



Solving high-dimensional eigenvalue problems using deep neural networks: A diffusion Monte Carlo like approach



Jiequn Han ^{a,*}, Jianfeng Lu ^{b,c}, Mo Zhou ^b

^a Department of Mathematics, Princeton University, Princeton, NJ 08544, USA

^b Department of Mathematics, Duke University, Durham, NC 27708, USA

^c Department of Physics and Department of Chemistry, Duke University, Durham, NC 27708, USA

ARTICLE INFO

Article history:

Received 15 February 2020

Received in revised form 16 July 2020

Accepted 16 August 2020

Available online 20 August 2020

Keywords:

Diffusion Monte Carlo

Deep neural networks

Eigenvalue problem

Schrödinger equation

ABSTRACT

We propose a new method to solve eigenvalue problems for linear and semilinear second order differential operators in high dimensions based on deep neural networks. The eigenvalue problem is reformulated as a fixed point problem of the semigroup flow induced by the operator, whose solution can be represented by Feynman-Kac formula in terms of forward-backward stochastic differential equations. The method shares a similar spirit with diffusion Monte Carlo but augments a direct approximation to the eigenfunction through neural-network ansatz. The criterion of fixed point provides a natural loss function to search for parameters via optimization. Our approach is able to provide accurate eigenvalue and eigenfunction approximations in several numerical examples, including Fokker-Planck operator and the linear and nonlinear Schrödinger operators in high dimensions.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

Many fundamental problems in scientific computing can be reduced to the computation of eigenvalues and eigenfunctions of an operator. One primary example is the electronic structure calculations, namely, computing the leading eigenvalue and eigenfunction of the Schrödinger operator. If the dimension of the state variable is low, one can use classical approaches, such as the finite difference method or spectral method, to discretize the operator and to solve the eigenvalue problem. However, these conventional, deterministic approaches suffer from the so-called curse of dimensionality, when the underlying dimension becomes high, since the number of degrees of freedom grows exponentially as the dimension increases.

For high-dimensional problems, commonly arising from quantum mechanics, statistical mechanics, and finance applications, stochastic methods become more attractive and in many situations the only viable option. In the context of quantum mechanics, two widely used approaches for high-dimensional eigenvalue problems are the variational Monte Carlo (VMC) and diffusion Monte Carlo (DMC) methods [1–6]. These two approaches deal with the high dimensionality via different strategies. VMC relies on leveraging chemical knowledge to propose an ansatz of the eigenfunction (wavefunction in the context of quantum mechanics) with parameters to be optimized under the variational formulation of the eigenvalue problem. The Monte Carlo approach is used to approximate the gradient of the energy with respect to the parameters at each optimization iteration step. On the other hand, DMC represents the density of the eigenfunction with a collection of particles that follow the imaginary time evolution given by the Schrödinger operator, via a Feynman-Kac representation of

* Corresponding author.

E-mail address: jiequnh@princeton.edu (J. Han).

the semigroup. It can be understood as a generalization of the classical power method from finite-dimensional matrices to infinite-dimensional operators. In electronic structure calculations, DMC usually can give more accurate eigenvalues compared with VMC, which relies on the quality of the proposed ansatz, while the particle representation of DMC often falls short of providing other information of the eigenfunction, such as its derivatives, unlike VMC.

As discussed above, one key to solving high-dimensional eigenvalue problems is the choice of function approximation to the targeted eigenfunction, ranging from the grid-based basis, spectral basis, to nonlinear parametrizations used in VMC, and to particle representations in DMC. Given the recent compelling success of neural networks in representing high-dimensional functions with remarkable accuracy and efficiency in various computational disciplines, it is fairly attempting to introduce neural networks to solve high-dimensional eigenvalue problems. This idea has been recently investigated under the variational formulation by [7–12]. Particularly [9–12] has shown the exciting potential of solving the many-electron Schrödinger equation with neural networks within the framework of VMC. On the other hand, how to apply neural networks in the formalism of DMC has not been explored in the literature, which leaves a natural open direction to investigate.

In this paper, we propose a new algorithm to solve high-dimensional eigenvalue problem for the second-order differential operators, in a similar spirit of DMC while based on the neural network parametrization of the eigenfunction. The eigenvalue problem is reformulated as a parabolic equation, whose solution can be represented by (nonlinear) Feynman-Kac formula in terms of forward-backward stochastic differential equations. Then we leverage the recently proposed deep BSDE method [13,14] to seek optimal eigenpairs. Specifically, two deep neural networks are constructed to represent the eigenfunction and its scaled gradient. Then the neural network is propagated according to the semigroup generated by the operator. The loss function is defined as the difference between the neural networks before and after the propagation. Compared to conventional DMC, the proposed algorithm provides a direct approximation to the target eigenfunction, which overcomes the shortcoming in providing the gradient information. Moreover, since the BSDE formulation is valid for nonlinear operators, our approach can be extended to high-dimensional nonlinear eigenvalue problems, as validated in our numerical examples.

The rest of this paper is organized as follows. In Section 2, we introduce the algorithm to solve the eigenvalue problem. In Section 3, numerical examples are presented. We conclude in Section 4 with an outlook for future work.

2. Numerical methods

2.1. The method for a linear operator

We consider the eigenvalue problem

$$\mathcal{L}\psi = \lambda\psi, \tag{1}$$

on $\Omega = [0, 2\pi]^d$ with periodic boundary condition where \mathcal{L} is a linear operator of the form

$$\mathcal{L}\psi(x) = -\frac{1}{2}\text{Tr}\left(\sigma\sigma^\top \text{Hess}(\psi)(x)\right) - b(x) \cdot \nabla\psi(x) + f(x)\psi(x). \tag{2}$$

σ is a $d \times d$ constant invertible matrix such that $\sigma\sigma^\top$ is positive definite, $\nabla\psi$ denotes the gradient of ψ , $b(x)$ is a d -dimensional vector field and $\text{Hess}(\psi)$ denotes the Hessian matrix of ψ .

To solve this eigenvalue problem, we augment a time variable and consider the following backward parabolic partial differential equation (PDE) in the time interval $[0, T]$:

$$\begin{cases} \partial_t u(t, x) - \mathcal{L}u(t, x) + \lambda u(t, x) = 0 & \text{in } [0, T] \times \Omega, \\ u(T, x) = \Psi(x) & \text{on } \Omega. \end{cases} \tag{3}$$

This is essentially a continuous time analog of the power iteration for matrix eigenvalue problem. Let us denote the solution of (3) as $u(T-t, \cdot) = \mathcal{P}_t^\lambda \Psi$ (note that the backward propagator $\{\mathcal{P}_t^\lambda\}_{t \leq T}$ forms a semigroup, i.e., $\mathcal{P}_{t_1}^\lambda \circ \mathcal{P}_{t_2}^\lambda = \mathcal{P}_{t_1+t_2}^\lambda$). According to the spectral theory of the elliptic operator, if Ψ is a stationary solution of (3), i.e., $\mathcal{P}_T^\lambda \Psi = \Psi$, then (λ, Ψ) must be an eigenpair of \mathcal{L} . Therefore, we can minimize the “loss function” $\|\mathcal{P}_T^\lambda \Psi - \Psi\|^2$ with respect to (λ, Ψ) to solve the eigenvalue problem. While this is a non-convex optimization problem, we expect local convergence to a valid eigenpair with appropriate initialization. Note that, unlike power iteration in which λ is determined by the eigenvector, in our algorithm λ is treated as a variational parameter and optimized jointly with the eigenfunction, as in the DMC method.

The reformulation above turns the eigenvalue problem into solving a parabolic PDE in high dimensions. For the latter, we can leverage the recently developed deep BSDE method [13–15] (which is why the parabolic PDE (3) is written backward in time). Let X_t solve the stochastic differential equation (SDE)

$$dX_t = \sigma dW_t, \tag{4}$$

or in the integral form

$$X_t = X_0 + \int_0^t \sigma dW_s, \tag{5}$$

where W_t is a d -dimensional Brownian motion, and X_0 is sampled from some initial distribution ν . Then according to the Itô's formula, the solution to (3), $u(t, x)$ satisfies

$$u(t, X_t) = u(0, X_0) + \int_0^t (f(X_s)u(s, X_s) - \lambda u(s, X_s) - b(X_s)\nabla u(s, X_s)) ds + \int_0^t \sigma^\top \nabla u(s, X_s) dW_s. \quad (6)$$

Note that simulating the two SDEs (5) and (6) is relatively simple even in high dimensions, while directly solving the PDE (3) is intractable. We remark that it is possible to add a drift term $b(X_t) dt$ to the SDE (4) and modify (6) accordingly (see the discussion below).

Of course, a priori in (6) for both $u(s, \cdot)$ and $\nabla u(s, \cdot)$ are unknown, while we know that if we set $u(s, \cdot) = \Psi(\cdot)$, the eigenfunction we look for, and $\nabla u(s, \cdot) = \nabla \Psi(\cdot)$, the solution $u(t, \cdot)$ remains $\Psi(\cdot)$ for all t . The idea of our method is then to use two neural networks, \mathfrak{N}_Ψ and $\mathfrak{N}_{\sigma^\top \nabla \Psi}$ as ansatz for the eigenfunction Ψ and its scaled gradient $\sigma^\top \nabla \Psi$, respectively. Assigning $u(0, X_0) = \mathfrak{N}_\Psi(X_0)$ and $\sigma^\top \nabla u(s, X_s) = \mathfrak{N}_{\sigma^\top \nabla \Psi}(X_s)$ in (6), the discrepancy for the propagated solution, i.e.,

$$\mathbb{E}_{X_0 \sim \nu} \left[\eta_1 |\mathfrak{N}_\Psi(X_T) - u(T, X_T)|^2 + \eta_2 |\mathfrak{N}_{\sigma^\top \nabla \Psi}(X_T) - \sigma^\top \nabla \mathfrak{N}_\Psi(X_T)|^2 \right] \quad (7)$$

then indicates the accuracy of the approximation. Here, we use $|\cdot|$ to denote the absolute value of a number or the Euclidean norm of a vector according to the context. Note that the second term above penalizes the discrepancy between the approximation of Ψ and its gradient, where η_1, η_2 are two weight hyperparameters. Therefore, using the above discrepancy as a loss function to optimize the triple $(\lambda, \mathfrak{N}_\Psi, \mathfrak{N}_{\sigma^\top \nabla \Psi})$ gives us a scheme to solve the eigenvalue problem. The above procedure can be directly extended to semilinear case where f depends on Ψ and $\nabla \Psi$, as we will discuss in Section 2.3.

To employ the above framework in practice, we numerically discretize the SDEs (5) and (6) using Euler–Maruyama method with a given partition of interval $[0, T]$: $0 = t_0 < t_1 < \dots < t_N = T$:

$$\mathcal{X}_0 = X_0, \quad \mathcal{X}_{t_{n+1}} = \mathcal{X}_{t_n} + \sigma \Delta W_n \quad (8)$$

and

$$\mathcal{U}_0 = \mathfrak{N}_\Psi(\mathcal{X}_0), \quad \mathcal{U}_{t_{n+1}} = \mathcal{U}_{t_n} + \left(f(\mathcal{X}_{t_n})\mathcal{U}_{t_n} - \lambda \mathcal{U}_{t_n} - (b\sigma^{-1\top} \mathfrak{N}_{\sigma^\top \nabla \Psi})(\mathcal{X}_{t_n}) \right) \Delta t_n + \mathfrak{N}_{\sigma^\top \nabla \Psi}(\mathcal{X}_{t_n}) \Delta W_n, \quad (9)$$

for $n = 0, 1, \dots, N - 1$. Here $\Delta t_n = t_{n+1} - t_n$, $\Delta W_n = W_{t_{n+1}} - W_{t_n}$, and we use X_t and $\mathcal{X}_t/\mathcal{U}_t$ to represent the continuous and discretized stochastic process, respectively. The noise terms ΔW_n have the same realization in (8) and (9), as in the forward-backward SDEs (5) and (6).

The loss function (7) then corresponds to the discrete counterpart:

$$\mathbb{E}_{X_0 \sim \nu} \left[\eta_1 |\mathfrak{N}_\Psi(\mathcal{X}_T) - \mathcal{U}_T|^2 + \eta_2 |\mathfrak{N}_{\sigma^\top \nabla \Psi}(\mathcal{X}_T) - \sigma^\top \nabla \mathfrak{N}_\Psi(\mathcal{X}_T)|^2 \right], \quad (10)$$

where $\nabla \mathfrak{N}_\Psi$ is the gradient of neural network \mathfrak{N}_Ψ with respect to its input. In practice, the expectation in (10) is further approximated by Monte Carlo sampling, which is similar to the empirical loss often used in the supervised learning context. For a given batch size K , we sample K points $\{\mathcal{X}_0^k\}_{k=1}^K$ of the initial state from the distribution ν at each training step and estimate the gradient of the loss with respect to the trainable parameters using the empirical Monte Carlo average of (10):

$$\frac{1}{K} \sum_{k=1}^K \left[\eta_1 |\mathfrak{N}_\Psi(\mathcal{X}_T^k) - \mathcal{U}_T^k|^2 + \eta_2 |\mathfrak{N}_{\sigma^\top \nabla \Psi}(\mathcal{X}_T^k) - \sigma^\top \nabla \mathfrak{N}_\Psi(\mathcal{X}_T^k)|^2 \right]. \quad (11)$$

We remark that the definition of the dynamic (5) is not unique and implicitly affects the detailed computation of the loss function (10) and (11). Specifically, in (5), the diffusion term σ is determined by the operator (2) while the choice of initial distribution and the drift term has some flexibility. If the drift in (5) changes, one can change the associated drift in (6) according to Itô's formula and define the loss again as (10) and (11). In this work, we choose the form of (5) without the drift and ν being the uniform distribution on Ω to ensure that the whole region is reasonably sampled for the optimization of eigenfunction. Some importance sampling can also be used if some prior knowledge of the eigenfunction is available, which we will not go into further details in this work.

At a high level, our algorithm is in a similar vein as the power method for solving the eigenvalue problem in linear algebra. Both algorithms seek for a solution that is stationary under the propagation. However, one distinction is that our algorithm may also be used for general eigenvalues depending on the initialization of λ and Ψ . Using matrix notation, this is similar to using the objective function $\|(A - \lambda)\nu\|^2$ to find non-dominant eigenpair λ and ν for the matrix A . The actual performance of solving for non-dominant eigenvalue depends on the initialization and spectral gap between eigenvalues, as will be illustrated by numerical results in Section 3. On the other hand, we find in the numerical experiments that if λ is initialized small enough, it will always converge to the first eigenvalue.

In practice, we use fully-connected feed-forward neural networks for the approximation of Ψ and $\sigma^\top \nabla \Psi$, respectively. To ensure periodicity of the neural network outputs, the input vector $x = (x_1, \dots, x_d)$ is first mapped into a fixed trigonometric




Fig. 1. Illustration of the neural network with 3 layers and several nodes in each layer. For $x = (x_1, \dots, x_d)$, $\sin(kx) = (\sin(kx_1), \sin(kx_2), \dots, \sin(kx_d))$ is a d -dimensional vector, and $\cos(kx)$ is defined similarly, $k = 1, 2, \dots, M$. The output for \mathfrak{N}_Ψ is a scalar and the output for $\mathfrak{N}_{\sigma^T \nabla \Psi}$ is a d -dimensional vector.

basis $\{\sin(jx_i), \cos(jx_i)\}_{i=1, j=1}^{d, M}$ of order M . Then the vector consisting of all basis components is fed into fully-connected neural networks with some hidden layers, each with several nodes. See Fig. 1 for illustration of the involved network structure. We use ReLU as the activation function and optimize the parameters with the Adam optimizer [16].

2.2. Normalization

The above loss has one caveat though, as the trivial solution ($\mathfrak{N}_\Psi = 0, \mathfrak{N}_{\sigma^T \nabla \Psi} = 0$) is a global minimizer. Therefore, normalization is required to exclude such a trivial case. We seek for eigenfunctions Ψ such that $\int_\Omega \Psi^2 = |\Omega|$, i.e., $\frac{1}{|\Omega|} \|\Psi\|_{L^2}^2 = 1$.¹ To proceed, we define the normalization constant

$$Z_\Psi = \text{sign}\left(\int_\Omega \mathfrak{N}_\Psi(x) dx\right) \left(\frac{1}{|\Omega|} \int_\Omega \mathfrak{N}_\Psi(x)^2 dx\right)^{1/2}. \quad (12)$$

Thus dividing \mathfrak{N}_Ψ by Z_Ψ , we enforce that the parametrized function ensures the normalization condition. Note that the first factor on the right hand side of (12) is introduced to fix the global sign ambiguity of the eigenfunction.

In computation, given the parameters of \mathfrak{N}_Ψ , we do not have direct access to Z_Ψ . Instead, at the ℓ -th step of training, we use our batch of K data samples to approximate Z_Ψ via

$$\hat{Z}_\Psi^\ell = \text{sign}\left(\sum_{k=1}^K \mathfrak{N}_\Psi(\mathcal{X}_0^{k,\ell})\right) \left(\frac{1}{K} \sum_{k=1}^K \mathfrak{N}_\Psi(\mathcal{X}_0^{k,\ell})^2\right)^{1/2}, \quad (13)$$

where the superscripts in $\mathcal{X}_0^{k,\ell}$ serve as the index of batch (k) and index of the training step (ℓ). The above is a Monte Carlo estimation of (12) if \mathcal{X}_0 is sampled from the uniform distribution (which we assume in this work). Due to the normalization procedure, \hat{Z}_Ψ^ℓ will enter into the loss function, and thus the stochastic gradient based on the empirical average over the batch becomes biased (since $\mathbb{E}(A/B) \neq \mathbb{E}A/\mathbb{E}B$ in general). To reduce the bias and make the training more stable, we introduce an exponential moving average scheme to the normalization constant in order to reduce the dependence of the loss on the estimated normalization constant of the current batch. In our implementation, we use

$$Z_\Psi^\ell = \gamma_\ell Z_\Psi^{\ell-1} + (1 - \gamma_\ell) \hat{Z}_\Psi^\ell. \quad (14)$$

Here $\gamma_\ell \in (0, 1)$ is the moving average coefficient for decay. It is observed that small γ_ℓ at the beginning makes training efficient, and later on its value is increased such that the gradient is less biased.

Given the introduced normalization factor, the neural network approximation \mathcal{U}_0 in the updating scheme (9) is replaced by (we suppress the training step index in \mathcal{X})

$$\mathcal{U}_0^k = \frac{\mathfrak{N}_\Psi(\mathcal{X}_0^k)}{Z_\Psi^\ell}, \quad (15)$$

and we would hope to reduce the discrepancy between the solution of (9) at time T and the normalized neural network approximation $\mathfrak{N}_\Psi(\mathcal{X}_T^k)/Z_\Psi^\ell$ through training. The associated batch approximation of loss function used for the computation of stochastic gradient is as follows

¹ The reason we set $\frac{1}{|\Omega|} \|\Psi\|_{L^2}^2 = 1$ instead of $\|\Psi\|_{L^2}^2 = 1$ is because we want to consider the problem in high dimensions in the domain $\Omega = [0, 2\pi]^d$. Consider the trivial case when $\mathcal{L} = -\Delta$, whose smallest eigenvalue is $\lambda = 0$ and any constant function is a corresponding eigenfunction. If $\|\Psi\|_{L^2}^2 = 1$, the constant function becomes $\Psi = (\frac{1}{2\pi})^{d/2}$, which vanishes as $d \rightarrow \infty$; instead the normalization $\|\Psi\|_{L^2}^2 = |\Omega|$ keeps the pointwise-value of Ψ as order 1, which benefits the training process.

$$\frac{1}{K} \sum_{k=1}^K \left(\eta_1 \left| \frac{\mathfrak{N}_\Psi(\mathcal{X}_T^k)}{Z_\Psi^\ell} - \mathcal{U}_T^k \right|^2 + \eta_2 \left| \mathfrak{N}_{\sigma^\top \nabla \Psi}(\mathcal{X}_T^k) - \frac{\sigma^\top \nabla \mathfrak{N}_\Psi(\mathcal{X}_T^k)}{Z_\Psi^\ell} \right|^2 \right) + \eta_3 (Z_0 - Z_\Psi^\ell)^+. \quad (16)$$

In the last term above, Z_0 is a hyperparameter and η_3 is the associated weight. This term is introduced to prevent Z_Ψ^ℓ being too small; otherwise the normalization would become unstable. In each training step, we calculate the gradient of (16) with respect to all the parameters to be optimized, including the eigenvalue λ and parameters in the neural network ansatz \mathfrak{N}_Ψ and $\mathfrak{N}_{\sigma^\top \nabla \Psi}$. Note that in (16) we do not normalize $\mathfrak{N}_{\sigma^\top \nabla \Psi}$ since its scale has been determined implicitly. When K is reasonably large and if we neglect the discretization error of simulating the SDEs, the empirical sum in (16) can be interpreted as a Monte Carlo approximation to the loss (ignoring the sign ambiguity)

$$\mathbb{E}_{\mathcal{X}_0 \sim \nu} \left[\eta_1 \left| \frac{\mathfrak{N}_\Psi(X_T)}{\|\mathfrak{N}_\Psi\|_2 / |\Omega|^{\frac{1}{2}}} - u(T, X_T) \right|^2 + \eta_2 \left| \mathfrak{N}_{\sigma^\top \nabla \Psi}(X_T) - \frac{\sigma^\top \nabla \mathfrak{N}_\Psi(X_T)}{\|\mathfrak{N}_\Psi\|_2 / |\Omega|^{\frac{1}{2}}} \right|^2 \right], \quad (17)$$

where $u(T, X_T)$ is defined as (6) except that $u(0, X_0) = |\Omega|^{\frac{1}{2}} \mathfrak{N}_\Psi(X_0) / \|\mathfrak{N}_\Psi\|_2$. We remark that the normalization procedure introduced here shares a similar spirit with Batch Normalization [17], which is widely used in the training of neural networks.

We summarize our algorithm as pseudocode in Algorithm 1.

Algorithm 1 Neural network based eigensolver.

Input: operator \mathcal{L} , terminal time T , number of time intervals N , loss weights $\eta_1, \eta_2, \eta_3, Z_0$, neural network structures, number of iterations, learning rate, batch size K , moving average coefficient γ_ℓ in (14)

Output: eigenvalue λ , eigenfunction \mathfrak{N}_Ψ and rescaled gradient $\mathfrak{N}_{\sigma^\top \nabla \Psi}$
 initialization: eigenvalue λ , \mathfrak{N}_Ψ , $\mathfrak{N}_{\sigma^\top \nabla \Psi}$ and normalization factor Z_Ψ^0

for $\ell = 1$ **to** the number of iterations **do**
 sample K points of \mathcal{X}_0 and sample K Wiener processes W_t
 compute \mathcal{X}_{t_n} via (8)
 compute the normalization factor Z_Ψ^ℓ via (13) and (14)
 normalize and propagate via (15) and (9)
 compute the gradient of loss (16) with respect to the trainable parameters
 update the trainable parameters by the Adam method

end for

2.3. The method for a semilinear operator

Our algorithm can be generalized to solve eigenvalue problems for semilinear operator

$$\mathcal{L}\psi(x) = -\frac{1}{2} \text{Tr} \left(\sigma \sigma^\top \text{Hess}(\psi)(x) \right) - b(x) \cdot \nabla \psi(x) + f(x, \psi(x), \sigma^\top \nabla \psi(x)). \quad (18)$$

The method for semilinear problems is almost the same to previous sections, except for a few modifications. The SDE for X_t is the same as (5) while equation (6) that the solution of the PDE (3) satisfies becomes

$$u(t, X_t) = u(0, X_0) + \int_0^t \left(f(X_s, u(s, X_s), \sigma^\top \nabla u(s, X_s)) - b(X_s) \cdot \nabla u(s, X_s) - \lambda u(s, X_s) \right) ds + \int_0^t \sigma^\top \nabla u(s, X_s) dW_s. \quad (19)$$

The discretization of X_t , equation (8), remains unchanged while equation (9) needs modification according to (19):

$$\begin{aligned} \tilde{\mathcal{U}}_{t_{n+1}} &= \mathcal{U}_{t_n} + \left(f(\mathcal{X}_{t_n}, \mathcal{U}_{t_n}, \mathfrak{N}_{\sigma^\top \nabla \Psi}) - \lambda \mathcal{U}_{t_n} - (b \sigma^{-\top} \mathfrak{N}_{\sigma^\top \nabla \Psi})(\mathcal{X}_{t_n}) \right) \Delta t_n + \mathfrak{N}_{\sigma^\top \nabla \Psi}(\mathcal{X}_{t_n}) \Delta W_n, \\ \mathcal{U}_{t_{n+1}} &= \text{Clip}(\tilde{\mathcal{U}}_{t_{n+1}}, P, Q), \end{aligned} \quad (20)$$

where (for $P < Q$) Clip is a clipping function given by

$$\text{Clip}(u, P, Q) = \begin{cases} P, & \text{if } u < P; \\ u, & \text{if } P \leq u \leq Q; \\ Q, & \text{if } u > Q. \end{cases} \quad (21)$$

Here we introduce the clipping function ($-P = Q > 0$) to prevent numerical instability caused by the nonlinearity of f in (19) (for instance, the cubic term in the nonlinear Schrödinger equation (25) below), especially at the early stage of training.

It checks \mathcal{U}_{t_n} and replaces those whose absolute values are larger than Q with $\text{sign}(\mathcal{U}_{t_n})Q$, where $Q > 0$ is an upper bound of the absolute value of the true normalized eigenfunction. Given the modified forward dynamics (20), the loss function for the semilinear operators are defined the same as (16), and the training algorithm is the same too.

3. Numerical results

In this section, we report the performance of the proposed eigensolver in three examples: the Fokker-Planck equation, the linear Schrödinger equation, and the nonlinear Schrödinger equation. The domain Ω is always $[0, 2\pi]^d$ with periodic boundary condition. In each example we consider the dimension $d = 5$ and $d = 10$. We also solve the second eigenpair of the linear Schrödinger equation with $d = 10$ to illustrate that our algorithm is able to get non-dominant eigenpairs. All the hyperparameters are given in Appendix B. We examine the errors of the prescribed eigenvalue, the associated eigenfunction, and the gradient of the eigenfunction. The errors for eigenfunctions and gradients of eigenfunctions are computed in the L^2 and L^∞ sense, approximated through a set of validation points. Given a set of validation points $\{\mathcal{X}^k\}_{k=1}^K$, we use the quantity

$$\text{err}_\Psi =: \left[\frac{1}{K} \sum_{k=1}^K \left(\frac{\mathfrak{N}_\Psi(\mathcal{X}^k)}{Z_\Psi^\ell} - \frac{\Psi(\mathcal{X}^k)}{\left(\frac{1}{K} \sum_{m=1}^K \Psi(\mathcal{X}^m)^2\right)^{\frac{1}{2}}} \right)^2 \right]^{\frac{1}{2}} \quad (22)$$

and

$$\text{err}_\Psi^\infty =: \max_k \left| \frac{\mathfrak{N}_\Psi(\mathcal{X}^k)}{Z_\Psi^\ell} - \frac{\Psi(\mathcal{X}^k)}{\left(\frac{1}{K} \sum_{m=1}^K \Psi(\mathcal{X}^m)^2\right)^{\frac{1}{2}}} \right| \quad (23)$$

to measure the error for eigenfunction where Z_Ψ^ℓ is computed via equation (14), with a known reference eigenfunction Ψ . We use

$$\text{err}_{\sigma^\top \nabla \Psi} =: \left[\frac{1}{Kd} \sum_{k=1}^K \left| \frac{\mathfrak{N}_{\sigma^\top \nabla \Psi}(\mathcal{X}^k)}{\left(\frac{1}{Kd} \sum_{m=1}^K |\mathfrak{N}_{\sigma^\top \nabla \Psi}(\mathcal{X}^m)|^2\right)^{\frac{1}{2}}} - \frac{\sigma^\top \nabla \Psi(\mathcal{X}^k)}{\left(\frac{1}{Kd} \sum_{m=1}^K |\sigma^\top \nabla \Psi(\mathcal{X}^m)|^2\right)^{\frac{1}{2}}} \right|^2 \right]^{\frac{1}{2}} \quad (24)$$

to quantify the error for the gradient approximation. We record and plot the error every 100 steps in the training process, with a smoothed moving average of window size 10. The final error reported is based on the average of the last 1000 steps. Besides the errors above, we also visualize and compare the density of the true eigenfunction and its neural network approximation (since it is hard to visualize the high-dimensional eigenfunction directly). The density of a function Ψ is defined as the probability density function of $\Psi(X)$ where X is a uniformly distributed random variable on Ω . In practice, the density is approximated by Monte Carlo sampling. As shown below, in all three examples, we find that the eigenpairs (with gradients) are solved accurately and the associated densities match well.

3.1. Fokker-Planck equation

In this subsection we consider the linear Fokker-Planck operator

$$\mathcal{L}\psi = -\Delta\psi - \nabla \cdot (\psi \nabla V) = -\Delta\psi - \nabla V \cdot \nabla\psi - \Delta V \psi,$$

where $V(x)$ is a potential function. The smallest eigenvalue of \mathcal{L} is $\lambda_1 = 0$ and the corresponding eigenfunction is $\Psi(x) = e^{-V(x)}$, which can be used to compute the error. We consider an example $V(x) = \sin(\sum_{i=1}^d c_i \cos(x_i))$, where x_i is the i -th coordinate of x , and c_i takes values in $[0.1, 1]$. The function V is periodic by construction. Fig. 2 shows the density and error curves for the Fokker-Planck equation in $d = 5$ and $d = 10$. For $d = 5$, the final errors of the eigenvalue, eigenfunction (in L^2 and L^∞) and the scaled gradient are $3.08\text{e-}3$, $2.91\text{e-}2$, $1.25\text{e-}1$ and $4.91\text{e-}2$. For $d = 10$, the final errors are $3.58\text{e-}3$, $1.62\text{e-}2$, $1.08\text{e-}1$ and $4.10\text{e-}2$.

3.2. Linear Schrödinger equation

In this subsection we consider the Schrödinger operator

$$\mathcal{L}\psi = -\Delta\psi + V\psi,$$

where $V(x)$ is a potential function. Here we choose $V(x) = \sum_{i=1}^d c_i \cos(x_i)$, in which c_i takes values in $[0, 0.2]$. With potential function being such a form, the problem is essentially decoupled. Therefore we are able to compute the eigenvalues and eigenfunctions in each dimension through the spectral method and obtain the final first eigenpair with high accuracy

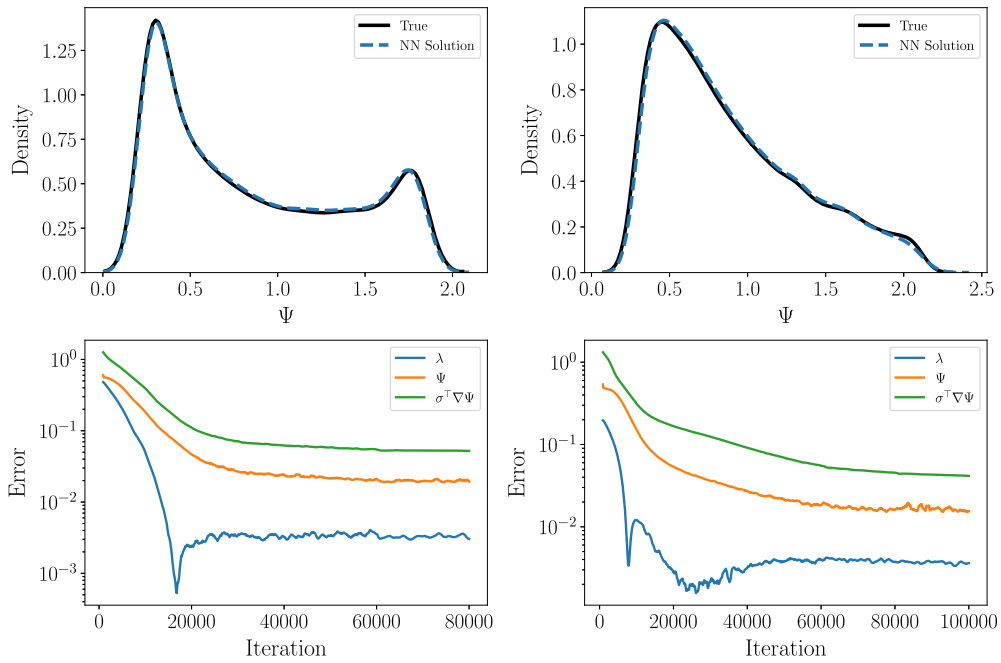


Fig. 2. Top: density of Ψ for Fokker-Planck equation with $d = 5$ (left) and $d = 10$ (right). Bottom: associated error curves in the training process with $d = 5$ (left) and $d = 10$ (right).




Fig. 3. Top: density of Ψ for linear Schrödinger equation with $d = 5$ (left) and $d = 10$ (right). Bottom: associated error curves in the training process with $d = 5$ (left) and $d = 10$ (right).

for comparison. The computation details are provided in Appendix A. Fig. 3 shows the density and error curves for the Schrödinger equation in $d = 5$ and $d = 10$. For $d = 5$, the final errors of the eigenvalue, eigenfunction (in L^2 and L^∞) and the scaled gradient are $8.84e-4$, $7.87e-3$, $4.78e-2$ and $2.30e-2$. For $d = 10$, the final errors are $1.40e-3$, $1.19e-2$, $7.48e-2$ and $3.14e-2$.




Fig. 4. Top: density of Ψ for nonlinear Schrödinger equation with $d = 5$ (left) and $d = 10$ (right). Bottom: associated error curves in the training process with $d = 5$ (left) and $d = 10$ (right).

3.3. Nonlinear Schrödinger equation

We finally consider a nonlinear Schrödinger operator with a cubic term, arising from the Gross-Pitaevskii equation [18,19] for the single-particle wavefunction in a Bose-Einstein condensate:

$$\mathcal{L}\psi = -\Delta\psi + \epsilon\psi^3 + V\psi. \tag{25}$$

Here we assume $\epsilon = 1$ and consider a specific external potential

$$V(x) = -\frac{1}{c^2} \exp\left(\frac{2}{d} \sum_{i=1}^d \cos x_i\right) + \sum_{i=1}^d \left(\frac{\sin^2 x_i}{d^2} - \frac{\cos x_i}{d}\right) - 3, \tag{26}$$

such that $\lambda = -3$, $\Psi(x) = \exp(\frac{1}{d} \sum_{j=1}^d \cos(x_j))/c$ is an eigenpair of the operator (25). Here c is a positive constant such that $\int_{\Omega} \Psi^2(x) dx = |\Omega|$. In this example we used $P = -5$ and $Q = 5$ in the clip function (21). Fig. 4 shows the density and error curves for the nonlinear Schrödinger equation in $d = 5$ and $d = 10$. For $d = 5$, the final errors of the eigenvalue, eigenfunction in L^2 and L^∞ and the scaled gradient are $1.53e-3$, $8.07e-3$, $6.58e-2$ and $4.36e-2$. For $d = 10$, the final errors are $2.6e-4$, $4.60e-3$, $3.83e-2$ and $3.55e-2$.

3.4. An example for the second eigenpair

In this subsection we reconsider the Schrödinger operator

$$\mathcal{L}\psi = -\Delta\psi + V\psi,$$

with the additional goal of finding the second eigenpairs. The potential function is double-well in each dimension: $V(x) = \sum_{i=1}^d A_i \cos(2x_i)$. The reference solutions are solved by the same numerical method as in Section 3.2.

We first consider the cases when the eigenvalues are well-separated. Assuming $d = 10$, $A_1 = 1.5$ and $A_i = 0.2$ for $2 \leq i \leq 10$, the associated eigen-gaps are $\lambda_2 - \lambda_1 = 4.52e-1$ and $\lambda_3 - \lambda_2 = 4.52e-1$. If we follow the training procedure described previously, we are able to find the first eigenpair with errors of $2.34e-3$, $3.95e-3$, $2.37e-2$ and $1.21e-2$. The left column in Fig. 5 shows the density and error curves.

On the other hand, our method is also able to find the second eigenpair given some mild prior estimate of the eigenvalue. Suppose that we have an approximation $\bar{\lambda}$ of the true second eigenvalue λ_2 . We can firstly fix this approximated eigenvalue and train the neural networks only with some steps. With such a pre-training procedure, we expect that the neural networks

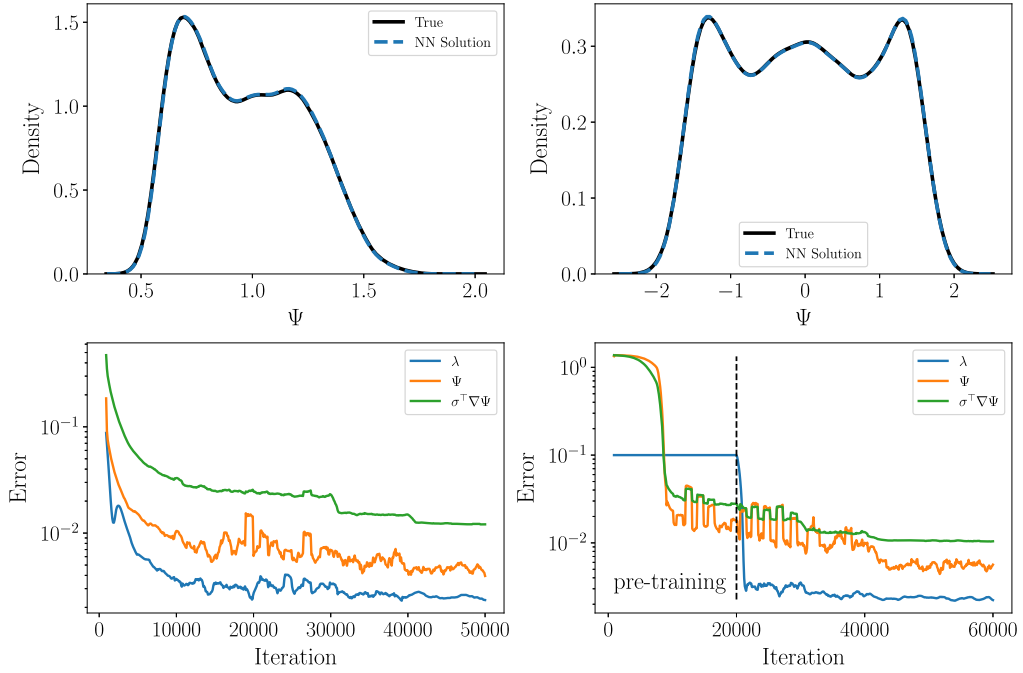


Fig. 5. Top: density of the first (left) and second (right) eigenfunction Ψ for linear Schrödinger equation ($d = 10$) with double-well potential. Bottom: associated error curves in the training process for the first (left) and second (right) eigenfunction. In the error curve on the right, the error for eigenvalue keeps unchanged in the first 20000 steps as we fixed it in this pre-training procedure.




Fig. 6. Plot of the eigenfunctions for one-dimensional degenerate linear Schrödinger problem with double-well potential. Left: first eigenfunction. Right: second eigenfunction.

would reach a reasonable approximation of the second eigenfunction and its scaled gradient. Then, we train both the eigenvalue and the neural networks for the eigenfunction until convergence. The right column in Fig. 5 shows the density and error curves of the second eigenpair, obtained by following the described procedure. In this example, $\bar{\lambda} = \lambda_2 + 0.1$, and the final errors are $2.23e-3$, $5.65e-3$, $5.08e-2$ and $1.04e-2$.

We remark that the eigenvalue problem becomes more challenging when it is nearly degenerate, i.e., the first and second eigenvalues are close to each other. This is a common phenomenon for various numerical methods, and our algorithm is no exception. Suppose we consider a one-dimensional problem with potential $V(x) = 5 \cos(2x)$. The first and second eigenvalues are -2.153 and -2.076 respectively, with a small gap $8.7e-2$. If we train the model for the first eigenpair like before, the obtained eigenfunction is plotted on the left of Fig. 6. If we solve the second eigenpair with the additional pre-training procedure introduced above, we still cannot get the second eigenfunction approximately, even with a fixed true second eigenvalue. The results can be improved if we furthermore have certain approximation to the second eigenfunction. For instance, if we initialize the neural networks for the eigenfunction as $\psi_{init} = \psi_2 + \epsilon \psi_1$ where ψ_1 and ψ_2 are the first and second eigenfunctions respectively and ϵ denotes the degree of perturbation. We take $\epsilon = 0.3$ in our numerical experiment, and we remark that the direction of perturbation is not necessarily in ψ_1 . Such initialization can be achieved through a supervised learning procedure. With this initialization, the final solution to the second eigenfunction is plotted on the right of Fig. 6.

According to the numerical results of both the well-separated and degenerate problems, we find that better initialization of eigenvalue or eigenfunctions can usually improve results. For the degenerate problem, even in the one-dimensional case, both the first and second eigenpairs are difficult to solve accurately since the final solution may involve a mix of two eigenfunctions.

4. Conclusion and future works

In this paper, we propose a new method to solve eigenvalue problems in high dimensions using neural networks. Our method is able to compute both eigenvalues and corresponding eigenfunctions (with gradients) with high accuracy.

There are several natural directions for future work. First, to apply our methodology to quantum many-body systems, we need to respect the permutation symmetry in our ansatz for the wavefunctions. Previous works [9,12,11,10] have proposed various flexible neural-network ansatz to incorporate in the symmetry, which can be combined with our approach. Moreover, in DMC, importance sampling techniques are often essential to improve the accuracy. In our context, this means to choose a better underlying diffusion process guided by a trial wavefunction depending on the problem. Last, the scalability of the method has to be tested on larger systems beyond the toy numerical examples in this work.

On the theoretical aspects, the understanding of the stability and convergence of the proposed method is a fascinating future direction. While the general analysis might be quite difficult given the highly nonlinear approximation induced by the neural networks and also the complicated optimization strategy, some perturbative analysis, especially in the linearized regime, might be possible. We will leave these to future works.

CRedit authorship contribution statement

Jiequn Han: Conceptualization, Methodology, Software, Writing - original draft, Writing - review & editing. **Jianfeng Lu:** Conceptualization, Funding acquisition, Methodology, Software, Writing - review & editing. **Mo Zhou:** Conceptualization, Methodology, Software, Writing - original draft, Writing - review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

The work of JL and MZ is supported in part by National Science Foundation via grant DMS-1454939 and DMS-2012286. The authors are grateful for computing time at the Terascale Infrastructure for Groundbreaking Research in Science and Engineering (TIGRESS) of Princeton University.

Appendix A. Spectrum method for linear Schrödinger equation

Suppose that the potential function in the linear Schrödinger operator $\mathcal{L} = -\Delta + V$ is decoupled with the form $V(x) = \sum_{j=1}^d c_j \cos(x_j)$, then we can solve the corresponding eigenvalue problem in a decoupled way. Specifically, assume we can solve the one-dimensional eigenvalue problem

$$-\Psi_j''(x) + c_j \cos(x) \Psi_j(x) = \lambda_j \Psi_j(x), \quad x \in [0, 2\pi]. \tag{A.1}$$

Then one can easily verify that $\lambda = \sum_{j=1}^d \lambda_j$ and

$$\Psi(x) = \prod_{j=1}^d \Psi_j(x_j) \tag{A.2}$$

together define an eigenpair of the original high-dimensional Schrödinger operator.

To solve (A.1), we can employ the classical spectrum method. For a fixed $N \in \mathbb{N}$, assume that

$$\Psi_j(x) = \sum_{m=-N}^N a_m^j e^{imx}, \tag{A.3}$$

then

$$\Psi_j'(x) = \sum_{m=-N}^N ma_m^j e^{imx}i. \tag{A.4}$$

Table B.1

Parameters for Fokker-Planck eigenvalue problems in Section 3.1.

Parameters	$d = 5$	$d = 10$
terminal time T	0.2	0.2
number of time intervals N	80	120
structure of neural networks	[80, 80, 80]	[300, 300, 300]
number of iterations	80000	100000
learning rate	[1e-4, 5e-5, 1e-5]	[5e-5, 2e-5, 1e-5]
moving average decay γ_ℓ	[0.2, 0.5, 0.9]	[0.2, 0.5, 0.9]
piecewise constant boundaries	[30000, 60000]	[60000, 80000]
batch size K	1024	1024

Table B.2

Parameters for linear Schrödinger eigenvalue problems in Section 3.2.

Parameters	$d = 5$	$d = 10$
terminal time T	0.3	0.3
number of time intervals N	80	120
structure of neural networks	[80, 80, 80]	[300, 300, 300]
number of iterations	80000	80000
learning rate	[1e-4, 5e-5, 1e-5]	[5e-5, 5e-5, 1e-5]
moving average decay γ_ℓ	[0.2, 0.5, 0.9]	[0.2, 0.5, 0.9]
piecewise constant boundaries	[30000, 60000]	[40000, 60000]
batch size K	1024	1024

Table B.3

Parameters for nonlinear Schrödinger eigenvalue problems in Section 3.3.

Parameters	$d = 5$	$d = 10$
terminal time T	0.2	0.3
number of time intervals N	120	200
structure of neural networks	[80, 80, 80]	[300, 300, 300]
number of iterations	60000	80000
learning rate	[5e-5, 2e-5, 1e-5]	[5e-5, 2e-5, 1e-5]
moving average decay γ_ℓ	[0.2, 0.9, 0.99]	[0.1, 0.9, 0.99]
piecewise constant boundaries	[20000, 40000]	[40000, 60000]
batch size K	2048	2048

Table B.4

Parameters for well-separated linear Schrödinger eigenvalue problems with $d = 10$ in Section 3.4.

state	first eigenpair	second eigenpair
terminal time T	0.2	0.2
number of time intervals N	320	320
structure of neural networks	[200, 200, 200]	[200, 200, 200]
number of iterations	50000	50000
learning rate	[5e-4, 1e-4, 1e-5]	[5e-4, 1e-4, 1e-5]
moving average decay γ_ℓ	[0.1, 0.2, 0.9]	[0.1, 0.9, 0.99]
piecewise constant boundaries	[30000, 40000]	[30000, 40000]
batch size K	2048	2048

Table B.5

Parameters for degenerate linear Schrödinger eigenvalue problems with $d = 1$ in Section 3.4.

state	first eigenpair	second eigenpair
terminal time T	0.2	0.2
number of time intervals N	80	80
structure of neural networks	[40, 40]	[40, 40]
number of iterations	6000	6000
learning rate	[5e-4, 1e-4, 1e-5]	[5e-4, 1e-4, 1e-5]
moving average decay γ_ℓ	[0.1, 0.2, 0.9]	[0.1, 0.2, 0.9]
piecewise constant boundaries	[2000, 4000]	[2000, 4000]
batch size K	512	512

derivatives, without a sacrifice of accuracy. The learning rates are chosen to be non-increasing piecewise constant according to common practice.

References

- [1] W.L. McMillan, Ground state of liquid He 4, *Phys. Rev. A* 138 (2) (1965) A442.
- [2] D. Ceperley, G.V. Chester, M. Kalos, Monte Carlo simulation of a many-fermion study, *Phys. Rev. B* 16 (7) (1977) 3081–3099.
- [3] R. Blankenbecler, D. Scalapino, R. Sugar, Monte Carlo calculations of coupled boson-fermion systems. I, *Phys. Rev. D* 24 (8) (1981) 2278.
- [4] S. Zhang, J. Carlson, J.E. Gubernatis, Constrained path Monte Carlo method for fermion ground states, *Phys. Rev. B* 55 (12) (1997) 7464.
- [5] W. Foulkes, L. Mitas, R.J. Needs, G. Rajagopal, Quantum Monte Carlo simulations of solids, *Rev. Mod. Phys.* 73 (1) (2001) 33.
- [6] R.J. Needs, M.D. Towler, N.D. Drummond, P.L. Ríos, Continuum variational and diffusion quantum Monte Carlo calculations, *J. Phys. Condens. Matter* 22 (2) (2009) 023201.
- [7] W. E, B. Yu, The deep Ritz method: a deep learning-based numerical algorithm for solving variational problems, *Commun. Math. Stat.* 6 (1) (2018) 1–12.
- [8] D. Pfau, S. Petersen, A. Agarwal, D.G.T. Barrett, K.L. Stachenfeld, Spectral inference networks: unifying deep and spectral learning, in: *International Conference on Learning Representations*, 2019.
- [9] J. Han, L. Zhang, W. E, Solving many-electron Schrödinger equation using deep neural networks, *J. Comput. Phys.* 399 (2019) 108929.
- [10] D. Pfau, J.S. Spencer, A.G. de, G. Matthews, W.M.C. Foulkes, Ab-initio solution of the many-electron Schrödinger equation with deep neural networks, preprint, arXiv:1909.02487, 2019.
- [11] J. Hermann, Z. Schätzle, F. Noé, Deep neural network solution of the electronic Schrödinger equation, preprint, arXiv:1909.08423, 2019.
- [12] K. Choo, A. Mezzacapo, G. Carleo, Fermionic neural-network states for ab-initio electronic structure, preprint, arXiv:1909.12852, 2019.
- [13] W. E, J. Han, A. Jentzen, Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations, *Commun. Math. Stat.* 5 (4) (2017) 349–380.
- [14] J. Han, A. Jentzen, W. E, Solving high-dimensional partial differential equations using deep learning, *Proc. Natl. Acad. Sci.* 115 (34) (2018) 8505–8510.
- [15] J. Han, J. Long, Convergence of the deep BSDE method for coupled FBSDEs, preprint, arXiv:1811.01165, 2018.
- [16] D. Kingma, J. Ba, Adam: a method for stochastic optimization, in: *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- [17] S. Ioffe, C. Szegedy, Batch normalization: accelerating deep network training by reducing internal covariate shift, in: *International Conference on Machine Learning*, 2015, pp. 448–456.
- [18] E.P. Gross, Structure of a quantized vortex in boson systems, *Nuovo Cimento* 20 (3) (1961) 454–477.
- [19] L.P. Pitaevskii, Vortex lines in an imperfect Bose gas, *Sov. Phys. JETP* 13 (2) (1961) 451–454.