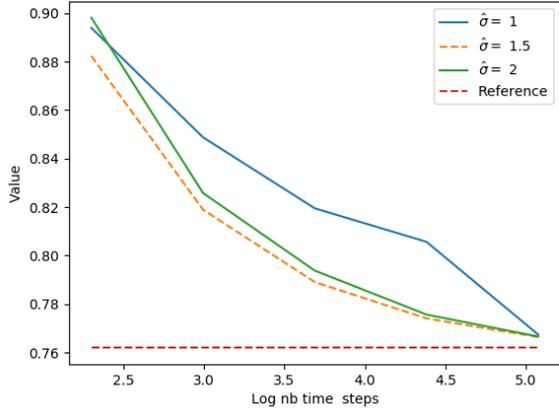


Convergence of $D_x u$ (first component) depending on $\hat{\sigma}$

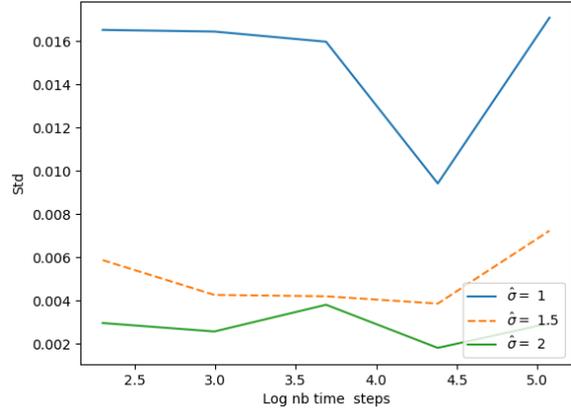


Standard deviation of $D_x u$ (first component)

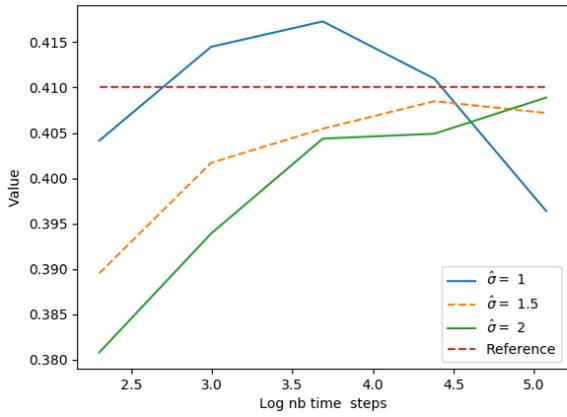
Figure 3: Convergence in 3D of the case one, number of neurons par layer equal to 20, 2 layers, $p = 0.999$



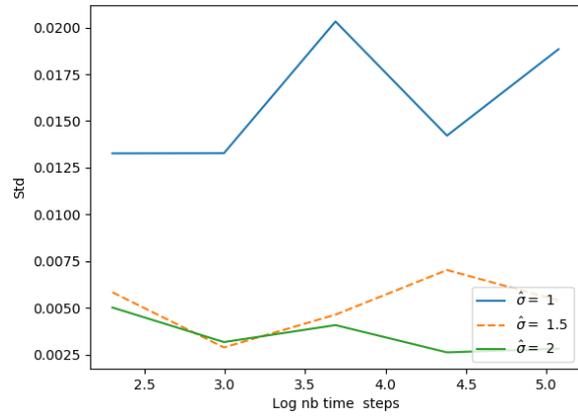
Convergence of u depending on $\hat{\sigma}$



Standard deviation of u



Convergence of $D_x u$ depending on $\hat{\sigma}$ (first component)



Standard deviation of $D_x u$ (first component)

Figure 4: Convergence in 10D of the case one, number of neurons par layer equal to 20, 2 layers, $p = 0.999$

On this simple test case, the dimension is not a problem and very good results are obtained in dimension 20 or above with only 20 neurons and 2 layers.

3.2 A linear quadratic stochastic test case.

In this example, we consider a controlled process $\mathcal{X} = \mathcal{X}^\alpha$ with dynamics in \mathbb{R}^d according to

$$d\mathcal{X}_t = (A\mathcal{X}_t + B\alpha_t)dt + D\alpha_t dW_t, \quad 0 \leq t \leq T, \quad \mathcal{X}_0 = x,$$

where W is a real Brownian motion, the control process α is valued in \mathbb{R} , and the constant coefficients $A \in \mathbb{M}^d$, $B \in \mathbb{R}^d$, $D \in \mathbb{R}^d$. The quadratic cost functional to be minimized is

$$J(\alpha) = \mathbb{E} \left[\int_0^T (\mathcal{X}_t^\top Q \mathcal{X}_t + \alpha_t^2 N) dt + \mathcal{X}_T^\top P \mathcal{X}_T \right],$$

where P, Q are non negative $d \times d$ symmetric matrices and $N \in \mathbb{R}$ is strictly positive.

The Bellman equation associated to this stochastic control problem is:

$$\begin{aligned} \frac{\partial u}{\partial t} + \inf_{a \in \mathbb{R}} [(Ax + Ba) \cdot D_x u + \frac{a^2}{2} \text{tr}(DD^\top D_x^2 u) + x^\top Q x + Na^2] &= 0, \quad (t, x) \in [0, T) \times \mathbb{R}^d, \\ u(T, x) &= x^\top P x, \quad x \in \mathbb{R}^d, \end{aligned}$$

which can be rewritten as a fully nonlinear equation in the form (1.1) with

$$f(t, x, y, z, \gamma) = x^\top Q x + Ax \cdot z - \frac{1}{2} \frac{|B^\top z|^2}{\text{tr}(DD^\top \gamma) + 2N}.$$

An explicit solution to this PDE is given by

$$u(t, x) = x^\top K(t)x,$$

where $K(t)$ is non negative $d \times d$ symmetric matrix function solution to the Riccati equation:

$$\dot{K} + A^\top K + KA + Q - \frac{KBB^\top K}{N + D^\top KD} = 0, \quad K(T) = P.$$

We take $T = 1$. The coefficients of the forward process used to solve the equation are

$$\sigma = \frac{\hat{\sigma}}{\sqrt{d}} \mathbf{I}_d, \quad \mu(t, x) = Ax.$$

In our numerical example we take the following parameters for the optimization problem:

$$A = B = D = \mathbf{I}_d, \quad Q = P = \frac{1}{d} \mathbf{I}_d, \quad N = d$$

and we want to estimate the solution at $x = \mathbf{I}_d$.

In this example, the truncation operator (indexed by p between 0 and 1 and close to 1) is as follows:

$$\mathcal{T}_p(X_t^x) = (xe^{\hat{A}t} - \sigma \sqrt{\frac{e^{2\hat{A}t} - \hat{1}}{2\hat{A}}} \phi_p) \vee X_t^x \wedge (xe^{\hat{A}t} + \sigma \sqrt{\frac{e^{2\hat{A}t} - \hat{1}}{2\hat{A}}} \phi_p),$$

where $\phi_p = \mathcal{N}^{-1}(p)$, \hat{A} is a vector so that $\hat{A}_i = A_{ii}$, $i = 1, \dots, d$, $\hat{1}$ is a unit vector, and the square root is taken componentwise.

On figure 5 we give the solution of the PDE using $\hat{\sigma} = 1.5$ obtained for two dates: at $t = 0.5$ and at t close to zero. We observe that we have a very good estimation of the function value and a correct one of the Γ value at date $t = 0.5$. The precision remains good for Γ close to $t = 0$ and very good for u and $D_x u$.

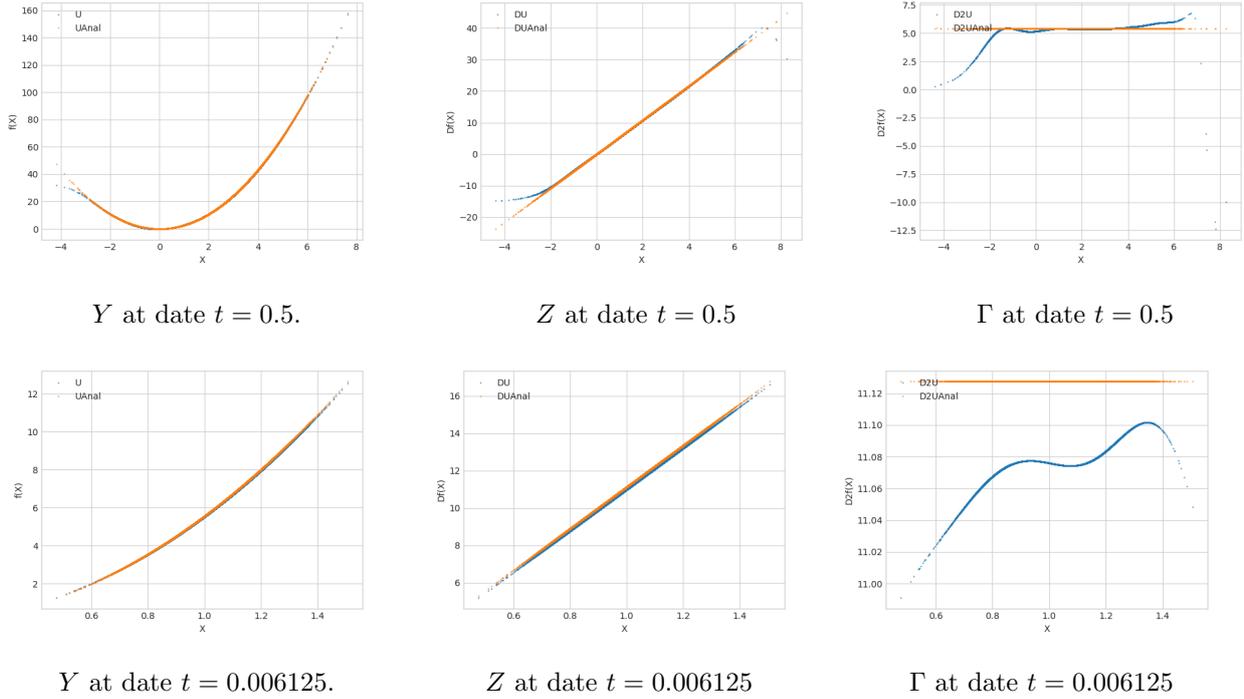


Figure 5: Test case linear quadratic 1D using 160 time steps, $\hat{\sigma} = 1.5$, $p = 0.999$, 100 neurons.

On figures 6, we give the results obtained in dimension one by varying $\hat{\sigma}$. For a value of $\hat{\sigma} = 2$, the standard deviation of the result becomes far higher than with $\hat{\sigma} = 0.5$ or 1.

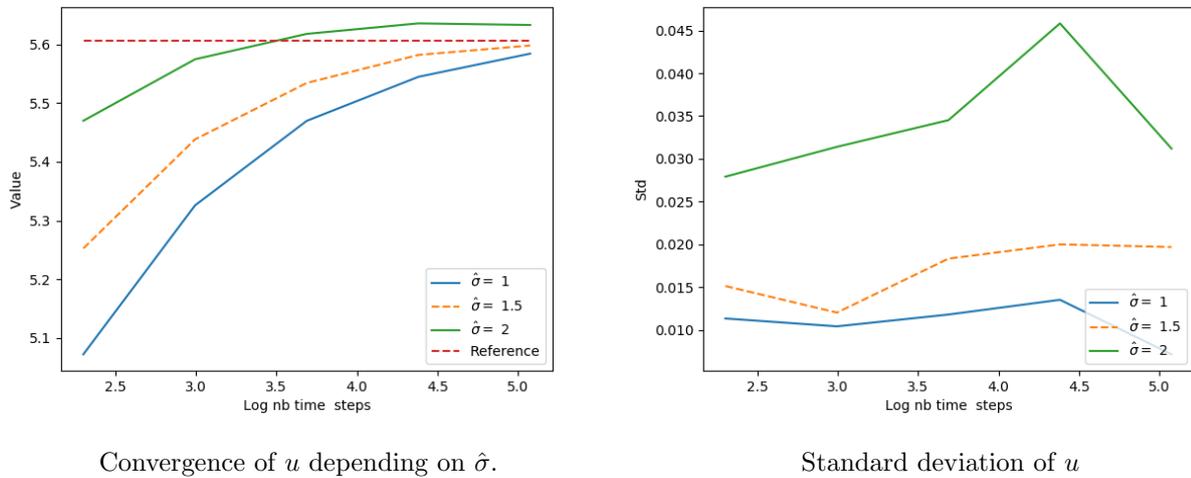
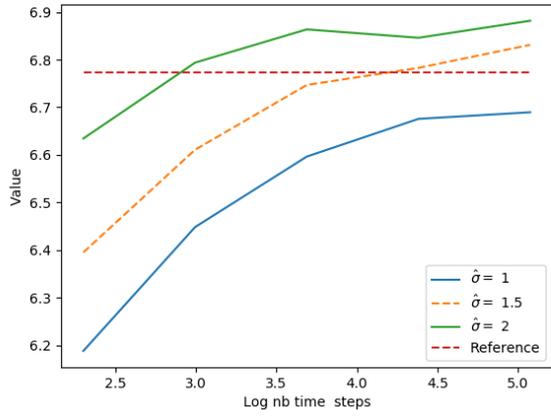
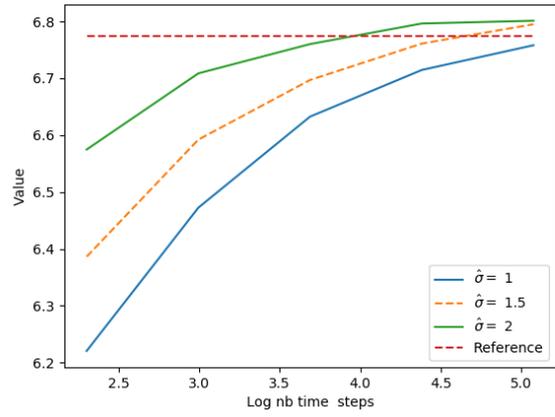


Figure 6: Convergence in 1D of the linear quadratic case, number of neurons par layer equal to 50, 2 layers, $p = 0.999$

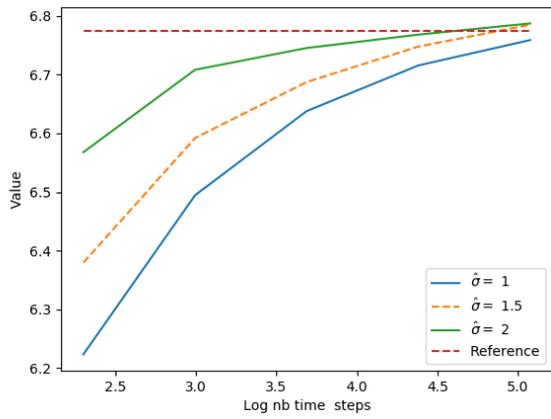
On figure 7, we take a quite low truncation factor $p = 0.95$ and observe that the number of neurons to take has to be rather high. We have also checked that taking a number of hidden layers equal to 3 does not improve the results.



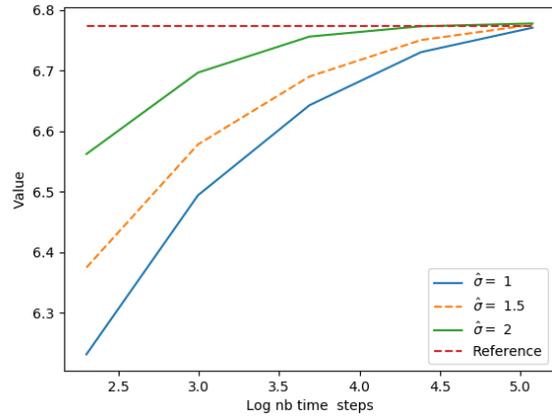
10 neurons



20 neurons



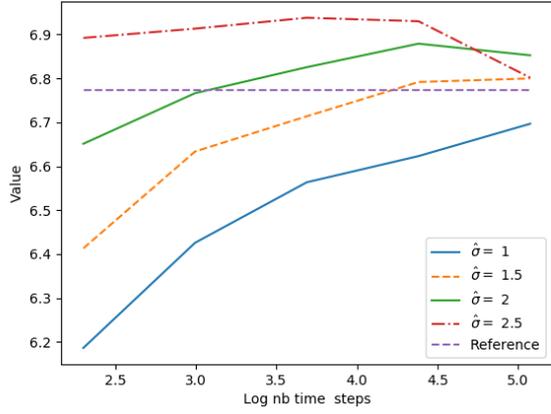
30 neurons



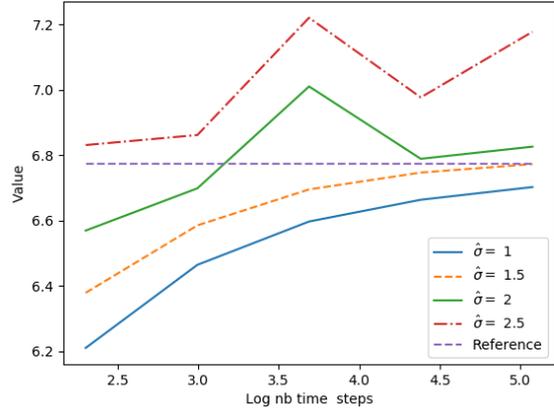
50 neurons

Figure 7: Convergence in 3D of the linear quadratic case, 2 layers, testing the influence of the number of neurons, truncation $p = 0.95$.

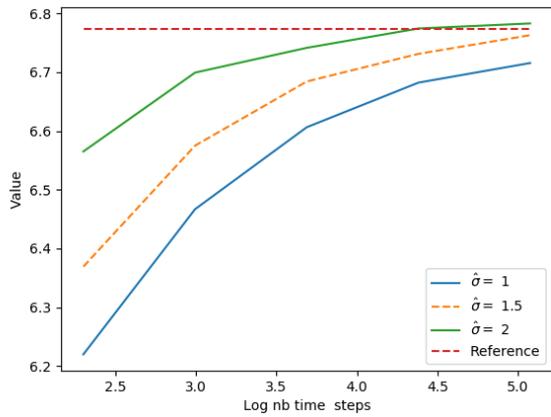
On figure 8, we give the same graphs for a truncation factor higher. As we take a higher truncation factor the number of neurons to use has to be increased to 100.



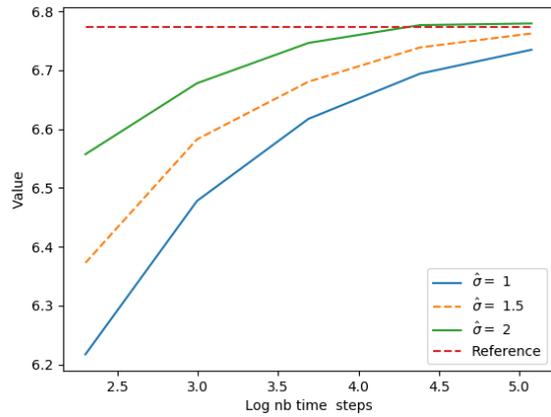
10 neurons



20 neurons



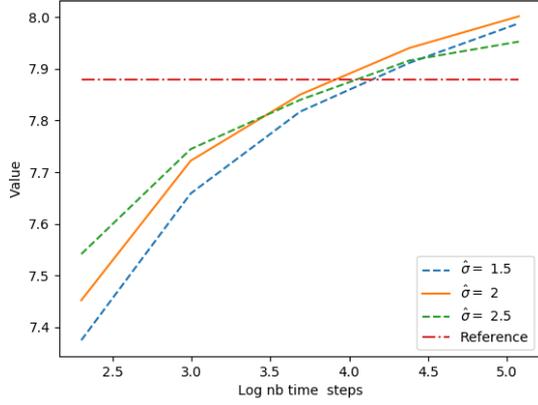
50 neurons



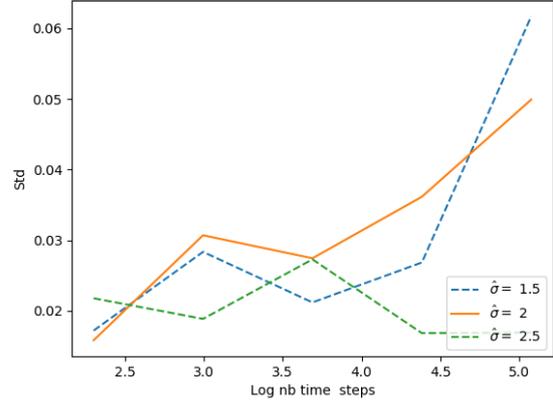
100 neurons

Figure 8: Convergence in 3D of the linear quadratic case, 2 layers, testing the influence of the number of neurons, truncation $p = 0.99$.

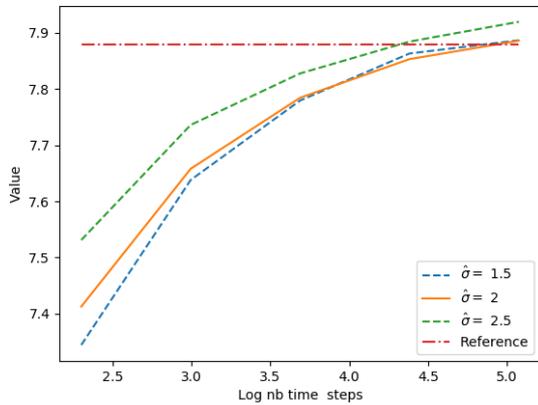
On figure 9, we observe in dimension 7 the influence of the number of neurons on the result for a high truncation factor $p = 0.999$. With a number of neurons equal to 50, we clearly have a bias disappearing with a number of neurons equal to 100. We had to take higher values of $\hat{\sigma}$ to get good results.



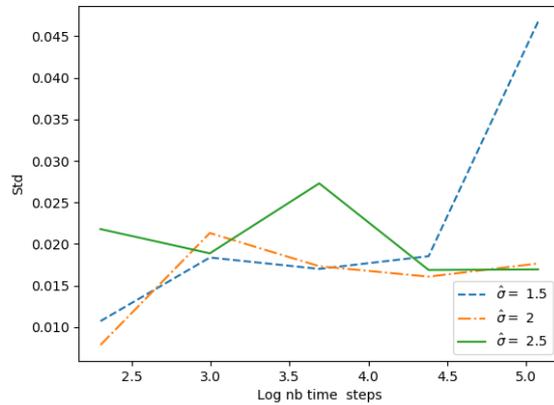
Convergence with 50 neurons



Standard deviation with 50 neurons



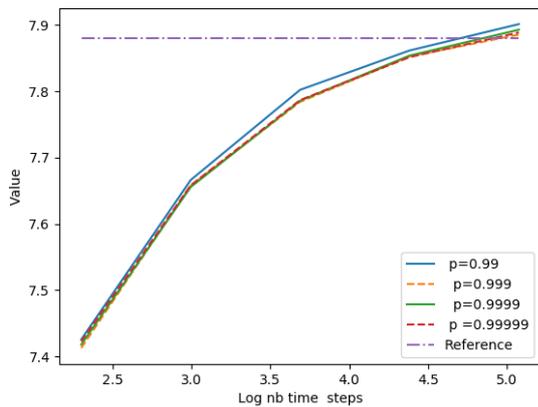
Convergence with 100 neurons



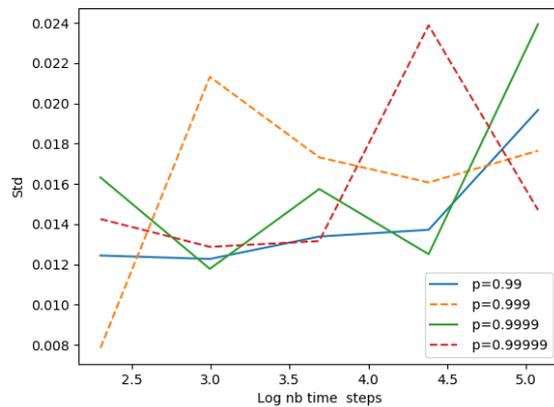
Standard deviation with 100 neurons

Figure 9: Convergence in 7D of the linear quadratic case, 2 layers, $p = 0.999$

On figure 10, we check that influence of the truncation factor appears to be slow for higher dimensions.



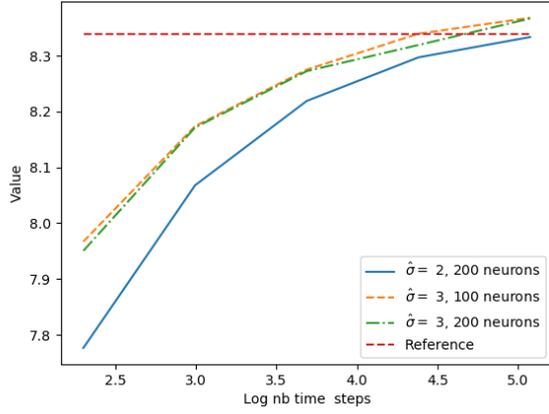
Function value



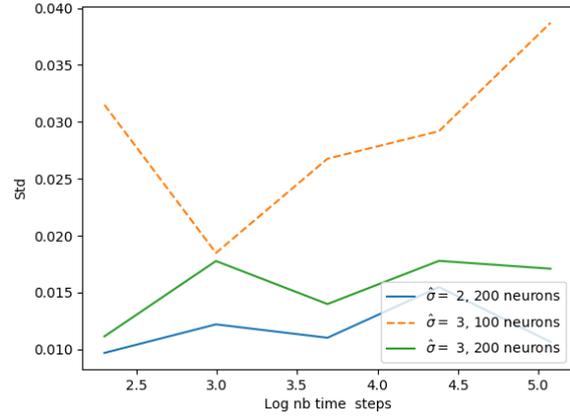
Standard deviation

Figure 10: Convergence in 7D of the linear quadratic case with 2 layers, 100 neurons, testing p

Finally, we give results in dimension 10, 15 and 20 for $p = 0.999$ on figures 11, 12, 13. We observe that the number a neurons with 2 hidden layers has to increase with the dimension but also that the increase is rather slow in contrast with the case of one hidden layer as theoretically shown in [Pin99]. For $\hat{\sigma} = 5$ we had to take 300 neurons to get very accurate results.

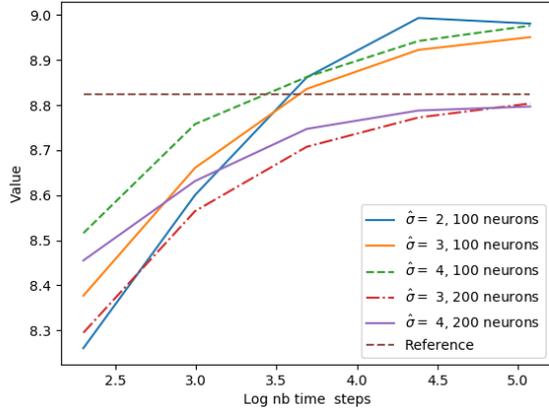


Function value

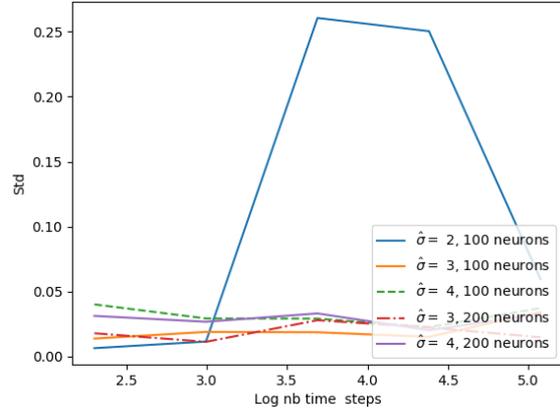


Standard deviation

Figure 11: Convergence in 10D of the linear quadratic case with 2 layers, $p = 0.999$

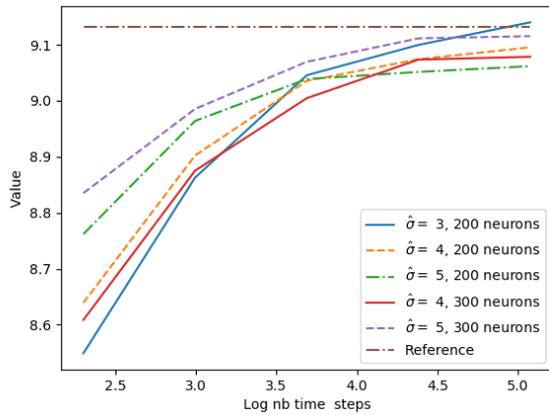


Function value

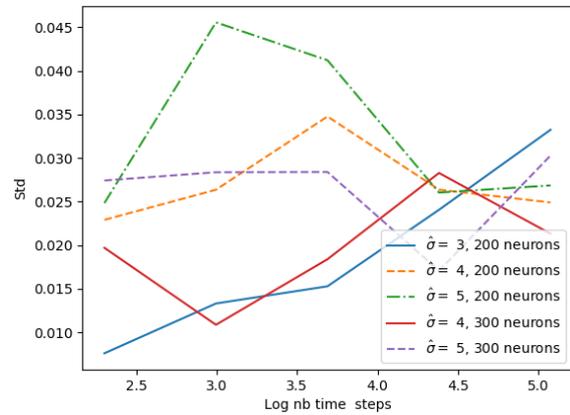


Standard deviation

Figure 12: Convergence in 15D of the linear quadratic case with 2 layers, $p = 0.999$



Function value



Standard deviation

Figure 13: Convergence in 20D of the linear quadratic case with 2 layers, $p = 0.999$

References

- [AA+18] Ali Al-Aradi et al. “Solving Nonlinear and High-Dimensional Partial Differential Equations via Deep Learning”. In: *arXiv preprint arXiv:1811.08782* (2018).
- [Aba+15] Martin Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. URL: <https://www.tensorflow.org/>.
- [Bec+19] Christian Beck et al. “Deep splitting method for parabolic PDEs”. In: *arXiv preprint arXiv:1907.03452* (2019).
- [BEJ17] C. Beck, W. E, and Arnulf Jentzen. “Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations”. In: *CoRR* abs/1709.05963 (2017).
- [Che+07] Patrick Cheridito et al. “Second-order backward stochastic differential equations and fully nonlinear parabolic PDEs”. In: *Communications on Pure and Applied Mathematics* 60.7 (2007), pp. 1081–1110.
- [FTW11] A. Fahim, N. Touzi, and X. Warin. “A probabilistic numerical method for fully nonlinear parabolic PDEs”. In: *The Annals of Applied Probability* (2011), pp. 1322–1364.
- [HPW19] Côme Huré, Huyên Pham, and Xavier Warin. “Some machine learning schemes for high-dimensional nonlinear PDEs”. In: *arXiv preprint arXiv:1902.01599* (2019).
- [Pin99] Allan Pinkus. “Approximation theory of the MLP model in neural networks”. In: *Acta numerica* 8 (1999), pp. 143–195.
- [SS18] J. Sirignano and K. Spiliopoulos. “DGM: A deep learning algorithm for solving partial differential equations”. In: *Journal of Computational Physics* 375 (2018), pp. 1339–1364.
- [Tan+13] Xiaolu Tan et al. “A splitting method for fully nonlinear degenerate parabolic PDEs”. In: *Electronic Journal of Probability* 18 (2013).
- [War18] X. Warin. “Monte Carlo for high-dimensional degenerated Semi Linear and Full Non Linear PDEs”. In: *arXiv preprint arXiv:1805.05078* (2018).

Finance for Energy Market Research Centre

Institut de Finance de Dauphine, Université Paris-Dauphine

1 place du Maréchal de Lattre de Tassigny

75775 PARIS Cedex 16

www.fime-lab.org