



HAL
open science

DeepSets and their derivative networks for solving symmetric PDEs *

Maximilien Germain, Mathieu Laurière, Huyên Pham, Xavier Warin

► To cite this version:

Maximilien Germain, Mathieu Laurière, Huyên Pham, Xavier Warin. DeepSets and their derivative networks for solving symmetric PDEs *. *Journal of Scientific Computing*, 2022, <10.1007/s10915-022-01796-w>. <hal-03154116v2>

HAL Id: hal-03154116

<https://hal.science/hal-03154116v2>

Submitted on 3 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

DeepSets and their derivative networks for solving symmetric PDEs *

Maximilien GERMAIN[†] Mathieu LAURIÈRE[‡] Huyên PHAM[§]
Xavier WARIN[¶]

January 3, 2022

to appear in Journal of Scientific Computing

Abstract

Machine learning methods for solving nonlinear partial differential equations (PDEs) are hot topical issues, and different algorithms proposed in the literature show efficient numerical approximation in high dimension. In this paper, we introduce a class of PDEs that are invariant to permutations, and called *symmetric* PDEs. Such problems are widespread, ranging from cosmology to quantum mechanics, and option pricing/hedging in multi-asset market with exchangeable payoff. Our main application comes actually from the particles approximation of mean-field control problems. We design deep learning algorithms based on certain types of neural networks, named PointNet and DeepSet (and their associated derivative networks), for computing simultaneously an approximation of the solution and its gradient to symmetric PDEs. We illustrate the performance and accuracy of the PointNet/DeepSet networks compared to classical feedforward ones, and provide several numerical results of our algorithm for the examples of a mean-field systemic risk, mean-variance problem and a min/max linear quadratic McKean-Vlasov control problem.

Keywords: Permutation-invariant PDEs, symmetric neural networks, exchangeability, deep backward scheme, mean-field control.

1 Introduction

The numerical resolution of partial differential equations (PDEs) in high dimension is a major challenge in various areas of science, engineering, and finance. PDEs that appear in the applications are often non linear and of very high dimension (number of particles in physics, number of agents in large population control problems, number of assets and factors in financial markets, etc), and are subject to the so-called curse of dimensionality, which makes infeasible the implementation of classical grid methods and Monte-Carlo approaches.

*This work was supported by FiME (Finance for Energy Market Research Centre) and the “Finance et Développement Durable - Approches Quantitatives” EDF - CACIB Chair.

[†]EDF R&D, and LPSM, Université de Paris mgermain at lpsm.paris

[‡]ORFE, Princeton University lauriere at princeton.edu

[§]LPSM, Université de Paris, and FiME, and CREST ENSAE pham at lpsm.paris

[¶]EDF R&D, and FiME xavier.warin at edf.fr

A breakthrough with deep learning based-algorithms has been made in the last five years towards this computational challenge, and we mention the recent survey papers by [Bec+20] and [GPW22b]. The main interest in the use of machine learning techniques for PDEs is the ability of deep neural networks to efficiently represent high dimensional functions without using spatial grids, and with no curse of dimensionality (see e.g. [Hut+20]). Although the use of neural networks for solving PDEs is not new, the approach has been successfully revived with new ideas and directions. Moreover, recently developed open source libraries like Tensorflow and Pytorch offer an accessible framework to implement these algorithms.

In this paper, we introduce a class of PDEs that are invariant by permutation, and called here *symmetric* PDEs. Such PDEs occur naturally in the modelling of systems dealing with sets that are invariant by permutation of their elements. Applications range from models in general relativity and cosmology, to quantum mechanics and chemistry, see e.g. [Wyk08], [Smu11]. Symmetric PDEs also appear in the pricing/hedging of basket option and options on the maximum of multiple assets. Our main motivation for introducing this general class of symmetric PDEs comes from the control of large population of interacting indistinguishable agents, which leads in the asymptotic regime of infinite population to the theory of mean-field games (MFG) and mean-field type control, also called McKean-Vlasov (MKV) control. These topics have attracted an increasing and large interest since the seminal papers [LL07] and [HCM06] with important mathematical developments and numerous applications in various fields over the last decade. We refer to the two-volume monographs [CD18a]-[CD18b] for an exhaustive exposition of this research domain, where it is known that the solution to MFG or MKV control problem are characterized in terms of a Master equation or a Bellman equation, which are PDEs in the Wasserstein space of probability measures. It turns out that the finite-dimensional approximation of these equations are formulated as *symmetric* non linear PDEs, and the convergence of this approximation has been recently obtained in [GMS21], and [GPW22a] (for a rate of convergence), see also [Lac17] and [Dje20].

Symmetric PDEs are often in very high dimension, typically of the order of one thousand in the case of particles approximation of Master and Bellman equations, and it is tempting to apply machine learning algorithms in this framework. For that purpose, we shall furthermore exploit the symmetric structure of the PDEs in order to design deep learning-based algorithms with a suitable class of neural networks. We first observe that the solution to symmetric PDEs is invariant by permutation (also called *exchangeable*), and we shall then consider a class of symmetric neural networks, named PointNet and DeepSets, aiming to approximate our solution. Such class of neural networks has been recently introduced in the machine learning community, see [RQ+17], [Zah+17], [BRT20], for dealing with tasks involving some invariant data sets, and it turns out that they provide much better accuracy than classical feedforward neural networks (NN in short) in the approximation of symmetric functions. Indeed, feedforward NN have too many degrees of freedom, and the optimization over parameters in (stochastic) gradient descent algorithm may be trapped away in the approximation of a symmetric function, as illustrated in several examples and comparison tests presented in this paper. We shall also introduce different classes of derivative symmetric network, named DeepDerSet and AD-DeepSet, for the approximation of the gradient of the solution to symmetric PDEs.

By relying on the class of symmetric NN, and their derivative networks, we next adapt the deep backward dynamic programming scheme [HPW20], [PWG21] for numerically solving symmetric PDEs, i.e., finding approximations of the solution and its gradient. We also

explain in the case of mean-field control problem how our scheme provides an approximation for the solution to a Bellman equation in the Wasserstein space of probability measures. This yields alternative deep learning schemes for mean-field control problems to the ones recently designed in [GMW19], [CL19], [FZ20], or [Rut+20]. We test our algorithms on several examples arising from different McKean-Vlasov control problem, for which we have explicit or benchmarked solutions: a systemic risk model as in [CFS15], the classical mean-variance, i.e., Markowitz portfolio allocation problem, and a min/max linear quadratic mean-field control problem as in [SMLN15].

Outline of the paper. The rest of the paper is organized as follows. We introduce in Section 2 the class of *symmetric* PDEs with some examples, and show exchangeability properties of the solution and its gradient to such PDEs. Section 3 is devoted to the exposition of the class of symmetric neural networks, as well as its derivative networks, and we provide several comparison tests with respect to classical feedforward NN. We describe in Section 4 the deep learning schemes for solving symmetric PDEs, and finally provide several numerical examples in Section 5.

Notations. Given $N \in \mathbb{N}^*$, \mathcal{X}^N denotes the set of all elements $\mathbf{x} = (x_i)_{i \in [1, N]}$ with coefficients x_i valued in \mathcal{X} and $[1, N] = \{1, \dots, N\}$. When $\mathcal{X} = \mathbb{R}^d$, one usually identifies $(\mathbb{R}^d)^N$ with $\mathbb{R}^{d \times N}$ the set of $d \times N$ -matrices with real-valued coefficients. $\mathbb{S}^N(\mathcal{X})$ is the set of $N \times N$ -symmetric matrices with coefficients valued in \mathcal{X} , and is simply denoted by \mathbb{S}^N when $\mathcal{X} = \mathbb{R}$. For a real-valued C^2 function φ defined on $(\mathbb{R}^d)^N$, its gradient $D\varphi(\mathbf{x}) = (D_{x_i}\varphi(\mathbf{x}))_{i \in [1, N]}$ is valued in $(\mathbb{R}^d)^N$, while its Hessian $D^2\varphi(\mathbf{x}) = (D_{x_i x_j}^2\varphi(\mathbf{x}))_{i, j \in [1, N]}$ is valued in $\mathbb{S}^N(\mathbb{S}^d)$.

We denote by \mathfrak{S}_N the set of permutations on $\{1, \dots, N\}$. For any $\mathbf{x} = (x_i)_{i \in [1, N]} \in \mathcal{X}^N$, $\pi \in \mathfrak{S}_N$, we denote by $\pi[\mathbf{x}] = (x_{\pi(i)})_{i \in [1, N]} \in \mathcal{X}^N$. For any $\mathbf{\Gamma} = (\Gamma_{ij})_{i, j \in [1, N]} \in \mathbb{S}^N(\mathcal{X})$, we denote by $\pi[\mathbf{\Gamma}] = (\Gamma_{\pi(i)\pi(j)})_{i, j \in [1, N]} \in \mathbb{S}^N(\mathcal{X})$.

We say that a function φ defined on \mathcal{X}^N is exchangeable to the order N on \mathcal{X} if it is invariant by permutation, i.e., $\varphi(\mathbf{x}) = \varphi(\pi[\mathbf{x}])$, for any $\mathbf{x} \in \mathcal{X}^N$, and $\pi \in \mathfrak{S}_N$. We may simply say exchangeable when it is clear from the context. The notations $0_d, 1_d$ refer respectively to d -dimensional vectors full of 0 and 1. With two vectors $a, b \in \mathbb{R}^d$, the notation $a.b = \sum_{i=1}^d a_i b_i$ refers to the canonical scalar product.

2 Symmetric PDEs

We consider a so-called *symmetric* class of parabolic second-order partial differential equations (PDEs):

$$\begin{cases} \partial_t v + F(t, \mathbf{x}, v, D_{\mathbf{x}}v, D_{\mathbf{x}}^2v) = 0, & (t, \mathbf{x}) \in [0, T] \times (\mathbb{R}^d)^N \\ v(T, \mathbf{x}) = G(\mathbf{x}), & \mathbf{x} \in (\mathbb{R}^d)^N, \end{cases} \quad (2.1)$$

where F is a real-valued function defined on $[0, T] \times (\mathbb{R}^d)^N \times \mathbb{R} \times (\mathbb{R}^d)^N \times \mathbb{S}^N(\mathbb{S}^d)$, G is defined on $(\mathbb{R}^d)^N$, and satisfying the permutation-invariance condition:

(HI) For any $t \in [0, T]$, $\mathbf{x} \in (\mathbb{R}^d)^N$, $y \in \mathbb{R}$, $\mathbf{z} \in (\mathbb{R}^d)^N$, $\gamma \in \mathbb{S}^N(\mathbb{S}^d)$,

$$\begin{aligned} F(t, \mathbf{x}, y, \mathbf{z}, \gamma) &= F(t, \pi[\mathbf{x}], y, \pi[\mathbf{z}], \pi[\gamma]) \\ G(\mathbf{x}) &= G(\pi[\mathbf{x}]), \quad \forall \pi \in \mathfrak{S}_N. \end{aligned}$$

We assume that PDE (2.1) is well-posed in the sense that there exists a unique classical solution satisfying a suitable growth condition.

We list below some examples of symmetric PDEs in the form (2.1). We start with an example of pricing in a “symmetric” multi-asset model.

Example 2.1 (*Multi-asset pricing*). Let us consider a model with N risky assets of price process $\mathbf{X} = (X^1, \dots, X^N)$ governed by

$$dX_t^i = \tilde{b}_i(\mathbf{X}_t)dt + \sum_{j=1}^N \sigma_{ij}(\mathbf{X}_t)dW_t^j,$$

where the diffusion coefficients satisfy the property: for all $\pi \in \mathfrak{S}_N$,

$$\sigma_{ij}(\pi[\mathbf{x}]) = \sigma_{\pi(i)\pi(j)}(\mathbf{x}), \quad \mathbf{x} = (x_i)_{i \in \llbracket 1, N \rrbracket}, \quad i, j = 1, \dots, N. \quad (2.2)$$

Notice that \tilde{b}_i is the drift of the asset price under the historical probability measure, and does not appear in the pricing equations below. The symmetry condition (2.2) is satisfied for example when $\sigma_{ii}(\mathbf{x}) = \sigma(\mathbf{x})$, and $\sigma_{ij}(\mathbf{x}) = \tilde{\sigma}(\mathbf{x})$, $i, j = 1, \dots, N$, $i \neq j$, with $\sigma, \tilde{\sigma}$ exchangeable functions. Another example is when $\sigma_{ii}(\mathbf{x}) = \sigma(x_i)$, and $\sigma_{ij}(\mathbf{x}) = \tilde{\vartheta}(x_i)\bar{\vartheta}(x_j)$, $i, j = 1, \dots, N$, $i \neq j$, for some functions $\sigma, \tilde{\vartheta}, \bar{\vartheta}$ defined on \mathbb{R} , which means that all the assets have the same marginal volatility coefficient, and the correlation function between any pair of assets is identical. We consider an option of maturity T with payoff $G(X_T^1, \dots, X_T^N)$, where G is an exchangeable function, for example:

$$G(\mathbf{x}) = \begin{cases} (\max(x_1, \dots, x_N) - K)_+, & \text{(call on max)} \\ (\sum_{i=1}^N x_i - K)_+, & \text{(call on sum),} \\ \sum_{i=1}^N 1_{x_i \geq K}, & \text{(sum of binary options),} \end{cases}$$

for $\mathbf{x} = (x_1, \dots, x_N) \in \mathbb{R}^N$. In a frictionless market with constant interest rate r , the option price $(t, \mathbf{x}) \in [0, T] \times \mathbb{R}^N \mapsto v(t, \mathbf{x})$ satisfies a linear PDE (2.1) with terminal condition given by the exchangeable function G and

$$F(t, \mathbf{x}, y, \mathbf{z}, \boldsymbol{\gamma}) = -ry + r \sum_{i=1}^N x_i z_i + \frac{1}{2} \sum_{i,j=1}^N \sigma_{ij}^2(\mathbf{x}) \gamma_{ij},$$

for $t \in [0, T]$, $\mathbf{x} = (x_i)_{i \in \llbracket 1, N \rrbracket} \in \mathbb{R}^N$, $y \in \mathbb{R}$, $\mathbf{z} = (z_i)_{i \in \llbracket 1, N \rrbracket} \in \mathbb{R}^N$, and $\boldsymbol{\gamma} = (\gamma_{ij})_{i,j \in \llbracket 1, N \rrbracket} \in \mathbb{S}^N$. In the case of counterparty risk, the pricing of CVA leads to a quasi-linear PDE (2.1) with F in the form (see [Lab12] for the details of the PDE derivation):

$$F(t, \mathbf{x}, y, \mathbf{z}, \boldsymbol{\gamma}) = \beta(y^+ - y) + r \sum_{i=1}^N x_i z_i + \frac{1}{2} \sum_{i,j=1}^N \sigma_{ij}^2(\mathbf{x}) \gamma_{ij},$$

where $\beta > 0$ is the intensity of default. Another case of non-linearity occurs when lending rate $r > 0$ is smaller than borrowing rate $R > 0$, which leads to a super-replication price solution to (2.1) with F given by

$$F(t, \mathbf{x}, y, \mathbf{z}, \boldsymbol{\gamma}) = \sup_{b \in [r, R]} [-by + b \sum_{i=1}^N x_i z_i] + \frac{1}{2} \sum_{i,j=1}^N \sigma_{ij}^2(\mathbf{x}) \gamma_{ij}.$$

In the above three cases, and under (2.2), the generator function F clearly satisfies the permutation-invariance condition in **(HI)**. \square

The second example is actually our main motivation for considering symmetric PDEs, and comes from mean-field models.

Example 2.2 (*McKean-Vlasov control problem with common noise*). Let us consider N interacting indistinguishable agents with controlled state process $\mathbf{X} = (X^1, \dots, X^N)$ valued in $(\mathbb{R}^d)^N$, and driven by

$$\begin{aligned} dX_t^i &= \beta(t, X_t^i, \bar{\mu}(\mathbf{X}_t), \alpha_t^i)dt + \sigma(t, X_t^i, \bar{\mu}(\mathbf{X}_t), \alpha_t^i)dW_t^i \\ &\quad + \sigma_0(t, X_t^i, \bar{\mu}(\mathbf{X}_t))dW_t^0, \quad 0 \leq t \leq T, \quad i = 1, \dots, N, \end{aligned}$$

where $\mathbf{x} = (x_i)_{i \in \llbracket 1, N \rrbracket} \mapsto \bar{\mu}(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}$ is the empirical measure (exchangeable) function, W^i , $i = 1, \dots, N$, are independent Brownian motions representing idiosyncratic noises, and W^0 is a Brownian motion independent of $\mathbf{W} = (W^i)_{i \in \llbracket 1, N \rrbracket}$, representing a common noise. Moreover, α^i is a control process (valued in some Polish space A) applied by the agent i who follows in a cooperative equilibrium a social planner aiming to minimize a social cost in the form

$$J(\alpha^1, \dots, \alpha^N) = \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\int_0^T e^{-rt} f(X_t^i, \bar{\mu}(\mathbf{X}_t), \alpha_t^i) dt + e^{-rT} g(X_T^i, \bar{\mu}(\mathbf{X}_T)) \right].$$

The Bellman equation to this N -cooperative agents control problem is in the form (2.1) with a Hamiltonian function F given by

$$\begin{aligned} F(t, \mathbf{x}, y, \mathbf{z}, \gamma) &= \sum_{i=1}^N \inf_{a \in A} [\beta(t, x_i, \bar{\mu}(\mathbf{x}), a) \cdot z_i + \frac{1}{2} \text{tr}(\Sigma(t, x_i, \bar{\mu}(\mathbf{x}), a) \gamma_{ii}) + \frac{1}{N} f(x_i, \bar{\mu}(\mathbf{x}), a)] \\ &\quad + \frac{1}{2} \sum_{i \neq j} \text{tr}(\sigma_0(t, x_i, \bar{\mu}(\mathbf{x})) \sigma_0^\top(t, x_j, \bar{\mu}(\mathbf{x})) \gamma_{ij}) - ry, \end{aligned}$$

where $\Sigma = \sigma \sigma^\top + \sigma_0 \sigma_0^\top$, and a terminal condition given by

$$G(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N g(x_i, \bar{\mu}(\mathbf{x})).$$

Such functions F and G clearly satisfy condition **(HI)**. Let us point out that in the limiting regime when the number N of agents goes to infinity, it is proved in [Lac17], [Dje20], [GMS21] that the solution to this cooperative-agents problem converges to the McKean-Vlasov control problem with state process $X = X^\alpha$ of dynamics

$$dX_t = \beta(t, X_t, \mathbb{P}_{X_t}^0, \alpha_t)dt + \sigma(t, X_t, \mathbb{P}_{X_t}^0, \alpha_t)dW_t + \sigma_0(t, X_t, \mathbb{P}_{X_t}^0)dW_t^0, \quad (2.3)$$

and cost functional

$$J_{MKV}(\alpha) = \mathbb{E} \left[\int_0^T e^{-rt} f(X_t, \mathbb{P}_{X_t}^0, \alpha_t) dt + e^{-rT} g(X_T, \mathbb{P}_{X_T}^0) \right].$$

(Here $\mathbb{P}_{X_t}^0$ denotes the conditional law of X_t given the common noise W^0). Moreover, the corresponding Bellman equation in the Wasserstein space of square-integrable probability measures $\mathcal{P}_2(\mathbb{R}^d)$ is given by (see [PW17])

$$\begin{cases} \partial_t v + \mathcal{F}(t, \mu, v, \partial_\mu v, \partial_x \partial_\mu v, \partial_\mu^2 v) = 0, & (t, \mu) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \\ v(T, \mu) = \mathcal{G}(\mu), & \mu \in \mathcal{P}_2(\mathbb{R}^d), \end{cases} \quad (2.4)$$

where $\partial_\mu \varphi(\mu)(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d$, $\partial_x \partial_\mu \varphi(\mu)(\cdot) : \mathbb{R}^d \rightarrow \mathbb{S}^d$, $\partial_\mu^2 \varphi(\mu)(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{S}^d$, are the L -derivatives of a function φ on $\mathcal{P}_2(\mathbb{R}^d)$ (see [CD18a]) and

$$\begin{aligned} \mathcal{F}(t, \mu, y, Z(\cdot), \Gamma(\cdot), \Gamma_0(\cdot, \cdot)) &= -ry + \int_{\mathbb{R}^d} h(t, x, \mu, Z(x), \Gamma(x)) \mu(dx) \\ &\quad + \int_{\mathbb{R}^d \times \mathbb{R}^d} \frac{1}{2} \text{tr}(\sigma_0(t, x, \mu) \sigma_0^\top(t, x', \mu) \Gamma_0(x, x')) \mu(dx) \mu(dx'), \\ \mathcal{G}(\mu) &= \int_{\mathbb{R}^d} g(x, \mu) \mu(dx), \end{aligned}$$

with

$$h(t, x, \mu, z, \gamma) = \inf_{a \in A} \left[\beta(t, x, \mu, a) \cdot z + \frac{1}{2} \text{tr}(\Sigma(t, x, \mu, a) \gamma) + f(x, \mu, a) \right].$$

□

We end this section by showing some exchangeability properties of the solution to the symmetric PDE (2.1). Let us introduce the notion of D -exchangeability where D stands for derivative.

Definition 2.3. *A function $(\mathbf{x}, x) \in (\mathbb{R}^d)^N \times \mathbb{R}^d \mapsto z(\mathbf{x}, x) \in \mathcal{X}$ is D -exchangeable if for any fixed $x \in \mathbb{R}^d$, $z(\cdot, x)$ is exchangeable. Given a D -exchangeable function z , we denote by \mathbf{z} the function defined on $(\mathbb{R}^d)^N$ by $\mathbf{z}(\mathbf{x}) = (z(\mathbf{x}, x_i))_{i \in [1, N]} \in \mathcal{X}^N$.*

This definition is actually motivated by the exchangeability property of the solution to the PDE (2.1), and by a structural property on the gradient of an exchangeable function that is differentiable.

Lemma 2.4. *The solution v to the PDE (2.1) with F and G satisfying **(HI)** is exchangeable, i.e., for all $\pi \in \mathfrak{S}_N$,*

$$v(t, \mathbf{x}) = v(t, \pi[\mathbf{x}]), \quad (t, \mathbf{x}) \in [0, T] \times (\mathbb{R}^d)^N.$$

Proof. Let $\pi \in \mathfrak{S}_N$. We set $v_\pi(t, \mathbf{x}) = v(t, \pi[\mathbf{x}])$, and observe that $\partial_t v(t, \pi[\mathbf{x}]) = \partial_t v_\pi(t, \mathbf{x})$, while

$$D_{\mathbf{x}} v(t, \pi[\mathbf{x}]) = \pi[D_{\mathbf{x}} v_\pi(t, \mathbf{x})], \quad D_{\mathbf{x}}^2 v(t, \pi[\mathbf{x}]) = \pi[D_{\mathbf{x}}^2 v_\pi(t, \mathbf{x})].$$

By writing the PDE (2.1) at $(t, \pi[\mathbf{x}])$, it follows under **(HI)** that v_π satisfies

$$\begin{cases} \partial_t v_\pi + F(t, \mathbf{x}, v_\pi, D_{\mathbf{x}} v_\pi, D_{\mathbf{x}}^2 v_\pi) = 0, & (t, \mathbf{x}) \in [0, T] \times (\mathbb{R}^d)^N \\ v_\pi(T, \mathbf{x}) = G(\mathbf{x}), & \mathbf{x} \in (\mathbb{R}^d)^N. \end{cases}$$

By uniqueness of the solution to PDE (2.1), we conclude that $v_\pi = v$, i.e., the exchangeability property of v . □

Lemma 2.5. *Let w be an exchangeable, and differentiable function on $(\mathbb{R}^d)^N$. Then there exists a D -exchangeable function z such that*

$$D_{x_i} w(\mathbf{x}) = z(\mathbf{x}, x_i), \quad i = 1, \dots, N, \tag{2.5}$$

for all $\mathbf{x} = (x_i)_{i \in [1, N]} \in (\mathbb{R}^d)^N$, i.e., $Dw = \mathbf{z}$.

Proof. Since w is exchangeable, it is clear that for fixed $i \in \llbracket 1, N \rrbracket$, and $x_i \in \mathbb{R}^d$,

$$\mathbf{x}_{-i} := (x_j)_{j \neq i} \in (\mathbb{R}^d)^{N-1} \mapsto D_{x_i} w(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_N) \quad \text{is exchangeable,}$$

and we shall then write:

$$z^i(\mathbf{x}_{-i}, x) := D_{x_i} w(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_N), \quad x \in \mathbb{R}^d.$$

By exchangeability of w , we also note that

$$z^i(\mathbf{x}_{-i}, x) = z^\ell(\mathbf{x}_{-i}, x), \quad \forall i, \ell \in \llbracket 1, N \rrbracket. \quad (2.6)$$

Let us now define the function z on $(\mathbb{R}^d)^N \times \mathbb{R}^d$ by:

$$z(\mathbf{x}, x) := \frac{1}{N} \sum_{p=0}^{N-1} (-1)^p \sum_{1 \leq i_1 < \dots < i_{p+1} \leq N} \sum_{\ell=1}^N z^\ell(\underbrace{(x, \dots, x)}_{p \text{ times}}, (x_j)_{j \neq i_1, \dots, i_{p+1}}, x).$$

for $\mathbf{x} = (x_1, \dots, x_N) \in (\mathbb{R}^d)^N$, and $x \in \mathbb{R}^d$, $(\underbrace{x, \dots, x}_{p \text{ times}}, (x_j)_{j \neq i_1, \dots, i_{p+1}})$ is the vector in $(\mathbb{R}^d)^{N-1}$ consisting of p components x , and the $N - p - 1$ components x_j , for $j \neq i_1, \dots, i_{p+1}$. By construction, it is clear that for fixed $x \in \mathbb{R}^d$, $z(\cdot, x)$ is exchangeable, i.e., z is a D -exchangeable function. Let us now show (2.5), i.e., that for fixed $\mathbf{x} = (x_i)_{i \in \llbracket 1, N \rrbracket} \in (\mathbb{R}^d)^N$,

$$z(\mathbf{x}, x_i) = z^i(\mathbf{x}_{-i}, x_i), \quad i \in \llbracket 1, N \rrbracket.$$

It suffices to check this property for $i = 1$. We set for $p = 0, \dots, N - 1$:

$$S^p := \sum_{1 \leq i_1 < \dots < i_{p+1} \leq N} \sum_{\ell=1}^N z^\ell(x_1, \underbrace{(x_1, \dots, x_1)}_{p \text{ times}}, (x_j)_{j \neq i_1, \dots, i_{p+1}}).$$

and see that

$$\left\{ \begin{array}{l} S^0 = \sum_{\ell=1}^N z^\ell(\mathbf{x}_{-1}, x_1) + \sum_{i_1=2}^N \sum_{\ell=1}^N z^\ell((x_1, x_j)_{j \neq 1, i_1}, x_1) \\ S^1 = \sum_{i_2=2}^N \sum_{\ell=1}^N z^\ell((x_1, x_j)_{j \neq 1, i_2}, x_1) + \sum_{2 \leq i_1 < i_2 \leq N} \sum_{\ell=1}^N z^\ell((x_1, x_j)_{j \neq i_1, i_2}, x_1) \\ \vdots \\ S^{N-2} = \sum_{2 \leq i_2 < \dots < i_{N-1} \leq N} \sum_{\ell=1}^N z^\ell(\underbrace{(x_1, \dots, x_1)}_{N-2 \text{ times}}, (x_j)_{j \neq 1, i_2, \dots, i_{N-1}}, x_1) + \sum_{\ell=1}^N z^\ell((x_1, \dots, x_1), x_1) \\ S^{N-1} = \sum_{\ell=1}^N z^\ell((x_1, \dots, x_1), x_1). \end{array} \right.$$

The telescopic sum then yields

$$z((x_1, \dots, x_N), x_1) = \frac{1}{N} \sum_{p=0}^{N-1} (-1)^p S^p = \frac{1}{N} \sum_{\ell=1}^N z^\ell((x_j)_{j \neq 1}, x_1) = z^1((x_j)_{j \neq 1}, x_1),$$

where the last equality follows from (2.6). This shows the property (2.5). \square

3 Symmetric neural networks

3.1 DeepSets and variants

In view of Lemma 2.4 and 2.5, we shall consider a class of neural networks (NN in short) that satisfy the exchangeability and D -exchangeability properties for approximating the solution (and its gradient) to the PDE (2.1).

We denote by

$$\mathcal{L}_{d_1, d_2}^\rho = \left\{ \phi : \mathbb{R}^{d_1} \rightarrow \mathbb{R}^{d_2} : \exists (\mathcal{W}, \beta) \in \mathbb{R}^{d_2 \times d_1} \times \mathbb{R}^{d_2}, \phi(x) = \rho(\mathcal{W}x + \beta) \right\},$$

the set of layer functions with input dimension d_1 , output dimension d_2 , and activation function $\rho : \mathbb{R} \rightarrow \mathbb{R}$. Here, the activation is applied component-wise, i.e., $\rho(x_1, \dots, x_{d_2}) = (\rho(x_1), \dots, \rho(x_{d_2}))$, to the result of the affine map $x \in \mathbb{R}^{d_1} \mapsto \mathcal{W}x + \beta \in \mathbb{R}^{d_2}$, with a matrix \mathcal{W} called weight, and vector β called bias. Standard examples of activation functions are the sigmoid, the ReLU, the Elu (see [CUH16]), or tanh. When ρ is the identity function, we simply write \mathcal{L}_{d_1, d_2} .

We then define

$$\mathcal{N}_{d_0, \ell, m, k}^\rho = \left\{ \varphi : \mathbb{R}^{d_0} \rightarrow \mathbb{R}^k : \exists \phi_0 \in \mathcal{L}_{d_0, m}^\rho, \exists \phi_i \in \mathcal{L}_{m, m}^\rho, i = 1, \dots, \ell - 1, \exists \phi_\ell \in \mathcal{L}_{m, k}, \right. \\ \left. \varphi = \phi_\ell \circ \phi_{\ell-1} \circ \dots \circ \phi_0 \right\},$$

as the set of feedforward (or artificial) neural networks with input layer dimension d_0 , output layer dimension k , and ℓ hidden layers with m neurons (or units). These numbers d_0, ℓ, m , and the activation function ρ , form the architecture of the network. When $\ell = 1$, one usually refers to shallow neural networks, as opposed to deep neural networks which have several hidden layers.

A *symmetric neural network* function, denoted $\mathcal{U} \in \mathcal{S}_{d, \ell, m, k, d'}^{\mathfrak{s}, N, \rho}$, is an $\mathbb{R}^{d'}$ -valued exchangeable function to the order N on \mathbb{R}^d , in the form:

$$\mathcal{U}(\mathbf{x}) = \psi(\mathfrak{s}((\varphi(x_i))_{i \in [1, N]})), \quad \text{for } \mathbf{x} = (x_i)_{i \in [1, N]} \in (\mathbb{R}^d)^N, \quad (3.1)$$

where $\varphi \in \mathcal{N}_{d, \ell, m, k}^\rho$, $\psi \in \mathcal{N}_{k, \ell, m, d'}^\rho$ (here, for simplicity of notations, we assume that the number of hidden layers and neurons of φ and ψ are the same but in practical implementation, they may be different), and \mathfrak{s} is a given \mathbb{R}^k -valued exchangeable function to the order N on \mathbb{R}^k , typically:

- Max-pooling (component-wise): $\mathfrak{s}(\mathbf{y}) = \max(y_i)_{i \in [1, N]}$,
- Sum: $\mathfrak{s}(\mathbf{y}) = \sum_{i=1}^N y_i$, or average: $\mathfrak{s}(\mathbf{y}) = \frac{1}{N} \sum_{i=1}^N y_i$,

for $\mathbf{y} = (y_i)_{i \in [1, N]} \in (\mathbb{R}^k)^N$. When \mathfrak{s} is the max-pooling function, $\mathcal{S}_{d, \ell, m, k, d'}^{\mathfrak{s}, N, \rho}$ is called *PointNet*, as introduced in [RQ+17], while for \mathfrak{s} equals to the sum/average function, it is called *DeepSet*, see [Zah+17]. The architecture is described in Figure 1, and k can be interpreted as a number of features describing the geometry of the set of points $\{x_i\}_{i \in [1, N]}$. For example in the context of mean-field control problem, k will be related to the moments for describing the law of the McKean-Vlasov SDE.

A given symmetric network function $\mathcal{U} \in \mathcal{S}_{d, \ell, m, k, d'}^{\mathfrak{s}, N, \rho}$ is determined by the weight/bias parameters $\theta = (\theta^{(1)}, \theta^{(2)})$ with $\theta^{(1)} = (\mathcal{W}_0^{(1)}, \beta_0^{(1)}, \dots, \mathcal{W}_\ell^{(1)}, \beta_\ell^{(1)})$ defining the layer functions

$\phi_0 \dots, \phi_\ell$ of φ , and $\theta^{(2)} = (\mathcal{W}_0^{(2)}, \beta_0^{(2)}, \dots, \mathcal{W}_\ell^{(2)}, \beta_\ell^{(2)})$ defining the layer functions $\psi_0 \dots, \psi_\ell$ of ψ . The number of parameters is $M = M_1 + M_2$, with $M_1 = m(d+1) + m(m+1)(\ell-1) + (m+1)k$, $M_2 = (k+1)m + m(m+1)(\ell-1) + (m+1)d'$, and we observe that it does not depend on the number N of inputs.

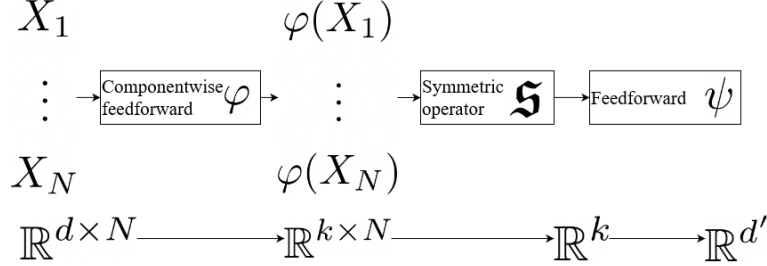


Figure 1: Architecture of a symmetric neural network.

Remark 3.1 (*Time dependent symmetric network*). A time-dependent symmetric in space neural network can be constructed as

$$\mathcal{U}(t, \mathbf{x}) = \psi(t, \mathfrak{s}((\varphi(x_i))_{i \in [1, N]})), \quad \text{for } t \in \mathbb{R}_+, \mathbf{x} = (x_i)_{i \in I} \in (\mathbb{R}^d)^N,$$

with φ a feedforward network from \mathbb{R}^d into \mathbb{R}^k , and ψ is a feedforward from \mathbb{R}^{k+1} into $\mathbb{R}^{d'}$, where we add time as an additional feature, see architecture in Figure 2.

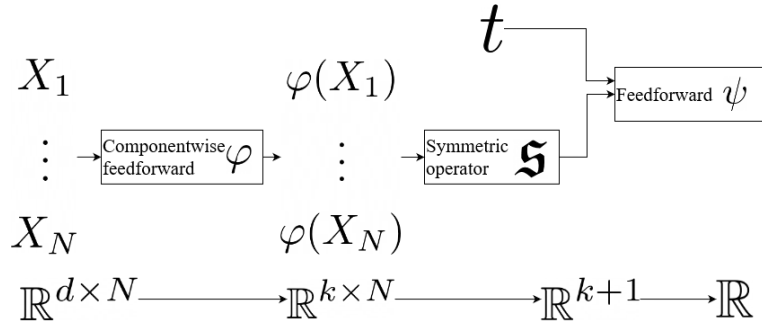


Figure 2: Architecture of time dependent symmetric network.

□

A D -symmetric neural network function, denoted $\mathcal{Z} \in DS_{d, \ell, m, k, d'}^{\mathfrak{s}, N, \rho}$, is an $\mathbb{R}^{d'}$ -valued D -exchangeable function in the form

$$\mathcal{Z}(\mathbf{x}, x) = \psi(\mathfrak{s}((\varphi(x_i))_{i \in [1, N]}), x), \quad \text{for } \mathbf{x} = (x_i)_{i \in [1, N]} \in (\mathbb{R}^d)^N, x \in \mathbb{R}^d,$$

where $\varphi \in \mathcal{N}_{d, \ell, m, k}^\rho$, $\psi \in \mathcal{N}_{k+d, \ell, m, d'}^\rho$, and \mathfrak{s} is a given \mathbb{R}^k -valued N -exchangeable function on \mathbb{R}^k . The number of parameters of a given $\mathcal{Z} \in DS_{d, \ell, m, k, d'}^{\mathfrak{s}, N, \rho}$ is $M' = M'_1 + M'_2$, with $M'_1 = m(d+1) + m(m+1)(\ell-1) + (m+1)k$, $M'_2 = m(k+d+1) + m(m+1)(\ell-1) + (m+1)d'$. We shall often take for \mathfrak{s} the average function, and call $DS_{d, \ell, m, k, d'}^{\mathfrak{s}, N, \rho}$ as *DeepDerSet*. Its architecture is given in Figure 3. Given a D -symmetric neural network \mathcal{Z} , we denote by \mathcal{Z} the function defined on $(\mathbb{R}^d)^N$ by $\mathcal{Z}(\mathbf{x}) = (\mathcal{Z}(\mathbf{x}, x_i))_{i \in [1, N]} \in (\mathbb{R}^d)^N$, and by misuse

of notation, we may also call \mathcal{Z} as a D -symmetric NN. By construction, these networks respect the representation given by Lemma 2.5, by being defined as a D -exchangeable function applied component by component. In that way we are able to enforce the correct symmetries for representing both a symmetric function and its derivative, which will be useful in Section 4 and Section 5 when looking for the gradient of PDE solutions.

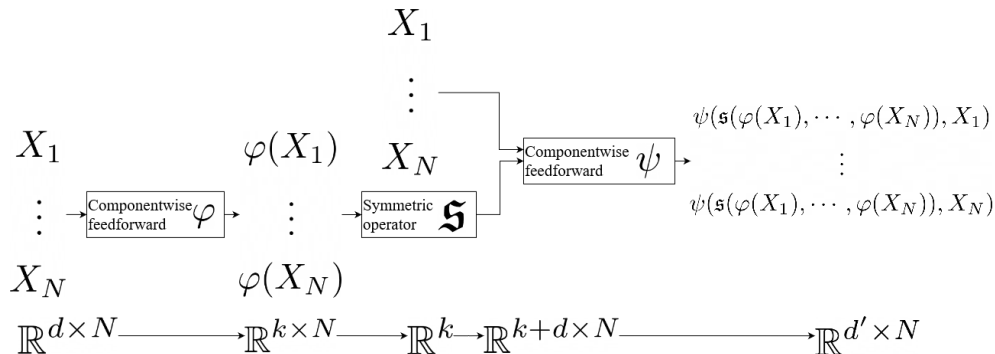


Figure 3: Architecture of DeepDerSet network

Alternatively, one can generate D -exchangeable functions as follows. Starting from a DeepSet element $\mathcal{U} \in \mathcal{S}_{d,\ell,m,k,d'}^{s,N,\rho}$ as in (3.1) with \mathfrak{s} the sum function, and network functions φ, ψ with smooth activation functions

$$D_{x_i} \mathcal{U}(\mathbf{x}) = D\mathcal{U}(\mathbf{x}, x_i), \quad \mathbf{x} = (x_i)_{i \in [1,N]} \in (\mathbb{R}^d)^N,$$

with

$$D\mathcal{U}(\mathbf{x}, x) := D\varphi(x) D\psi(\mathfrak{s}((\varphi(x_i))_{i \in I})), \quad x \in \mathbb{R}^d.$$

The set of D -exchangeable functions obtained from differentiation of DeepSet network functions is called AD -DeepSet, where AD stands for automatic differentiation, and Automatic refers to the implementation of the differentiation in software library, e.g. in TensorFlow. This alternative network for approximating the derivative of a differentiable symmetric function naturally respects the structure given by Lemma 2.5, by construction as the derivative of a symmetric DeepSet neural network.

Given a D -symmetric NN \mathcal{Z} in *DeepDerSet* or in *AD-DeepSet* with smooth activation functions, we denote by $D\mathcal{Z}$ its differentiation.

As for the well-known universal approximation theorem [Hor91] for neural networks, we have a similar result for symmetric neural networks, which states that any exchangeable function can be arbitrarily approximated by a PointNet or DeepSet given enough neurons. More precisely, by combining Theorem 2.9 of [Wag+19] with Theorem 2 of [Hor91], we obtain the following approximation theorem for DeepSets.

Universal approximation for DeepSets networks. Let \mathfrak{s} be the sum function. The set $\cup_{m=1}^{\infty} \mathcal{S}_{d,\ell,m,N+1,d'}^{s,N,\rho}$ approximates any N -exchangeable continuous function w arbitrary well on any compact set of $K \subset \mathbb{R}^d$, once ρ is continuous, bounded and non-constant: for all $\varepsilon > 0$, $N \in \mathbb{N}^*$, there exists $\mathcal{U} \in \cup_{m=1}^{\infty} \mathcal{S}_{d,\ell,m,N+1,d'}^{s,N,\rho}$ such that

$$|w(\mathbf{x}) - \mathcal{U}(\mathbf{x})| \leq \varepsilon \quad \forall \mathbf{x} \in K^N.$$

Note that a priori the latent space dimension k has to be chosen equal to $N + 1$.

Alternatively, by combining Theorem 1 of [RQ+17] with Theorem 2 of [Hor91], we obtain the following one-dimensional approximation theorem for PointNet.

Universal approximation for PointNet networks. Let \mathfrak{s} be the max function. The set $\cup_{m=1}^{\infty} \cup_{k=1}^{\infty} \mathcal{S}_{1,\ell,m,k,d'}^{\mathfrak{s},N,\rho}$ approximates any N -exchangeable Hausdorff continuous function w (seen as a function on sets) arbitrary well on any compact set of $K \subset \mathbb{R}$, once ρ is continuous, bounded, and non-constant: for all $\varepsilon > 0$, $N \in \mathbb{N}^*$, there exists $\mathcal{U} \in \cup_{m=1}^{\infty} \cup_{k=1}^{\infty} \mathcal{S}_{1,\ell,m,k,d'}^{\mathfrak{s},N,\rho}$ such that

$$|w(S) - \mathcal{U}(\mathbf{x})| \leq \varepsilon, \quad \forall S \subset K, S = \{x_1, \dots, x_N\}.$$

Note here that a priori the latent space dimension k has to be chosen as large as needed.

3.2 Comparison tests

In this paragraph, we test the accuracy of the approximation of exchangeable functions by *DeepSet* or *PointNet*, and also the approximation of D -exchangeable functions by *DeepDerSet* or *AD-DeepSet*, and compare numerically with classical feedforward approximations.

3.2.1 Approximation of some simple functions

We first test the approximation of the following simple symmetric functions:

1. $f(x) = \exp(2\bar{x} + 3\bar{x}^3)$, with $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$ (case 1)
2. $f(x) = \frac{1}{N} \sum_{i=1}^N [\sin(x_i)1_{x_i < 0} + x_i 1_{x_i \geq 0}]$ (case 2)
3. $f(x) = \bar{x} + 2\bar{x}^2 + 3\bar{x}^3$, with $\bar{x} = \max\{x_i, i = 1, \dots, N\}$ (case 3)
4. $f(x) = \cos(2\bar{x} + 3\bar{x}^2)$, with $\bar{x} = \sum_{i=1}^N x_i$ (case 4)

We use a symmetric neural network architecture as proposed in [RQ+17], [Zah+17]:

- First, a feedforward network φ with $\ell = 5$ hidden layers, and respectively 64, 64, 64, 128 and 1024 neurons such that each dimension i , $i = 1, \dots, N$, is treated with the same network in one dimension avoiding to break the symmetry.
- Two possible symmetric functions \mathfrak{s} to the order N on \mathbb{R}^k with $k = 1024$, the max-pooling (*PointNet*) and the sum function (*DeepSet*).
- At last, a feedforward network ψ from \mathbb{R}^{1024} to \mathbb{R} with $\ell = 2$ hidden layers, and respectively 512 and 256 neurons.

For the approximation with classical feedforward networks, we used three or four layers and a number of neurons constant per layer equal to $10 + d$, or $10 + 2d$ neurons.

We use the ADAM optimizer ([KB14]), with a batch size equal to 300 for solving the approximation problem with quadratic loss function:

$$\min_{\theta} \mathbb{E}[|f(X) - \mathcal{U}^{\theta}(X)|^2], \quad (3.2)$$

with training simulations from $X \sim \mathcal{N}(0_N, 1_N)$, and θ are the parameters of the network function \mathcal{U}^θ . The number of epochs (corresponding to the number of gradient descent iterations) used is equal to 100. After epoch iterations of the stochastic gradient, the error (3.2) is estimated with 20000 simulations. If the error is below a threshold equal to $1e-5$ the optimization is stopped, otherwise a counter for outer iterations is incremented. The number of outer iterations is blocked at $\text{epochExt} = 1000$ (meaning a maximal total number of stochastic gradient iterations equal to $\text{epoch} \times \text{epochExt} = 100000$).

In Tables 1, 2, 3, we report the accuracy reached (Error) and the number of iterations (Iter.) used to obtain this given accuracy: then a threshold equal to $1e-5$ means that the optimization has been successful and the relevant parameter is the number of iterations used. A number of iterations equal to $\text{epochExt} = 1000$ means that the optimization has not been successful and the error reached indicates how far we are from optimality. For the feedforward case, we report the best result (“minimum” in table) and the worse result (“maximum” in table) obtained changing the number of layers and the number of neurons used.

The initial learning rate is taken equal to $1e-3$ for first outer simulation in cases 1 and 2 with a linear decay to $1e-5$ for a number of outer iterations equal to 1000. For test case 3, the initial learning rate is taken equal to $1e-4$ with a linear decay to $1e-5$. The result obtained in Table 1 is similarly obtained with a large number of functions tested in dimension between $N = 10$ to 1000. It shows the following results:

- The classical feedforward, with dense layers, often permits to obtain optimally without forcing symmetry of the solution,
- Classical feedforward results do not depend a lot on the number of layers, the number of neurons tested and the activation function used,
- For symmetric approximations, DeepSets generally permits to get the best results and the ReLU activation function is the best out of the three tested.

	Symmetric				Feedforward			
	PointNet		DeepSet		Minimum		Maximum	
Activation	Error	Iter.	Error	Iter.	Error	Iter.	Error	Iter.
ReLU	0.008	1000	$1e-5$	10	$1e-5$	125	$1e-5$	166
tanh	0.016	1000	$1e-5$	288	$1e-5$	180	$1e-5$	308
ELU	0.015	1000	$1e-5$	176	$1e-5$	108	$1e-5$	130

Table 1: Approximation error (3.2) obtained for different networks on one run and number of iterations used depending on activation functions for approximation of the function f in case 1, dimension $N = 100$.

In the sequel, we drop the ELU activation function on other cases as shown in Tables 2 for cases 2 and 3. Notice that case 3, involving a max function is the only one where PointNet approximation gives the best result among the other tested networks. On cases 2 and 3 in dimension 100, the DeepSets approximation outperforms the classical feedforward network for all the number of layers and neurons tested. However, results on case 3 are not very good for the PointNet approximation even with the ReLU activation function.

Case	activ	Symmetric				Feedforward			
		PointNet		DeepSets		Minimum		Maximum	
		Error	Iter.	Error	Iter.	Error	Iter.	Error	Iter.
2	ReLU	0.001	1000	1e−5	5	1e−5	992	0.002	1000
2	tanh	0.03	1000	1e−5	342	0.0018	1000	6× 1e−5	1000
3	ReLU	0.001	1000	0.23	1000	88	1000	432	1000
3	tanh	0.002	1000	65	1000	933	1000	969	1000

Table 2: Approximation error (3.2) obtained for different networks on one run and number of iterations used for approximation of the function f in cases 2 and 3, dimension $N = 100$, activation function ReLU.

Results for test case 4 are given on Table 3 using an initial learning rate equal to $5e-5$ and a decay linear to $5e-6$ with the number of outer iterations. At last, considering case

	Symmetric				Feedforward			
	PointNet		DeepSets		Minimum		Maximum	
activ	Error	Iter.	Error	Iter.	Error	Iter.	Error	Iter.
ReLU	0.1910	1000	8.9e−5	1000	0.0045	1000	0.01	1000
tanh	0.19	1000	4e−5	1000	0.0006	1000	0.0012	1000

Table 3: Approximation error (3.2) obtained for different networks, activation functions for approximation of the function f in case 4 dimension 1000.

4 in dimension $N = 1000$, when the function is quickly changing, we see that the classical feedforward network functions have difficulty to converge while the DeepSets network approximation converges. The latter turns out to be a very good candidate to some very high dimensional PDEs when there is symmetry in the solution.

3.2.2 Approximation of a function of t and x with symmetry in x

We test the accuracy of our time dependent symmetric neural network by considering the following two cases of functions:

1. $f(x) = \exp(\bar{x}(t + 2t^2) + 3t\bar{x}^3)$ with $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$ (case 1)
2. $f(x) = t + \cos(t\bar{x})$, with $\bar{x} = \frac{1}{\sqrt{N}} \sum_{i=1}^N x_i$ (case 2)

The approximation is performed through the minimization problem

$$\min_{\theta} \mathbb{E}[|f(\tau, X) - \mathcal{U}^{\theta}(\tau, X)|^2], \quad (3.3)$$

with training simulations from $X \sim \mathcal{N}(0_N, 1_N)$, and an independent uniform law for τ on $[0, 1]$, and where \mathcal{U}^{θ} is a time-dependent DeepSet with parameters θ . We keep the same number of neurons and layers as in the previous section, and compare with a classical feedforward network composed of 3 layers of $d + 10$ neurons. In all experiments, we use a ReLU activation function.

Case	DeepSets		Feedforward	
	Error	Iter.	Error	Iter.
1	1e-5	67	0.008	1000
2	1e-5	344	0.048	1000

Table 4: Approximation error (3.3) obtained for different networks on one run and number of iterations used for approximation of the function f in case 1 and 2, dimension $N = 100$.

In Table 4, we give the results obtained in dimension 100. Surprisingly, the feedforward approximation seems to have difficulties to approximate the case 1 although it is quite similar to case one in the previous section. As for the second case, the result is not so surprising as the case is quite similar to case 4 in previous section, where the feedforward network has difficulties to converge.

3.2.3 Gradient approximation

We now focus on the approximation of the derivative of an exchangeable function by means of \mathcal{U}^θ a *DeepDerSet*, a *AD-DeepSet*, or a classical feedforward network .

The minimization problem is now:

$$\min_{\theta} \mathbb{E}[\|Df(X) - \mathcal{U}^\theta(X)\|^2], \quad (3.4)$$

where the norm $\|\cdot\|$ is the Euclidean norm.

The comparison is performed on the following test functions:

1. $f(x) = \exp(\bar{x} + \bar{x}^3)(1_N + 3x^2)$ where $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$ (case 1)
2. $f(x) = y$ where $y_i = 1_{x_i > 0} + \cos(x_i)1_{x_i < 0}$, $i = 1, \dots, N$ (case 2)
3. $f(x) = \frac{1}{\sqrt{N}} \sin(\bar{x})1_N$, where $\bar{x} = \frac{1}{\sqrt{N}} \sum_{i=1}^N x_i$ (case 3)

We compare the classical feedforward approximation to our network approximation in Tables 5 and 6, using a maximal number of iterations equal to 5000. Clearly using a ReLU activation function is superior to the tanh activation function and the *DeepDerSet* gives the best approximation while the *AD-DeepSet* or the feedforward may have difficulties to approximate the functions accurately.

4 Numerical schemes

We now adapt the deep backward dynamic programming (DBDP) schemes developed in [HPW20] and [PWG21] for solving nonlinear PDEs, by using symmetric neural networks and D -symmetric neural networks (instead of feedforward neural networks) for approximating the exchangeable solution v and its gradient $D_{\mathbf{x}}v$. We recall the main steps of the DBDP scheme, and distinguish the case of semi-linear and fully non-linear PDEs.

Case	N	Feedforward		DeepDerSet		AD-DeepSet	
		Error	Iter.	Error	Iter.	Error	Iter.
1	10	5e-4	5000	5e-4	5000	7e-3	5000
1	100	1e-5	300	1e-5	250	1e-5	50
2	10	0.03	5000	1e-5	550	2e-4	5000
2	100	0.12	5000	1e-5	450	1e-4	5000
3	10	1e-5	3800	1e-5	850	0.03	5000
3	100	1e-5	1850	1e-5	700	3e-3	5000

Table 5: Approximation error (3.4) with ReLU activation function obtained for different networks on one run and number of iterations used for approximation of the derivative of an exchangeable function.

Case	N	Feedforward		DeepDerSet		AD-DeepSet	
		Error	Iter.	Error	Iter.	Error	Iter.
1	10	5e-4	5000	1.7e-4	5000	3e-3	5000
1	100	1e-5	500	1e-5	100	1e-5	50
2	10	0.03	5000	1e-5	950	0.028	5000
2	100	0.12	5000	1e-5	700	0.56	5000
3	10	1.5e-5	5000	1e-5	1850	0.02	5000
3	100	1e-5	2100	1e-5	1350	4e-3	5000

Table 6: Approximation error (3.4) with tanh activation function obtained for different networks on one run and number of iterations used for approximation of the derivative of an exchangeable function.

4.1 Semi-linear PDE

We first consider the case where the generator F in (2.1) may be decomposed into the form

$$F(t, \mathbf{x}, y, \mathbf{z}, \boldsymbol{\gamma}) = H(t, \mathbf{x}, y, \mathbf{z}) + \sum_{i=1}^N b_i(t, \mathbf{x}) \cdot z_i + \frac{1}{2} \sum_{i,j=1}^N \text{tr}(\Sigma_{ij}(t, \mathbf{x}) \gamma_{ij}),$$

for $t \in [0, T]$, $\mathbf{x} = (x_i)_{i \in [1, N]} \in (\mathbb{R}^d)^N$, $y \in \mathbb{R}$, $\mathbf{z} = (z_i)_{i \in [1, N]} \in (\mathbb{R}^d)^N$, and $\boldsymbol{\gamma} = (\gamma_{ij})_{i, j \in [1, N]} \in \mathbb{S}^N(\mathbb{S}^d)$. Here, H is a function on $[0, T] \times (\mathbb{R}^d)^N \times \mathbb{R} \times (\mathbb{R}^d)^N$ satisfying the permutation-invariance condition:

$$H(t, \mathbf{x}, y, \mathbf{z}) = H(t, \pi[\mathbf{x}], y, \pi[\mathbf{z}]), \quad \forall \pi \in \mathfrak{S}_N,$$

the coefficients b_i , $i = 1, \dots, N$, are \mathbb{R}^d -valued functions on $[0, T] \times (\mathbb{R}^d)^N$ satisfying the condition

$$b_i(t, \pi[\mathbf{x}]) = b_{\pi(i)}(t, \mathbf{x}), \quad \forall \pi \in \mathfrak{S}_N, \quad (4.2)$$

and the coefficients Σ_{ij} , $i, j = 1, \dots, N$, are $d \times d$ -matrix valued functions on $[0, T] \times (\mathbb{R}^d)^N$ in the form

$$\Sigma_{ij}(t, \mathbf{x}) = \sigma_{ij} \sigma_{ij}^\top(t, \mathbf{x}) + \sigma_{i0} \sigma_{j0}^\top(t, \mathbf{x}), \quad (4.3)$$

for some $d \times d$ -matrix valued functions σ_{ij} , and $d \times q$ -matrix valued functions σ_{i0} , satisfying the invariance property: for all $\pi \in \mathfrak{S}_N$,

$$\sigma_{ij}(t, \pi[\mathbf{x}]) = \sigma_{\pi(i)\pi(j)}(t, \mathbf{x}), \quad \sigma_{i0}(t, \pi[\mathbf{x}]) = \sigma_{\pi(i)0}(t, \mathbf{x}). \quad (4.4)$$

In this case, the permutation-invariance condition **(HI)** on F is satisfied, and we observe that it includes Example 2.1 of multi-asset pricing with $H(t, \mathbf{x}, y, \mathbf{z}) = \beta(y^+ - y)$ (in the case of the CVA pricing), $b_i \equiv r$ and $\sigma_{i0} \equiv 0$. This also includes Example 2.2 of the McKean-Vlasov control problem under common noise with uncontrolled diffusion coefficient, where $\sigma_{ij}(t, \mathbf{x}) = \sigma(t, x_i, \bar{\mu}(\mathbf{x}))\delta_{ij}$, $\sigma_{i0}(t, \mathbf{x}) = \sigma_0(t, x_i, \bar{\mu}(\mathbf{x}))$, and

$$\begin{aligned} H(t, \mathbf{x}, y, \mathbf{z}) = & -ry + \sum_{i=1}^N \inf_{a \in A} \left[\beta(t, x_i, \bar{\mu}(\mathbf{x}), a) \cdot z_i + \frac{1}{N} f(x_i, \bar{\mu}(\mathbf{x}), a) \right] \\ & - \sum_{i=1}^N b_i(t, \mathbf{x}) \cdot z_i, \end{aligned} \quad (4.5)$$

for any function b_i satisfying (4.2). We shall discuss more in detail the relevant choice of the drift coefficient b_i in Section 4.3.

The starting point of the numerical scheme is the probabilistic representation of the PDE (2.1) with F as in (4.1) in terms of a forward backward stochastic differential equation (FBSDE), as in [PP90]. In our context, the forward system is described by the process $\mathbf{X} = (X^1, \dots, X^N)$ valued in $(\mathbb{R}^d)^N$ governed by the diffusion dynamics:

$$dX_t^i = b_i(t, \mathbf{X}_t)dt + \sum_{j=0}^N \sigma_{ij}(t, \mathbf{X}_t)dW_t^j, \quad (4.6)$$

where W^i , $i = 1, \dots, N$, are independent d -dimensional Brownian motions, independent of the q -dimensional Brownian motion W^0 . Given this forward diffusion process, we then consider the pair process $(Y, \mathbf{Z} = (Z^i)_{i \in [1, N]})$ valued in $\mathbb{R} \times (\mathbb{R}^d)^N$ solution to the BSDE

$$\begin{aligned} G(\mathbf{X}_T) - Y_t + \int_t^T H(s, \mathbf{X}_s, Y_s, \mathbf{Z}_s)ds \\ - \sum_{i=1}^N \sum_{j=0}^N \int_t^T (Z_s^i)^\top \sigma_{ij}(s, \mathbf{X}_s) dW_s^j = 0, \quad 0 \leq t \leq T, \end{aligned} \quad (4.7)$$

which is connected by Itô's formula to the solution of the PDE (2.1) via:

$$Y_t = v(t, \mathbf{X}_t), \quad Z_t^i = D_{x_i} v(t, \mathbf{X}_t), \quad i = 1, \dots, N, \quad 0 \leq t \leq T.$$

We next consider a time discretization of this FBSDE on a time grid $\{t_k, k = 0, \dots, N_T\}$, with $t_0 = 0$, $t_{N_T} = T$, $\Delta t_k := t_{k+1} - t_k > 0$, by defining the Euler scheme $\{\mathbf{X}_k^{N_T} = (X_k^{i, N_T})_{i \in [1, N]}, k = 0, \dots, N_T\}$ associated to the forward diffusion process $\{\mathbf{X}_t = (X_t^i)_{i \in [1, N]}, 0 \leq t \leq T\}$, which is used for the training simulations, together with the increments of the Brownian motions: $\Delta W_k^j := W_{t_{k+1}}^j - W_{t_k}^j$, $k = 0, \dots, N_T - 1$, $j = 0, \dots, N$, of our numerical backward scheme. The DBDP algorithm reads then as follows:

The output of the DBDP scheme provides approximations $\hat{U}_k(\mathbf{x})$ of $v(t_k, \mathbf{x})$, and $\hat{\mathbf{Z}}_k(\mathbf{x})$ of $D_{\mathbf{x}} v(t_k, \mathbf{x})$, $k = 0, \dots, N_T - 1$, for values of $\mathbf{x} \in (\mathbb{R}^d)^N$ that are well-explored by the training simulations of $\mathbf{X}_k^{N_T}$. We refer to [PWG21] (see their section 3.1) for a discussion on the choice of the algorithm hyperparameters.

Algorithm 1: DBDP scheme with symmetric NN

Initialization: Initialize from the exchangeable function: $\widehat{\mathcal{U}}_{N_T}(\cdot) = G(\cdot)$

for $k = N_T - 1, \dots, 0$ **do**

minimize over symmetric NN \mathcal{U}_k , and D -symmetric NN \mathcal{Z}_k , the quadratic loss function

$$J_1(\mathcal{U}_k, \mathcal{Z}_k) = \mathbb{E} \left| \widehat{\mathcal{U}}_{k+1}(\mathbf{X}_{k+1}^{N_T}) - \mathcal{U}_k(\mathbf{X}_k^{N_T}) + H(t_k, \mathbf{X}_k^{N_T}, \mathcal{U}_k(\mathbf{X}_k^{N_T}), \mathcal{Z}_k(\mathbf{X}_k^{N_T})) \Delta t_k - \sum_{i=1}^N \sum_{j=0}^N (\mathcal{Z}_k(\mathbf{X}_k^{N_T}, X_k^{i, N_T}))^\top \sigma_{ij}(t_k, \mathbf{X}_k^{N_T}) \Delta W_k^j \right|^2,$$

and update $(\widehat{\mathcal{U}}_k, \widehat{\mathcal{Z}}_k)$ as the solution to this minimization problem.

end

Remark 4.1. We stress that the neural networks do not take time as an input. Adding time would not make any difference because the training is done locally, time step per time step. So the neural networks approximating \mathcal{U}_k and \mathcal{Z}_k would be trained with only t_k as an input and hence they would not be able to learn the dependence on time. However, at time $k < N_T$, we initialize the parameters of \mathcal{U}_k and \mathcal{Z}_k respectively with the parameters of the neural networks for \mathcal{U}_{k+1} and \mathcal{Z}_{k+1} , which have been trained at the previous iteration. This gives a good initial guess for the neural networks at time t_k and leads to more efficient training.

4.2 Fully nonlinear PDE

We consider more generally the fully non-linear PDE (2.1) with a symmetric generator F satisfying **(HI)**. We adapt the machine learning scheme in [PWG21] for solving fully nonlinear PDEs by exploiting furthermore the exchangeability property of the solution by using again symmetric neural networks as in the semi-linear case.

We fix some arbitrary drift and diffusion coefficients $b_i, \sigma_{ij}, i = 1, \dots, N, j = 0, \dots, N$, satisfying invariance properties as in (4.2)-(4.4) (in practice, they should be chosen depending on the studied problem as for the semi-linear case, see a general discussion in Section 3.1 in [PWG21], and an application in Section 5.3), and introduce the forward diffusion system \mathbf{X} as in (4.6) and its discrete-time Euler scheme \mathbf{X}^{N_T} . We then consider the triple process $(Y, \mathbf{Z} = (Z^i)_{i \in [1, N]}, \mathbf{\Gamma} = (\Gamma^{ij})_{i, j \in [1, N]})$ valued in $\mathbb{R} \times (\mathbb{R}^d)^N \times \mathbb{S}^N(\mathbb{S}^d)$ solution to the BSDE

$$G(\mathbf{X}_T) - Y_t + \int_t^T F_{b, \sigma}(s, \mathbf{X}_s, Y_s, \mathbf{Z}_s, \mathbf{\Gamma}_s) ds - \sum_{i=1}^N \sum_{j=0}^N \int_t^T (Z_s^i)^\top \sigma_{ij}(s, \mathbf{X}_s) dW_s^j = 0, \quad 0 \leq t \leq T,$$

with

$$F_{b, \sigma}(t, \mathbf{x}, y, \mathbf{z}, \gamma) := F(t, \mathbf{x}, y, \mathbf{z}, \gamma) - \sum_{i=1}^N b_i(t, \mathbf{x}) \cdot z_i - \frac{1}{2} \sum_{i, j=1}^N \text{tr}(\Sigma_{ij}(t, \mathbf{x}) \gamma_{ij}),$$

and Σ_{ij} as in (4.3). It is connected by Itô's formula to the fully non-linear PDE (2.1) via the representation: $Y_t = v(t, \mathbf{X}_t)$, $\mathbf{Z}_t = D_{\mathbf{x}}v(t, \mathbf{X}_t)$, $\mathbf{\Gamma}_t = D_{\mathbf{x}}^2v(t, \mathbf{X}_t)$, $0 \leq t \leq T$.

Assuming that G is smooth, the algorithm is designed in Algorithm 2.

Algorithm 2: Fully nonlinear DPBD scheme with symmetric NN

Initialization: Initialize from the exchangeable function: $\widehat{\mathcal{U}}_{N_T}(\cdot) = G(\cdot)$ and the D -exchangeable function $\widehat{\mathcal{Z}}_{N_T}(\cdot) = DG(\cdot)$.

for $k = N_T - 1, \dots, 0$ **do**

 minimize over symmetric NN \mathcal{U}_k , and D -symmetric NN \mathcal{Z}_k , the quadratic loss function

$$\begin{aligned} J_2(\mathcal{U}_k, \mathcal{Z}_k) = & \mathbb{E} \left| \widehat{\mathcal{U}}_{k+1}(\mathbf{X}_{k+1}^{N_T}) - \mathcal{U}_k(\mathbf{X}_k^{N_T}) \right. \\ & + F_{b,\sigma}(t_k, \mathbf{X}_k^{N_T}, \mathcal{U}_k(\mathbf{X}_k^{N_T}), \mathcal{Z}_k(\mathbf{X}_k^{N_T}), D\widehat{\mathcal{Z}}_{k+1}(\mathbf{X}_{k+1}^{N_T}))\Delta t_k \\ & \left. - \sum_{i=1}^N \sum_{j=0}^N (\mathcal{Z}_k(\mathbf{X}_k^{N_T}, X_k^{i,N_T}))^\top \sigma_{ij}(t_k, \mathbf{X}_k^{N_T}) \Delta W_k^j \right|^2, \end{aligned}$$

 and update $(\widehat{\mathcal{U}}_k, \widehat{\mathcal{Z}}_k)$ as the solution to this minimization problem. Here $D\widehat{\mathcal{Z}}_{k+1}$ is the automatic differentiation of the D -symmetric NN $\widehat{\mathcal{Z}}_{k+1}$ computed previously at the time step $k + 1$.

end

4.3 The case of mean-field PDEs

We consider in this section the case where the PDE (2.1) is the particles approximation of a McKean-Vlasov control problem with a Bellman equation (2.4) in the Wasserstein space of probability measures as described in Example 2.2. To simplify the presentation, we consider that there is only control on the drift coefficient $\beta(t, x, \mu, a)$ but no control on the diffusion coefficient $\sigma(t, x, \mu)$ and $\sigma_0(t, x, \mu)$ of the McKean-Vlasov equation (2.3). In this case, recall that when the solution $v(t, \mu)$ to this Bellman equation is smooth, an optimal control is given in feedback form by (see [PW17]):

$$\alpha_t^* = \hat{a}(t, X_t^*, \mathbb{P}_{X_t^*}^0, \partial_\mu v(t, \mathbb{P}_{X_t^*}^0)(X_t^*)),$$

where $\hat{a}(t, x, \mu, z)$ is an argmin of $a \in A \mapsto \beta(t, x, \mu, a) \cdot z + f(x, \mu, a)$, and $X^* = X^{\alpha^*}$ is the optimal McKean-Vlasov state process.

Approximation of the optimal control by forward induction of the scheme. As proven in [GPW22a], the solution $(\mathbf{X}, Y, \mathbf{Z})$ to the FBSDE (4.6)-(4.7) provides an approximation with a rate of convergence $1/N$, when N goes to infinity, of the solution v to (2.4), and its L -derivative: $Y_t \simeq v(t, \bar{\mu}(\mathbf{X}_t))$, $N\mathbf{Z}_t^i \simeq \partial_\mu v(t, \bar{\mu}(\mathbf{X}_t))(X_t^i)$. The drift coefficients b_i of the forward particles system \mathbf{X} should be chosen in order to generate from training simulations a suitable exploration of the state space for getting a good approximation of the optimal feedback control. In practice, in a first step, one can choose $b_i(t, \mathbf{x}) = \beta(t, x_i, \bar{\mu}(\mathbf{x}), a_0)$, for some arbitrary value $a_0 \in A$ of the control. After a first implementation of Algorithm 1, we thus have an approximation of $\partial_\mu v(t, \mu)(x)$ at time $t = t_k$, and $\mu = \bar{\mu}(\mathbf{x})$, by $N\widehat{\mathcal{Z}}_k(\mathbf{x}, x)$. Notice however that we solved the PDE along the law of the forward

training process, which is different from the optimally controlled process law, except at the initial time t_0 , where we then get an approximation of the optimal feedback control with

$$(x, \bar{\mu}(\mathbf{x})) \longmapsto \hat{a}(t_0, x, \bar{\mu}(\mathbf{x}), N\hat{\mathcal{Z}}_0(\mathbf{x}, x)).$$

Next, by defining an updated initial drift coefficient as

$$\hat{b}_i(t_0, \mathbf{x}) := \beta(t_0, x_i, \bar{\mu}(\mathbf{x}), \hat{a}(t_0, x_i, \bar{\mu}(\mathbf{x}), N\hat{\mathcal{Z}}_0(\mathbf{x}, x_i))), \text{ for } \mathbf{x} = (x_i)_{i \in \llbracket 1, N \rrbracket}, i = 1, \dots, N,$$

and considering the N -particle discrete-time system $\{\hat{\mathbf{X}}_k^{N_T} = (\hat{X}_k^{i, N_T})_{i \in \llbracket 1, N \rrbracket}, k = 0, \dots, N_T\}$, starting from i.i.d. samples $X_0^i, i = 1, \dots, N$ distributed according to some distribution μ_0 on \mathbb{R}^d , and with dynamics

$$\begin{aligned} \hat{X}_1^{i, N_T} &= X_0^i + \hat{b}_i(t_0, \mathbf{X}_0) \Delta t_0 + \sigma(t_0, X_0^i, \bar{\mu}(\mathbf{X}_0)) \Delta W_0^i, \\ \hat{X}_{k+1}^{i, N_T} &= \hat{X}_k^{i, N_T} + b_i(t_k, \mathbf{X}_k^{N_T}) \Delta t_k + \sigma(t_k, \hat{X}_k^{i, N_T}, \bar{\mu}(\hat{\mathbf{X}}_k^{N_T})) \Delta W_k^i, \end{aligned}$$

for $k = 1, \dots, N_T - 1$, we obtain an approximation of the distribution of the optimal particle mean-field process at time t_1 . Applying the algorithm again between t_1 and t_{N_T} then allows to compute an approximation of the optimal feedback control $\hat{a}(t_1, x, \bar{\mu}(\mathbf{x}), \hat{\mathcal{Z}}_1(\mathbf{x}, x))$ at time t_1 and to update the simulation of \hat{X}_2^{i, N_T} . By induction, we can compute the optimal feedback control at every time step through N_T executions of the scheme.

Approximation of the solution by randomization of the training simulations.

Algorithm 1 provides actually an approximation of $v(t, \mu)$ (resp. $\partial_\mu v(t, \mu)(x)$) at time t_k , and for empirical measures $\mu = \bar{\mu}(\mathbf{x})$, by $\hat{U}_k(\mathbf{x})$ (resp. $N\hat{\mathcal{Z}}_k(\mathbf{x}, x)$). Thus, in order to get an approximation of $v(t_k, \cdot)$ (resp. $\partial_\mu v(t_k, \cdot)(x)$) on the whole Wasserstein space $\mathcal{P}_2(\mathbb{R}^d)$, we need a suitable exploration of $\bar{\mu}(\mathbf{X}_k^{N_T})$ when using the training simulations $\mathbf{X}_k^{N_T}, k = 0, \dots, N_T$. For that purpose, some randomization can first be implemented by randomizing the initial law μ_0 of the forward process. By sampling μ_0 in a compact set K of $\mathcal{P}_2(\mathbb{R}^d)$ for each batch element, such as a family of Gaussian measures for instance, our algorithm will be able to learn the value function $v(0, \mu)$ and its Lions derivative $\partial_\mu v(0, \mu)$ on K . Therefore, instead of solving the PDE several times for each initial law we can run the algorithm only once. This can be useful if we have an uncertainty in the initial law of the problem we aim to solve. It corresponds to learning the solution $v(t_k, \overline{\mu_{k, \ell}})$ on a family of empirical measures corresponding to forward processes $X^{i, (\ell), N_T}, i = 1, \dots, N$, with initial laws $\mu_0^\ell \in K$. Relying on the generalization properties of neural networks, we expect to approximate the value function at time $t_0 = 0$ on K . Furthermore, if the goal is to obtain an approximation of the PDE solution at any time step t_k , the task is more complex. A randomization needs to be performed at each time step t_k by sampling $\mathbf{X}_k^{N_T}$ according to a Gaussian mixture ν_k with random parameters. We then apply Algorithm 1, and expect to learn the solution over measures with regular densities. The updated method is presented in Algorithm 3. If the state space exploration is efficient, the feedback control will be directly available with only one execution of the algorithm, contrarily to the previously described procedure with N_T executions. We should explore the Wasserstein space well enough to learn the value function and its derivative on the unknown law of the optimal process.

Algorithm 3: DBDP scheme with symmetric NN and exploration of Wasserstein space

Initialization: Initialize from the exchangeable function: $\widehat{\mathcal{U}}_{N_T}(\cdot) = G(\cdot)$

for $k = N_T - 1, \dots, 0$ **do**
define random variables

$$L \sim U(1, L_{max}), \varphi_l \sim U(0, 1), \mu_i \sim U(-\mu_{max}, \mu_{max}), (\theta_i)^2 \sim U(0, \sigma_{max}^2)$$

define a random Gaussian mixture ν_k of random density

$$\frac{\sum_{l=1}^L \varphi_l \mathcal{N}(\mu_l, \theta_l^2)}{\sum_{l=1}^L \varphi_l}$$

define N i.i.d. particles X_k^{i, N_T} with law ν_k for $i = 1, \dots, N$,
perform one Euler-Maruyama scheme step

$$X_{k+1}^{i, N_T} = X_k^{i, N_T} + b_i(t_k, \mathbf{X}_k^{N_T}) \Delta t_k + \sigma(t_k, X_k^{i, N_T}, \bar{\mu}(\mathbf{X}_k^{N_T})) \Delta W_k^i,$$

minimize over symmetric NN \mathcal{U}_k , and D -symmetric NN \mathcal{Z}_k , the quadratic loss function (with H as in (4.5)):

$$\begin{aligned} J_1(\mathcal{U}_k, \mathcal{Z}_k) = & \mathbb{E} \left| \widehat{\mathcal{U}}_{k+1}(\mathbf{X}_{k+1}^{N_T}) - \mathcal{U}_k(\mathbf{X}_k^{N_T}) \right. \\ & + H(t_k, \mathbf{X}_k^{N_T}, \mathcal{U}_k(\mathbf{X}_k^{N_T}), \mathcal{Z}_k(\mathbf{X}_k^{N_T})) \Delta t_k \\ & - \sum_{i=1}^N (\mathcal{Z}_k(\mathbf{X}_k^{N_T}, X_k^{i, N_T}))^\top \sigma(t_k, X_k^{i, N_T}, \bar{\mu}(\mathbf{X}_k^{N_T})) \Delta W_k^i \\ & \left. - \sum_{i=1}^N (\mathcal{Z}_k(\mathbf{X}_k^{N_T}, X_k^{i, N_T}))^\top \sigma_0(t_k, X_k^{i, N_T}, \bar{\mu}(\mathbf{X}_k^{N_T})) \Delta W_k^0 \right|^2, \end{aligned}$$

and update $(\widehat{\mathcal{U}}_k, \widehat{\mathcal{Z}}_k)$ as the solution to this minimization problem.

end

5 Numerical results

In the different test cases, for the approximation of the solution v by means of symmetric neutral networks, we used DeepSets.

5.1 A toy example of symmetric PDE in very high dimension

We consider a symmetric semi-linear PDE:

$$\begin{cases} \partial_t v + b \cdot D_{\mathbf{x}} v + \frac{1}{2} \text{tr}(\sigma \sigma^\top D_{\mathbf{x}}^2 v) + f(\mathbf{x}, v, \sigma^\top D_{\mathbf{x}} v) = 0, & (t, \mathbf{x}) \in [0, T] \times \mathbb{R}^N \\ v(T, \mathbf{x}) = \cos(\bar{x}), & \text{with } \bar{x} = \sum_{i=1}^N x_i, \text{ for } \mathbf{x} = (x_1, \dots, x_N) \in \mathbb{R}^N, \end{cases}$$

with $b = 0.2/N$, $\sigma = \frac{I_N}{\sqrt{N}}$,

$$f(\mathbf{x}, y, \mathbf{z}) = (\cos(\bar{x}) + 0.2 \sin(\bar{x})) e^{\frac{T-t}{2}} - \frac{1}{2} (\sin(\bar{x}) \cos(\bar{x}) e^{T-t})^2 + \frac{1}{2N} (y(1_N \cdot \mathbf{z}))^2.$$

so that the PDE solution is exchangeable and given by

$$v(t, \mathbf{x}) = \cos(\bar{x}) \exp\left(\frac{T-t}{2}\right).$$

We solve this PDE in dimension $N = 1000$ by using the deep backward scheme (DBDP) in [HPW20] with 60 time steps, and estimate $U_0 = v(0, 1_N)$ and $Z_0 = D_{\mathbf{x}}v(0, 1_N)$. For the approximation of v , and its gradient $D_{\mathbf{x}}v$, we test with three classes of networks:

- (i) DeepSet \mathcal{U} for v , and AD-DeepSet $D\mathcal{U}$ for $D_{\mathbf{x}}v$ (DeepSets derivative case).
- (ii) DeepSet for v , and DeepDerSet for $D_{\mathbf{x}}v$ (DeepDerSet case)
- (iii) Feedforward for v and $D_{\mathbf{x}}v$ (Feedforward case)

For each of these cases, we use ReLU activation functions for all the networks, and for the feedforward network, we choose 3 layers of 1010 neurons.

Remark 5.1. An alternative to Case (i) is to consider an AD-DeepSet $D\mathcal{U}_1$ for $D_{\mathbf{x}}v$ with \mathcal{U}_1 another DeepSet independent of the one \mathcal{U} used for v . \square

We report the solution in Table 7.

Analytical		(i) AD-DeepSets		(ii) DeepDerSet		(iii) Feedforward	
U_0	Z_0	U_0	Z_0	U_0	Z_0	U_0	Z_0
0.9272	-1.3632	0.9289	-1.2973	0.90140	-1.304	0.6896	-1e-7

Table 7: PDE resolution in dimension 1000 with DBDP scheme [HPW20].

We observe that the results with the feedforward network are not good. This is due to the fact that the feedforward network is not able to approximate correctly the final condition whatever the initial learning rate and the number of epochExt are taken, as already shown in Table 3. In contrast, we see that the AD-DeepSets and DeepDerSet networks give good results but only when the initial learning rate is taken small enough (here we took $1e-5$). Finally, we have tested the Deep BSDE method in [HJW18] with the variation proposed in [CWNMW19] using a network reported in section 3.2.2. The results are unstable and so we do not report them. A direct use of [HJW18] method with a network per time step is impossible to test due the size of the problem but results in lower dimension also indicate some instability directly linked to the initialization of the network.

5.2 A mean-field control problem of systemic risk

We consider a mean-field model of systemic risk introduced in [CFS15]. This model was introduced in the context of mean field games but here we consider a cooperative version. The limiting problem (when the number of banks is large) of the social planner (central bank) is formulated as follows. The log-monetary reserve of the representative bank is governed by the mean-reverting controlled McKean-Vlasov dynamics

$$dX_t = [\kappa(\mathbb{E}[X_t] - X_t) + \alpha_t] dt + \sigma dW_t, \quad X_0 \sim \mu_0,$$

where $\alpha = (\alpha_t)_t$ is the control rate of borrowing/lending to a central bank that aims to minimize the functional cost

$$J(\alpha) = \mathbb{E} \left[\int_0^T \tilde{f}(X_t, \mathbb{E}[X_t], \alpha_t) dt + \tilde{g}(X_T, \mathbb{E}[X_T]) \right] \rightarrow V_0 = \inf_{\alpha} J(\alpha), \quad (5.1)$$

where the running and terminal costs are given by

$$\tilde{f}(x, \bar{x}, a) = \frac{1}{2}a^2 - qa(\bar{x} - x) + \frac{\eta}{2}(\bar{x} - x)^2, \quad \tilde{g}(x, \bar{x}) = \frac{c}{2}(x - \bar{x})^2,$$

for some positive constants $q, \eta, c > 0$, with $q^2 \leq \eta$.

The value function v to the mean-field type control problem (5.1) is solution to the Bellman (semi-linear PDE) equation (2.4) with σ constant, $\sigma_0 \equiv 0$, $r = 0$, and

$$\begin{aligned} h(t, x, \mu, z, \gamma) &= \inf_{a \in \mathbb{R}} \left\{ [\kappa(\mathbb{E}_\mu[\xi] - x) + a]z + \frac{1}{2}a^2 - qa(\mathbb{E}_\mu[\xi] - x) \right\} + \frac{\sigma^2}{2}\gamma + \frac{\eta}{2}(\mathbb{E}_\mu[\xi] - x)^2 \\ &= (\kappa + q)(\mathbb{E}_\mu[\xi] - x)z + \frac{\sigma^2}{2}\gamma + \frac{\eta - q^2}{2}(\mathbb{E}_\mu[\xi] - x)^2 - \frac{z^2}{2}, \end{aligned}$$

and $\mathcal{G}(\mu) = \frac{c}{2}\text{Var}(\mu) := \frac{c}{2}\mathbb{E}_\mu|\xi - \mathbb{E}_\mu[\xi]|^2$ is the variance of the distribution μ (up to $c/2$). Here, we use the notation: $\mathbb{E}_\mu[\varphi(\xi)] := \int \varphi(x)\mu(dx)$.

The finite-dimensional approximation of (5.1) with N -bank model corresponds to the symmetric Bellman semi-linear PDE on $[0, T] \times \mathbb{R}^N$:

$$\partial_t v^N + \sum_{i=1}^N (\kappa + q)(\bar{x} - x_i) \partial_{x_i} v^N + \frac{\sigma^2}{2} \Delta_{\mathbf{x}} v^N + \frac{\eta - q^2}{2N} \sum_{i=1}^N (\bar{x} - x_i)^2 - \frac{N}{2} \sum_{i=1}^N |\partial_{x_i} v^N|^2 = 0, \quad (5.2)$$

for $\mathbf{x} = (x_1, \dots, x_N) \in \mathbb{R}^N$, where we set $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$, and $\Delta_{\mathbf{x}} = \sum_{i=1}^N \partial_{x_i}^2$ is the Laplacian operator. We numerically solve (5.2) with Algorithm 1 described in Section 4. The algorithm is trained with the forward process in \mathbb{R}^N :

$$X_{k+1}^i = X_k^i + \sigma \Delta W_k^i, \quad X_0^i \sim \mu_0, \quad k = 0, \dots, N_T - 1, \quad i = 1, \dots, N.$$

The choice of a null drift for this training process is intuitively justified by the fact that the objective in (5.1) is to incite the log-monetary reserve of the banks to be close to the average of all the other banks, hence we formally expect their drift to be close to zero.

We test our algorithm by increasing N , and compare with the explicit solution of the limiting linear-quadratic McKean-Vlasov control problem (5.1), which is solved via the resolution of a Riccati equation (see [BP19]), and is analytically given by

$$v(t, \mu) = K_t \text{Var}(\mu) + \sigma^2 \int_t^T K_s \, ds,$$

where

$$K_t = -\frac{1}{2} \left[\kappa + q - \sqrt{\Delta} \frac{\sqrt{\Delta} \sinh(\sqrt{\Delta}(T-t)) + (\kappa + q + c) \cosh(\sqrt{\Delta}(T-t))}{\sqrt{\Delta} \cosh(\sqrt{\Delta}(T-t)) + (\kappa + q + c) \sinh(\sqrt{\Delta}(T-t))} \right],$$

with $\sqrt{\Delta} = \sqrt{(\kappa + q)^2 + \eta - q^2}$, and

$$\int_t^T K_s \, ds = \frac{1}{2} \ln \left[\cosh(\sqrt{\Delta}(T-t)) + \frac{\kappa + q + c}{\sqrt{\Delta}} \sinh(\sqrt{\Delta}(T-t)) \right] - \frac{1}{2}(\kappa + q)(T-t).$$

We have tested various approximation symmetric networks and different resolution methods. We list the methods that fail to solve the problem in high dimension:

- First we tried the global resolution method in [EHJ17] by using the network described in paragraph 3.2.2 for v .

In this case, we could not obtain exploitable results.

- Then we decided to use the local resolution method [HPW20] with a DeepDerSet approximation approach for \mathcal{Z} . We found that the method give accurate results in dimension below 100 but with a variance increasing with the dimension. Results were impossible to exploit in dimension 1000. We thus decided to not report the results.
- At last we tested the local resolution methods [HPW20] with classical feedforward networks using tanh or ReLU activation functions. Two variants were tested with $N = 500$: the first one using a network for v^N (we stress the dependence of the solution to the PDE on N) and another network for Dv^N giving values not exploitable, and a second version using a single network for v^N and using automatic differentiation to approximate Dv^N giving a very high bias and a high standard deviation.

Therefore we only report the case where we use a DeepSet network for \mathcal{U} and a second AD-DeepSet network to estimate \mathcal{Z} or a single DeepSet network for \mathcal{U} which is differentiated to approximate \mathcal{Z} .

We test the tanh and ReLU activation function on this test case using the parameters $\sigma = 1$, $\kappa = 0.6$, $q = 0.8$, $c = 2$, $\eta = 2$, $T = 1$. We report v^N estimated with different values of N_T and N at time $t = 0$ and $x = 0$ so using $\mu_0 = \delta_0$ on the figures below. The theoretical solution obtained when N goes to infinity is 0.29244.

We use a batch size equal to 200, a number of gradient iteration equal to 30000 for the resolution to project the terminal condition on the network and 6000 gradient iterations for other resolutions. The initial learning rate is taken equal to $1e - 4$ at the first resolution and $5e - 5$ for other resolutions. The learning is taken decreasing linearly with gradient iterations to $5e - 6$.

On figure 8, we give the results obtained with ReLU activation function using a DeepSet network for \mathcal{U} and a second AD-DeepSet network to approximate \mathcal{Z} . The convergence is steady as N_T grows and as the dimension grows leading to a very accurate result for $N_T = 60$ and $N = 1000$.

N_T	Dimension N	Averaged	Std	Relative error
15	10	0.259	0.0029	0.11
15	100	0.2871	0.0016	0.018
15	500	0.2866	0.00179	0.019
15	1000	0.2877	0.00238	0.016
30	10	0.265	0.004	0.09
30	100	0.2892	0.2892	0.010
30	500	0.2897	0.00153	0.009
30	1000	0.2899	0.00146	0.0084
60	10	0.2655	0.0045	0.092
60	100	0.2894	0.0012	0.010
60	500	0.2894	0.0027	0.010
60	1000	0.2916	0.0014	0.0025

Table 8: Systemic risk with ReLU activation function, a DeepSet network for \mathcal{U} and a second AD-DeepSet network to estimate \mathcal{Z} .

Using a ReLU activation function, a single network for \mathcal{U} which is differentiated to approximate \mathcal{Z} , we get the results in figure 9. The convergence is still steady but results are not as good as in table 8.

N_T	Dimension N	Averaged	Std	Relative error
15	10	0.2530	0.0074	0.1346
15	100	0.27968	0.0051	0.043
15	500	0.2938	0.0067	0.0049
15	1000	0.3084	0.0253	0.054
30	10	0.2494	0.0074	0.1471
30	100	0.2756	0.00677	0.057
30	500	0.2885	0.0127	0.013
30	1000	0.2860	0.009	0.02
60	10	0.2519	0.0037	0.138
60	100	0.28253	0.0047	0.033
60	500	0.28329	0.0108	0.03
60	1000	0.2881	0.0043	0.014

Table 9: Systemic risk with ReLU activation function, a single DeepSet network for \mathcal{U} which is differentiated to approximate \mathcal{Z} .

The replace the ReLU activation function by a tanh one using two networks and the results are given in table 10. The convergence is not steady and increasing to much N or N_T worsen to results : it shows the importance of the activation function in this method.

N_T	N	Averaged	Std	Relative error
15	10	0.2678	0.0061	0.08
15	100	0.28858	0.0144	0.013
15	500	0.2491	0.027	0.14
15	1000	0.27401	0.0127	0.063
30	10	0.2725	0.0052	0.068
30	100	0.2959	0.0161	0.012
30	500	0.2577	0.01568	0.118
30	1000	0.320	0.0030	0.096
60	10	0.2739	0.0049	0.063
60	100	0.2924	0.0309	0.0001
60	500	0.3158	0.00297	0.079
60	1000	0.2210	0.004	0.24

Table 10: Systemic with tanh activation function, a DeepSet network for \mathcal{U} and a second AD-DeepSet network to estimate \mathcal{Z} .

At last we do not report the test obtained using a ReLU activation function for the first network and a tanh one for the second network given results far better than in table 10 but not as good as in tables 8 and 9. We also test the accuracy of our algorithm for approximating the L -derivative of the solution, which is here explicitly given by

$$\partial_\mu v(t, \mu)(x) = 2K_t(x - \mathbb{E}_\mu[\xi]).$$

For this purpose, using N_T steps, we solve the same problem on $[t, T]$, starting at $t = \frac{T}{2}$ with a distribution μ_0 equal to real distribution of the solution of (5.1) taken at date t . After training, we plot $x \mapsto N\widehat{\mathcal{Z}}(\mathbf{X}_t, x)$, where $\mathbf{X}_t \sim \mu_0^{\otimes N}$, and compare to the analytic solution: $x \mapsto \partial_\mu v(t, \mu_0)(x)$. Some graphs are reported in figure 4, which shows the accuracy of the approximation.

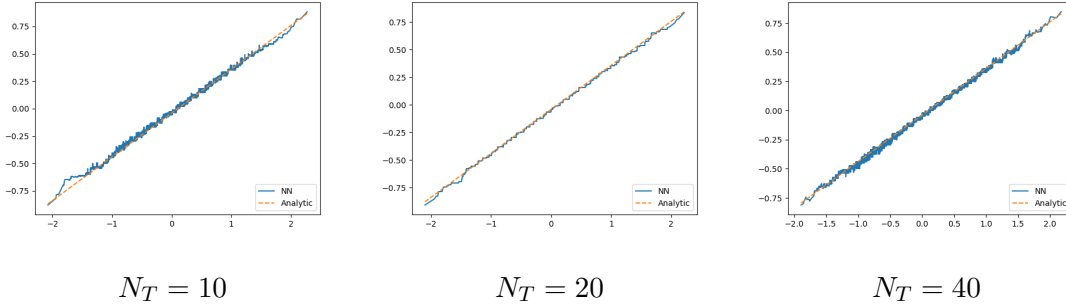
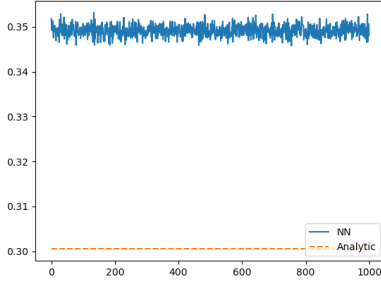


Figure 4: Resolution on $[0.5, 1]$ in dimension $N = 500$: analytic Lions derivatives versus $N\widehat{\mathcal{Z}}$ estimated by the network. DeepSet network for \mathcal{U} , AD-DeepSet for \mathcal{Z} . ReLU activation function.

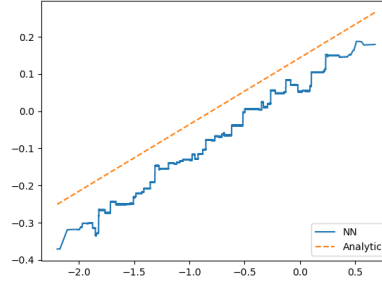
As mentioned in Section 4.3, in theory, the proposed methodology should learn the solution for any initial law μ_0 in the space of measures so that we should be able to solve the problem in infinite dimension. We test our algorithm by sampling μ_0 in the following way: for a sample j , we pick up a mean $\hat{M} \in [-1, 1]$ and a standard deviation $\sigma \in [0.2, 1]$ with a uniform law. Then $X_0^{i,j} \sim \mathcal{N}(\hat{M}, \sigma^2)$ $i = 1, \dots, N$ and as before we use the forward process:

$$X_{k+1}^{i,j} = X_k^{i,j} + \sigma \Delta W_k^{i,j}, \quad k = 0, \dots, N_T - 1, \quad i = 1, \dots, N.$$

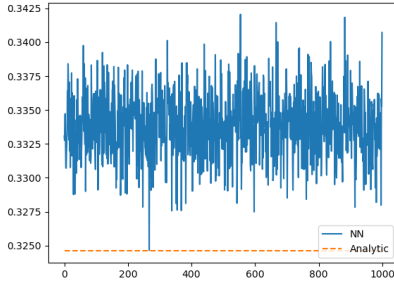
After the training part, we try to recover the initial solution and the initial Lions derivative for a given μ_0 following a gaussian law. Results are given on figure 5. The Lions derivative is relatively correctly calculated but the initial value can get an error around 15%.



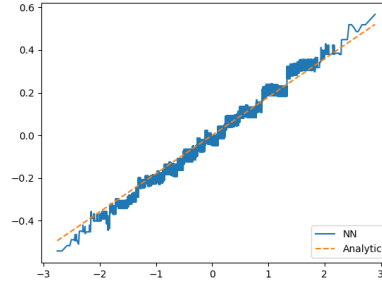
Solution for $\mu_0 = \mathcal{N}(-0.8, 0.09)$



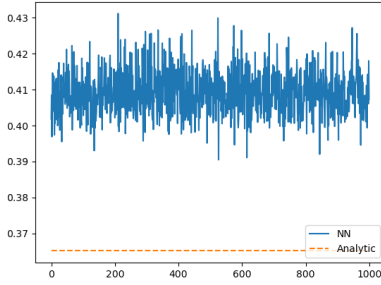
Lions derivative for $\mu_0 = \mathcal{N}(-0.8, 0.09)$



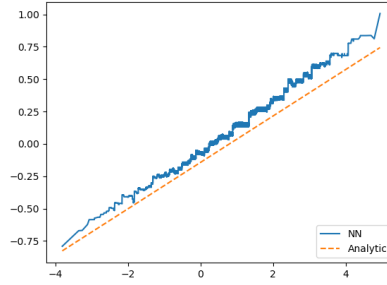
Solution for $\mu_0 = \mathcal{N}(0., 0.36)$



Lions derivative for $\mu_0 = \mathcal{N}(0., 0.36)$



Solution for $\mu_0 = \mathcal{N}(0.8, 0.81)$



Lions derivative for $\mu_0 = \mathcal{N}(0.8, 0.81)$

Figure 5: Solution and Lions derivative after a single training, with $N = 500$, $N_T = 30$, with ReLU activation function, a single DeepSet network for \mathcal{U} which is differentiated to approximate \mathcal{Z} . For the solution, the x -axis corresponds to the sample number and the y -axis is the value of the estimated solution. For the Lions derivative, the x -axis is the state space and the y -axis is the value of the derivative.

More generally, if we want to solve the PDE at each time step in the Wasserstein space, we can use Algorithm 3. In order to illustrate the exploration of the Wasserstein space we plot in Figure 6 the graphs of $(X^i, N\hat{\mathcal{Z}}(\mathbf{X}, X^i))$, $i = 1, \dots, N$, vs $X^i \mapsto \partial_\mu v(t, \mathbb{P}_{X^i})(X^i)$, when $X_t^i \rightsquigarrow$ random mixture of Gaussian laws, for $N = 300$, $N_T = 30$. We observe that we are able to estimate correctly the Lions derivative of the solution (and therefore the optimal control) on several probability measures through a randomized training. Concerning the solution itself, we observe similar behavior as in Figure 5 with an error of order 10-15% so

we do not show the plots. Further numerical studies are left to future research to improve the estimation of the solution with the randomization procedure.

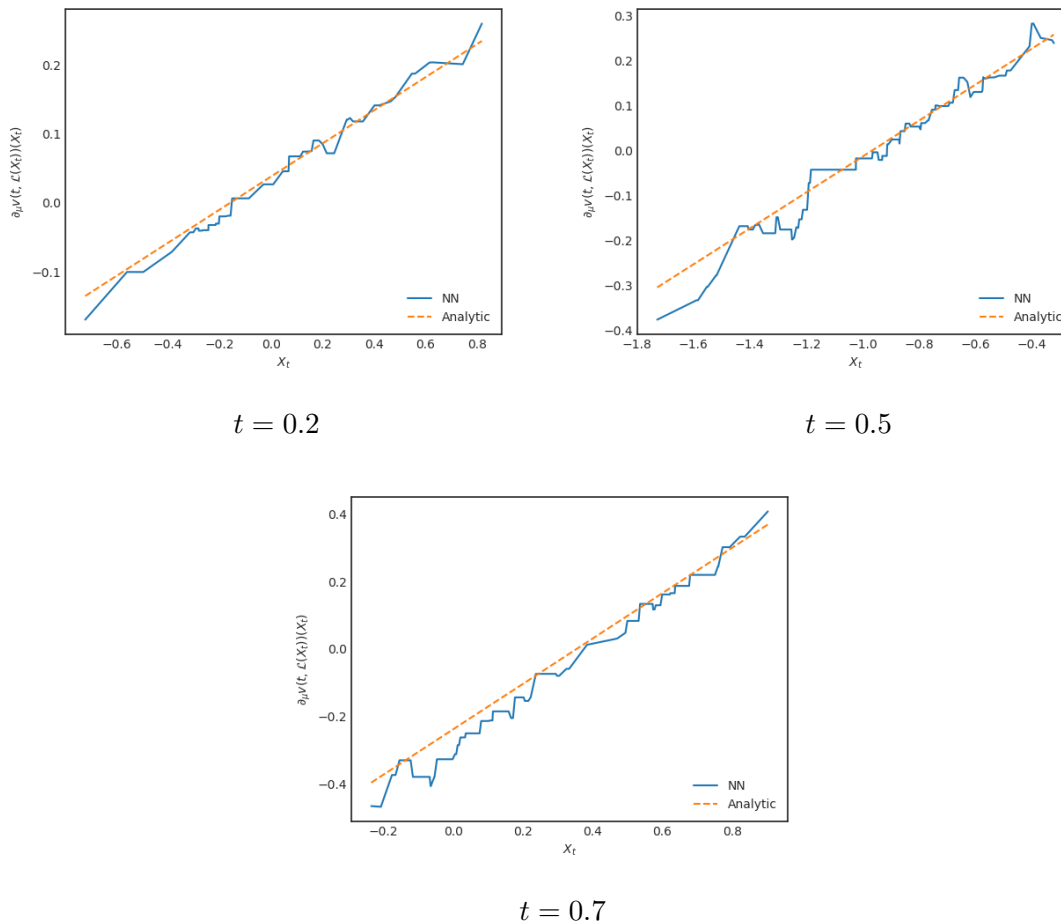


Figure 6: Analytic Lions derivative versus $N\mathcal{Z}$ estimated by the network. Dimension $N = 300$, number of time steps $N_T = 30$. We use a DeepSet network for \mathcal{U} with ReLU activation functions, and \mathcal{Z} its automatic derivative.

5.3 Mean-variance problem

We consider the celebrated Markowitz portfolio selection problem where an investor can invest at any time t an amount α_t in a risky asset (assumed for simplicity to follow a Black-Scholes model with constant rate of return β and volatility $\nu > 0$), hence generating a wealth process $X = X^\alpha$ with dynamics

$$dX_t = \alpha_t \beta dt + \alpha_t \nu dW_t, \quad 0 \leq t \leq T, \quad X_0 = x_0 \in \mathbb{R}.$$

The goal is then to minimize over portfolio control α the mean-variance criterion:

$$J(\alpha) = \lambda \text{Var}(X_T^\alpha) - \mathbb{E}[X_T^\alpha],$$

where $\lambda > 0$ is a parameter related to the risk aversion of the investor. Due to the presence of the variance term Var in the criterion, the Markowitz problem falls into the class of

McKean-Vlasov control problems, and the associated value function v satisfies the Bellman equation (2.4) on $[0, T] \times \mathcal{P}_2(\mathbb{R})$ with $r = 0$, $\sigma_0 \equiv 0$,

$$\begin{cases} h(x, \mu, z, \gamma) = \inf_{a \in \mathbb{R}} [za\beta + \frac{1}{2}\gamma a^2 \nu^2] = -\frac{Rz^2}{2\gamma}, & z \in \mathbb{R}, \gamma > 0, \\ \mathcal{G}(\mu) = \lambda \mathbb{E}_\mu |\xi - \mathbb{E}_\mu[\xi]|^2 - \mathbb{E}_\mu[\xi], & \mu \in \mathcal{P}_2(\mathbb{R}), \end{cases}$$

where we set $R := \beta^2/\nu^2$.

The associated finite-dimensional PDE with N particles is given by

$$\begin{cases} \partial_t v^N - \frac{R}{2} \sum_{i=1}^N \frac{(D_{x_i} v^N)^2}{D_{x_i}^2 v^N} = 0, & t \in [0, T], \mathbf{x} = (x_1, \dots, x_N) \in (\mathbb{R}^d)^N, \\ v^N(T, \mathbf{x}) = \mathcal{G}(\bar{\mu}(\mathbf{x})). \end{cases} \quad (5.3)$$

We refer to [IP19] for the McKean-Vlasov approach to Markowitz mean-variance problems (in a more general context), and we recall that the solution to the Bellman equation is given by

$$\begin{aligned} v(t, \mu) &= \lambda e^{-R(T-t)} \mathbb{E}_\mu |\xi - \mathbb{E}_\mu[\xi]|^2 - \mathbb{E}_\mu[\xi] - \frac{1}{4\lambda} [e^{R(T-t)} - 1] \\ \partial_\mu v(t, \mu)(x) &= 2\lambda e^{-R(T-t)} (x - \mathbb{E}_\mu[\xi]) - 1, \quad \partial_x \partial_\mu v(t, \mu)(x) = 2\lambda e^{-R(T-t)} \end{aligned} \quad (5.4)$$

and in particular $V_0 := \inf_\alpha J(\alpha) = v(0, \delta_{x_0}) = -x_0 - \frac{1}{4\lambda} [e^{RT} - 1]$. Moreover, the optimal portfolio strategy is given by

$$\begin{aligned} \alpha_t^* &= \hat{a}(t, X_t^*, \mathbb{E}[X_t^*]) := -\frac{\beta}{\nu^2} \left[X_t^* - \mathbb{E}[X_t^*] - \frac{e^{R(T-t)}}{2\lambda} \right] \\ &= -\frac{\beta}{\nu^2} \left[X_t^* - x_0 - \frac{e^{RT}}{2\lambda} \right], \quad 0 \leq t \leq T, \end{aligned} \quad (5.5)$$

where $X^* = X^{\alpha^*}$ is the optimal wealth process.

We test our Algorithm 2 described in Section 4.2 with the training of the forward process

$$X_{k+1}^{i,N,\pi} = X_k^{i,N,\pi} + \frac{R}{2\lambda} \Delta t_k + \frac{\sqrt{R}}{2\lambda} \Delta W_k^i, \quad X_0^i = x_0, \quad k = 0, \dots, N_T - 1, \quad i = 1, \dots, N,$$

which is the time discretization of the wealth process for a constant portfolio strategy $\alpha_t = \beta/(2\nu^2\lambda)$, which is known to be optimal for the exponential utility function $U(x) = -e^{-2\lambda x}$. This corresponds to the choice of $b_i = R/(2\lambda)$ and $\sigma_{ij} = \sqrt{R}/(2\lambda)$. Here, notice that $\partial_\mu G(\mu)(x) = 2\lambda(x - \mathbb{E}_\mu[\xi]) - 1$, and we then use for the initialization at terminal step N_T , the DeepDerSet function $\mathcal{Z}_{N_T}((x_i)_i, x) = 2\lambda(x - \frac{1}{N} \sum_i x_i) - 1$ (corresponding to the average function $\mathfrak{s}((x_i)_i) = \frac{1}{N} \sum_i x_i$), which yields the automatic differentiation $D\mathcal{Z}_{N_T}((x_i)_i, x) = 2\lambda(1 - \frac{1}{N})$.

We choose the parameters $\beta = 0.15$, $\nu = 0.35$, $\lambda = 1$, and the quantile at 99.9% for the truncation in scheme [PWG21], and report the results in Table 11. The optimization parameters are the same as in the semi linear case, except the batch size taken equal to 50 and the number of gradient iterations after first step taken equal to 4000. We use a ReLU Deepset for \mathcal{U} and a AD-Deepset with a tanh activation function for \mathcal{Z} . Remark that in this case it is not possible to use a ReLU activation function for the second network.

N_T	Dimension N	Averaged	Std	Relative error
10	10	-1.0561	0.001	0.005
10	100	-1.0522	0.0008	0.0017
20	10	-1.0570	0.0008	0.006
20	100	-1.0520	0.0007	0.0015
30	10	-1.0578	0.0011	0.007
30	100	-1.0535	0.0021	0.0029

Table 11: Estimate of $\mathbb{E}[v^N(0, X_0^1, \dots, X_0^N)]$ with a deterministic initial condition $X_0 = 1$, $T = 1$, $\sigma = 1$. Average and standard deviation observed over 10 independent runs are reported. The theoretical solution is -1.0504058 when $N, N_T \rightarrow +\infty$.

Moreover, we test the accuracy of the control approximation. We solve the PDE from $T/2$ to T starting with the optimal distribution of the wealth at $T/2$, which is given by:

$$\log\left(\frac{X_{T/2}}{\eta(T/2)} - x_0 - e^{\frac{RT}{2\lambda}}\right) \sim \mathcal{N}(0, \kappa(T/2)),$$

with

$$\eta(t) = -e^{\frac{R(T-t)}{2\lambda}}, \quad \kappa^2(t) = \log\left(\frac{e^{R(T-t)}(e^{RT} - e^{R(T-t)})}{4\lambda^2\eta(t)^2} + 1\right),$$

and we calculate the solution obtained at date $T/2$ and the control obtained solving the PDE (5.3) that we can compare to the analytical solution given by (5.4) and (5.5). After training, using $n_s = 50$ samples of $X \in (\mathbb{R}^N)^{n_s}$ following the law of $X_{T/2}$, we calculate the control obtained for each sample in each of the dimension. After sorting X in a one dimensional array, We plot the result obtained on Figures 7–8. For the solution, the x -axis corresponds to the sample number and the y -axis is the value of the estimated solution. For the other plots, the x -axis is the state space and the y -axis is the value of the corresponding function.

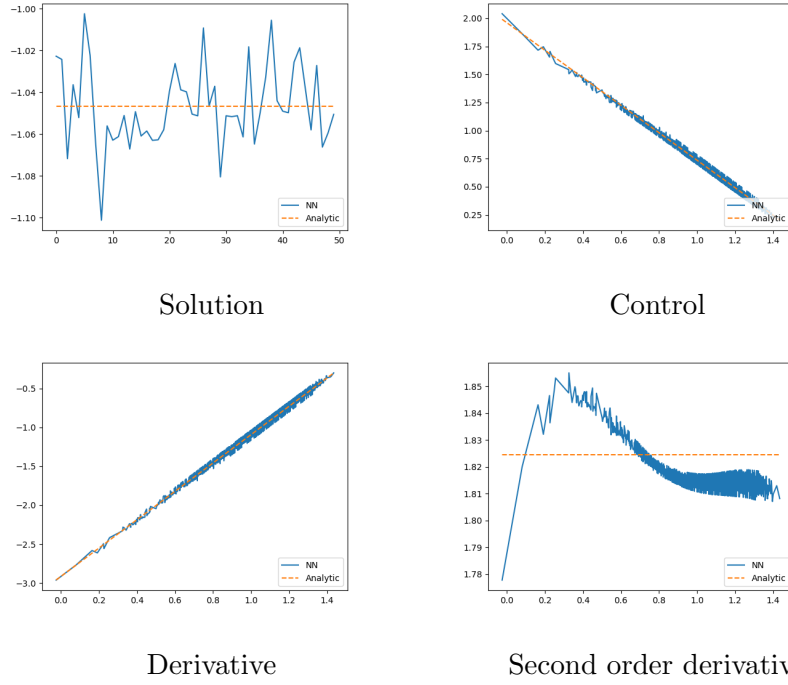


Figure 7: Solution and control obtained on the mean variance case at $\frac{T}{2}$ in dimension 100 with 20 time steps comparing analytic solution to the calculated one (NN). Truncation factor equal to 0.999.

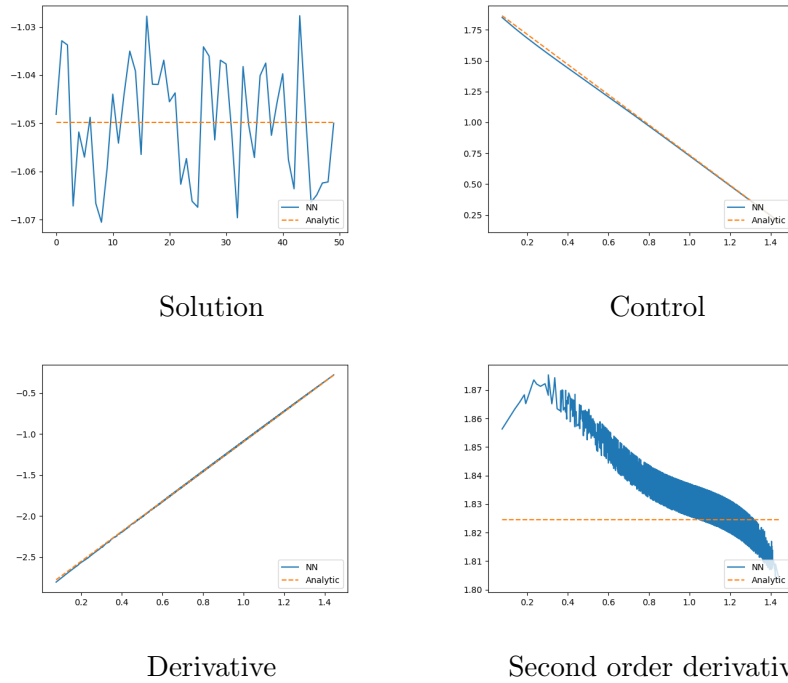


Figure 8: Solution and control obtained on the mean variance case at $\frac{T}{2}$ in dimension 300 with 20 time steps comparing analytic solution to the calculated one (NN). Truncation factor: 0.999.

5.4 A min/max Linear quadratic mean-field control problem

We consider a mean-field model in which the dynamics is linear and the running cost is quadratic in the position, the control and the expectation of the position. The terminal cost is encourages to be close to one of two targets. This type of model is inspired by the min-LQG problem of [SMLN15]. More precisely, we consider the following controlled McKean-Vlasov dynamics

$$dX_t = [AX_t + \bar{A}\mathbb{E}[X_t] + B\alpha_t] dt + \sigma dW_t, \quad X_0 \sim \mu_0,$$

where $\alpha = (\alpha_t)_t$ is the control, and the agent aims to minimize the functional cost

$$J(\alpha) = \mathbb{E} \left[\int_0^T f(X_t, \mathbb{E}[X_t], \alpha_t) dt + g(X_T) \right] \rightarrow V_0 = \inf_{\alpha} J(\alpha), \quad (5.6)$$

where the running and terminal costs are given by

$$f(x, \bar{x}, a) = \frac{1}{2} (Qx^2 + \bar{Q}(x - S\bar{x})^2 + Ra^2), \quad g(x) = \min \{ |x - \xi_1|^2, |x - \xi_2|^2 \},$$

for some non-negative constants Q, \bar{Q}, S, R , and two real numbers ξ_1 and ξ_2 .

The value function to the mean-field type control problem (5.6) is solution to the Bellman (semi-linear PDE) equation (2.4) with $r = 0$, and

$$\begin{aligned} h(x, \mu, z, \gamma) &= \inf_{a \in \mathbb{R}} \left\{ [Ax + \bar{A}\mathbb{E}_{\mu}[\xi] + Ba]z + \frac{1}{2} (Qx^2 + \bar{Q}(x - S\mathbb{E}_{\mu}[\xi])^2 + Ra^2) \right\} + \frac{\sigma^2}{2} \gamma \\ &= [Ax + \bar{A}\mathbb{E}_{\mu}[\xi]]z - \frac{B^2}{2R} z^2 + \frac{1}{2} (Qx^2 + \bar{Q}(x - S\mathbb{E}_{\mu}[\xi])^2) + \frac{\sigma^2}{2} \gamma, \end{aligned}$$

where the minimizer in the above inf is given by $a = -\frac{B}{R}z$, and the terminal condition $\mathcal{G}(\mu) = \mathbb{E}_{\mu} \min \{ |\xi - \xi_1|^2, |\xi - \xi_2|^2 \}$ is the expected minimal distance to one of the targets ξ_1, ξ_2 .

For the sake of illustration, we present several test cases.

The targets are at $\xi_1 = 0.25$ and $\xi_2 = 1.75$. Here we used $A = \bar{A} = 0, B = 1, Q = 0, \bar{Q} = S = R = 1$, and a time horizon $T = 0.5$. The initial distribution μ_0 is a Gaussian $\mathcal{N}(x_0, \vartheta_0^2)$. We consider the following test cases:

1. $\sigma = 0.3, x_0 = 1, \vartheta_0 = 0.2$,
2. $\sigma = 0.5, x_0 = 0.625, \vartheta_0 = \sqrt{0.2}$,
3. $\sigma = 0.3, x_0 = 0.625, \vartheta_0 = \sqrt{0.2}$,
4. $\sigma = 0.3, x_0 = 0.625, \vartheta_0 = \sqrt{0.4}$.

References are given in table 12: they are calculated by the PDE method in [ACD10] (in the context of mean field games; see [AL15] for the adaptation to the PDE system arising in mean field control) with step size in space and time of size 10^{-3} , and the neural network method referred to as Algorithm 1 in [CL19] with $N = 10000$ and $N_T = 50$.

Case	Benchmark	Global
1	0.2256	0.2273(0.004)
2	0.2085	0.2098(0.006)
3	0.1734	0.1742(0.005)
4	0.2276	0.2300(0.009)

Table 12: Min-LQC example reference solutions : benchmark solution estimated by finite difference scheme and Algorithm 1 in [CL19] with $N = 10000$, $N_T = 50$, 10 neurons and 3 hidden layers, tanh activation function, average on 10 runs.

In table 13, we give the results obtained with different time discretization and dimension for the DPBD scheme using ReLU activation functions with a DeepSet network for \mathcal{U} and a second AD-DeepSet network to estimate \mathcal{Z} . Results are very good except for test case 1 where a small bias appears.

Case	$N = 100, N_T = 30$	$N = 100, N_T = 60$	$N = 500, N_T = 30$	$N = 500, N_T = 60$
1	0.2370(0.013)	0.2382(0.012)	0.2446(0.013)	0.2495(0.09)
2	0.2088(0.002)	0.2092(0.001)	0.2106(0.003)	0.2105(0.003)
3	0.1774(0.007)	0.1784(0.005)	0.1819(0.005)	0.1785(0.008)
4	0.2279(0.005)	0.2264(0.005)	0.2292(0.006)	0.2274(0.006)

Table 13: Min-LQC example with DPBD scheme using ReLU activation functions with a DeepSet network for \mathcal{U} and a second AD-DeepSet network to estimate \mathcal{Z} , average on 10 runs, standard deviation in parenthesis.

In table 14, we give the same results using a single network. Here the results are very good for all test cases. Using two networks, the algorithm certainly face difficulties to approximate the derivatives near maturities which is not required using a single network.

Case	$N = 100, N_T = 30$	$N = 100, N_T = 60$	$N = 500, N_T = 30$	$N = 500, N_T = 60$
1	0.2289(0.0006)	0.2271(0.001)	0.2290(0.0004)	0.2271(0.0008)
2	0.2083(0.0008)	0.2086(0.0007)	0.2097(0.0008)	0.2089(0.0004)
3	0.1740(0.001)	0.1740(0.001)	0.1742(0.0004)	0.1729(0.0007)
4	0.2276(0.001)	0.2310(0.003)	0.2282(0.0008)	0.2278(0.001)

Table 14: Min-LQC example with DBDP scheme using ReLU activation functions and a single DeepSet network for \mathcal{U} which is differentiated to approximate \mathcal{Z} , average on 10 run, standard deviation in parenthesis.

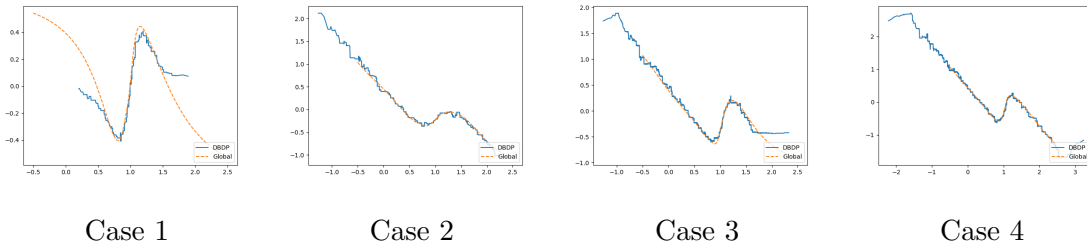


Figure 9: Control calculated at $t = 0$ for Min-LQC examples: comparison DBDP using a single DeepSet network with $N_T = 50$, $N = 500$ and global approximation.

References

- [ACD10] Y. Achdou and I. Capuzzo-Dolcetta. “Mean field games: numerical methods”. In: *SIAM J. Numer. Anal.* 48.3 (2010), pp. 1136–1162.
- [AL15] Y. Achdou and M. Laurière. “On the system of partial differential equations arising in mean field type control”. In: *Discrete Contin. Dyn. Syst.* 35.9 (2015), pp. 3879–3900.
- [Bec+20] C. Beck, M. Hutzenthaler, A. Jentzen, and B. Kuckuck. “An overview on deep learning-based approximation methods for partial differential equations”. In: *arXiv:2012.12348v1* (2020).
- [BP19] M. Basei and H. Pham. “A weak martingale approach to linear-quadratic McKean-Vlasov stochastic control problem”. In: *Journal of Optimization Theory and Applications* 181.2 (2019), pp. 347–382.
- [BRT20] B. Bloem-Reddy and Y.W. Teh. “Probabilistic symmetries and invariant neural networks”. In: *Journal of Machine Learning Research* 21 (2020), pp. 1–61.
- [CD18a] R. Carmona and F. Delarue. *Probabilistic Theory of Mean Field Games: vol. I, Mean Field FBSDEs, Control, and Games*, Springer, 2018.
- [CD18b] R. Carmona and F. Delarue. *Probabilistic Theory of Mean Field Games: vol. II, Mean Field FBSDEs, Control, and Games*, Springer, 2018.
- [CFS15] R. Carmona, J.-P. Fouque, and L. Sun. “Mean field games and systemic risk”. In: *Commun. Math. Sci.* 13.4 (2015), pp. 911–933.
- [CL19] R. Carmona and M. Laurière. “Convergence analysis of machine learning algorithms for the numerical solution of mean field control and games: II—the finite horizon case”. In: *arXiv:1908.01613, to appear in The Annals of Applied Probability* (2019).
- [CUH16] D.-A. Clevert, T. Unterthiner, and S. Hochreiter. “Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)”. In: *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*. Ed. by Y. Bengio and Y. LeCun. 2016.

- [CWNMW19] Q. Chan-Wai-Nam, J. Mikael, and X. Warin. “Machine Learning for Semi Linear PDEs”. In: *Journal of Scientific Computing* 79 (2019), pp. 1667–1712.
- [Dje20] F Djete. “Extended mean-field control problem: a propagation of chaos result”. In: *arXiv:2006.12996* (2020).
- [EHJ17] W. E, J. Han, and A. Jentzen. “Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations”. In: *Communications in Mathematics and Statistics* 5.4 (2017), pp. 349–380.
- [FZ20] J.-P. Fouque and Z. Zhang. “Deep Learning Methods for Mean Field Control Problems with Delay”. In: *Frontiers in Applied Mathematics and Statistics* 6 (2020).
- [GMS21] W. Gangbo, S. Mayorga, and A. Swiech. “Finite Dimensional Approximations of Hamilton–Jacobi–Bellman Equations in Spaces of Probability Measures”. In: *SIAM Journal on Mathematical Analysis* 53.2 (2021), pp. 1320–1356.
- [GMW19] M. Germain, J. Mikael, and X. Warin. “Numerical resolution of McKean-Vlasov FBSDEs using neural networks”. In: *arXiv:1909.12678* (2019).
- [GPW22a] M. Germain, H. Pham, and X. Warin. “Convergence analysis of particles approximation of PDEs in Wasserstein space”. In: *arXiv:2103.00837, to appear in Journal of Applied Probability* 59.4 (2022).
- [GPW22b] M. Germain, H. Pham, and X. Warin. “Neural networks based algorithms for stochastic control and PDEs in finance”. In: *arXiv:2101.08068 to appear in Machine Learning And Data Sciences For Financial Markets: A Guide To Contemporary Practices*. Ed. by A. Capponi and C.A. Lehalle. Cambridge University Press, 2022.
- [HCM06] M. Huang, P. Caines, and R. Malhamé. “Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle”. In: *Communication in Information and Systems* 3 (2006), pp. 221–252.
- [HJW18] J. Han, A. Jentzen, and E. Weinan. “Solving high-dimensional partial differential equations using deep learning”. In: *Proceedings of the National Academy of Sciences* 115.34 (2018), pp. 8505–8510.
- [Hor91] K. Hornik. “Approximation Capabilities of Multilayer Feedforward Networks”. In: *Neural Networks* 4 (1991), pp. 251–257.
- [HPW20] C. Huré, H. Pham, and X. Warin. “Deep backward schemes for high-dimensional nonlinear PDEs”. In: *Mathematics of Computation* 89.324 (2020), pp. 1547–1579.
- [Hut+20] M. Hutzenthaler, A. Jentzen, T. Kruse, and T.A. Nguyen. “A proof that rectified deep neural networks overcome the curse of dimensionality in the numerical approximation of semilinear heat equation”. In: *SN partial differential equations and applications* 1.10 (2020), pp. 1–34.

- [IP19] A. Ismail and H. Pham. “Robust Markowitz mean-variance portfolio selection under ambiguous covariance matrix”. In: *Mathematical Finance* 29.174-207 (2019).
- [KB14] D. P. Kingma and J. Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [Lab12] P.H. Labordère. “Counterparty risk valuation: a marked branching diffusion approach”. In: *Hal-00677348* (2012).
- [Lac17] D. Lacker. “Limit theory for controlled McKean-Vlasov dynamics”. In: *SIAM Journal on Control and Optimization* 55.3 (2017), pp. 1641–1672.
- [LL07] J.M. Lasry and P.L. Lions. “Mean field games”. In: *Japanese Journal of Mathematics* 2 (2007), pp. 229–260.
- [PP90] E. Pardoux and S. Peng. “Adapted solution of a backward stochastic differential equation”. In: *Systems & Control Letters* 14.1 (1990), pp. 55–61.
- [PW17] H. Pham and X. Wei. “Dynamic programming for optimal control of stochastic McKean-Vlasov dynamics”. In: *SIAM Journal on Control and Optimization* 55.2 (2017), pp. 1069–1101.
- [PWG21] H. Pham, X. Warin, and M. Germain. “Neural networks-based backward scheme for fully nonlinear PDEs.” In: *SN Partial Differential Equations and Applications* 2.16 (2021).
- [RQ+17] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. “PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation”. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017, pp. 77–85.
- [Rut+20] L. Ruthotto, S. J. Osher, W. Li, L. Nurbekyan, and S. W. Fung. “A machine learning framework for solving high-dimensional mean field game and mean field control problems”. In: *Proc. Natl. Acad. Sci. USA* 117.17 (2020), pp. 9183–9193.
- [SMLN15] R. Salhab, R. P. Malhamé, and J. Le Ny. “A dynamic game model of collective choice in multi-agent systems”. In: *2015 IEEE 54th Annual Conference on Decision and Control (CDC)*. 2015, pp. 4444–4449.
- [Smu11] J. Smulevici. “On the area of the symmetry orbits of cosmological spacetimes with toroidal or hyperbolic symmetry”. In: *Analysis and PDE* 4.2 (2011), pp. 191–245.
- [Wag+19] E. Wagstaff, F. Fuchs, M. Engelcke, I. Posner, and M. A. Osborne. “On the Limitations of Representing Functions on Sets”. In: ed. by Kamalika Chaudhuri and Ruslan Salakhutdinov. Vol. 97. Proceedings of Machine Learning Research. 2019, pp. 6487–6494.
- [Wyk08] S. Van Wyk. “Partial differential equations and quantum mechanics”. In: *Computer solution in Physics* (2008), pp. 99–139.
- [Zah+17] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. R Salakhutdinov, and A. J Smola. “Deep Sets”. In: *Advances in Neural Information Processing Systems* 30. Ed. by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Curran Associates, Inc., 2017, pp. 3391–3401.