Master in Quantum Science and Technology Barcelona

UNIVERSITAT DE BARCELONA

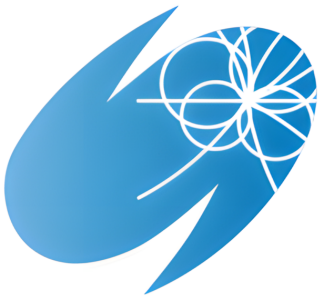

ICCUB

# Time-Dependent Neural Quantum States for Quantum Dynamics


Miguel Fernández Suárez

Supervisors

Arnau Rios Huguet

Javier Rozalén Sarmiento


7th of July 2026

# Time-Dependent Neural Quantum States for Quantum Dynamics

Miguel Fernández Suárez

Supervised by: Arnau Rios Huguet, Javier Rozalén Sarmiento

Institut de Ciències del Cosmos, 08028 Barcelona
29/06/2026

Neural Quantum States have emerged as powerful variational ansätze for modeling complex quantum systems, although their extension to real-time dynamics remains challenging. In this work, the wavefunctions of the three-dimensional Quantum Harmonic Oscillator and the deuteron are parametrized through two real-valued feed-forward neural networks, representing separately the amplitude and the phase. Expectation values are estimated using Markov chain Monte Carlo sampling, ground states are optimized through Stochastic Reconfiguration, and real-time evolution is performed using the McLachlan time-dependent variational principle. After accurately reproducing the ground states of the considered systems, the optimized Neural Quantum States are used as initial conditions for time evolution. Successful dynamics are obtained for the Quantum Harmonic Oscillator, a mass-rescaled deuteron and the physical deuteron in a periodic box, with fidelities against independent benchmarks remaining above 0.9986 in all cases considered. Ultimately, these results show that Neural Quantum States provide flexible variational ansätze capable of describing accurate real-time dynamics.

*Keywords*: Machine Learning, Neural Quantum States, Variational Monte Carlo, Time-Dependent Variational Principle, Dynamics, Deuteron.

## Acknowledgements

The author gratefully acknowledges the computer resources at Artemisa and the technical support provided by the Instituto de Física Corpuscular, IFIC(CSIC-UV). Artemisa is co-funded by the European Union through the 2014-2020 ERDF Operative Programme of Comunitat Valenciana, project IDIFEDER/2018/048.

I would like to thank Dr. Arnau Rios, Javier Rozalén and Dr. Alejandro Romero-Ros for their guidance, encouragement and help throughout this project, as well as for the many discussions that helped shape this work. I am also grateful to Dr. Alessandro Lovato and Dr. Noemi Rocco for their valuable suggestions, advice and technical insight. On a more personal note, I want to thank my parents for their constant support and trust in every decision I have made. Finally, I would like to acknowledge my roommates in Barcelona, who have become close friends and with whom I have shared, and hope to keep sharing, many great memories.

Miguel Fernández Suárez: mfernasu21@alumnes.ub.edu

# Contents

## Acronyms

**ANN** Artificial Neural Network. 1–3

**FFNN** Feed-Forward Neural Network. 1–3, 9, 20

**GD** Gradient Descent. 3, 5

**LO** Leading-Order. 14, 20

**NQS** Neural Quantum States. 1, 2, 4, 6, 10–20, 27

**QGT** Quantum Geometric Tensor. 5, 7, 28

**QHO** Quantum Harmonic Oscillator. 2, 8–14, 20, 26–29

**RK2** Second-Order Runge–Kutta. 29, 30

**RK4** Fourth-Order Runge–Kutta. 29, 30

**SGD** Stochastic Gradient Descent. 5

**SR** Stochastic Reconfiguration. 2, 3, 5–8, 10, 11, 15–17, 20

**TDVP** Time-Dependent Variational Principle. 1, 2, 7–9, 11–15, 17–20, 26–30

**VMC** Variational Monte Carlo. 1, 2, 4, 5, 20

## Introduction

Machine-learning techniques have become widespread both inside and outside the scientific domain. Their ability to identify patterns, compress information and solve high-dimensional optimization problems has made them useful in fields ranging from image recognition and natural-language processing to data analysis in experimental and theoretical physics [1, 2]. In recent years, these tools have also played an increasingly important role in quantum many-body physics, where the central object of interest, the wavefunction, is generally difficult to represent exactly [3].

A central challenge in quantum many-body problems is the exponential growth of the information required to describe a generic quantum state. Representations on a grid or in a complete many-body basis rapidly become impractical as the number of degrees of freedom increases, especially for interacting systems. A wide range of numerical methods has therefore been developed to make these problems tractable, including, among others, approaches that sample physically relevant configurations, as in Quantum Monte Carlo methods, and strategies that compress the wavefunction into more efficient representations. The latter class includes tensor-network ansätze, which have been particularly successful in one-dimensional systems [4–6], but whose real-time simulations are often limited by the growth of entanglement during the evolution [7, 8]. More generally, accurate ground state results alone do not guarantee that the same variational representation will remain reliable during time evolution, which is usually a more demanding task. These limitations motivate the search for alternative flexible variational representations of quantum many-body states.

Neural Quantum States (NQS) provide one such alternative. In this approach, the wavefunction is represented by an Artificial Neural Network (ANN) whose parameters are optimized according to the physical problem of interest, rather than by imposing a fixed functional form. This makes NQS flexible variational ansätze capable of encoding complex correlations between the degrees of freedom of a quantum system. Moreover, their natural combination with Variational Monte Carlo (VMC) allows one to avoid the explicit construction of the full wavefunction on an exponentially large grid or basis, thereby alleviating the *curse of dimensionality* inherent to many-body systems. Pioneering work by Carleo and Troyer [3] showed that NQS can achieve accurate results for both ground state properties and real-time dynamics in interacting spin systems. Since then, NQS have become a promising framework for the variational study of quantum many-body systems, with applications ranging from spin models to nuclear systems [9–11].

Most applications of NQS have focused on ground state optimization, where the variational energy provides a natural cost function. In nuclear physics, this strategy has been used to study the deuteron with Feed-Forward Neural Networks (FFNNs), showing that even relatively simple architectures can provide accurate variational representations of the bound-state wavefunction [12, 13]. NQS have also been combined with VMC to describe light nuclei beyond the deuteron, further illustrating their potential as flexible representations of many-body wavefunctions [14, 15]. Extending this framework to real-time dynamics is, however, challenging. Time evolution requires an especially accurate description of the full complex wavefunction, including its phase, and introduces the additional difficulty that numerical errors may accumulate over time. Recent works have addressed this problem using Time-Dependent Variational Principles (TDVPs) to project the exact Schrödinger evolution onto the variational manifold, exploring both the capabilities and the limitations of NQS dynamics in lattice and continuum systems [16–18].

Building on these developments, this thesis explores the use of NQS for ground state optimization and real-time dynamics in continuum quantum systems. It is structured as follows. Sec. 1 lays out the theoretical framework used throughout this work. After

introducing the NQS parametrization in terms of two FFNNs, representing separately the logarithm of the amplitude and the phase, this section presents the VMC sampling procedure, the SR method for ground state optimization and the McLachlan TDVP for real-time evolution. This setup is first validated on the three-dimensional Quantum Harmonic Oscillator (QHO) in Sec. 2, for which both the ground state and the dynamics are known analytically. The QHO provides not only a benchmark for the SR optimization and the TDVP evolution, but also a controlled setting to analyze how different numerical choices, such as the time integration scheme and the linear solver, affect the results. The method is then applied to the deuteron in Sec. 3, where the same NQS framework is tested in a more challenging weakly bound nuclear system with a spatially extended wavefunction.

The scripts developed for this thesis are available in a public GitHub repository [19]. The implementation relies on PyTorch for the NQS architectures and automatic differentiation, NumPy for certain numerical operations and Matplotlib for data visualization.

# 1 Theoretical Framework

## 1.1 Neural Quantum States as Variational Ansätze

ANNs are parametrized models designed to learn complex functional dependencies from data or from an optimization problem. In its simplest and most common architecture, the FFNN, information propagates in a single direction from an input layer, through one or more hidden layers, to an output layer [1].

Each layer consists of multiple neurons, each of which computes a weighted sum of its inputs followed by the application of a non-linear activation function. The output of a specific hidden layer $(l)$ is the vector

$$\mathbf{h}^{(l)}(\mathbf{z}) = \sigma(W^{(l)}\mathbf{z} + \mathbf{b}^{(l)}), \tag{1.1}$$

where $\sigma(\mathbf{x})$ is the activation function, applied element-wise, $\mathbf{z} \in \mathbb{R}^{n^{(l-1)}}$ is the input vector to the layer, and the parameters $W^{(l)} \in \mathbb{R}^{n^{(l)} \times n^{(l-1)}}$ and $\mathbf{b}^{(l)} \in \mathbb{R}^{n^{(l)}}$ are the weight matrix and bias vector, respectively. The non-linear activation function is crucial, as it enables the network to yield something other than a simple linear transformation.

One of the key motivations for the use of ANNs in physics is the Universal Approximation Theorem [20], which states that an FFNN with at least one hidden layer and a sufficient number of neurons can approximate any continuous function on a compact domain to an arbitrary level of accuracy $\varepsilon$:

$$\sup_{\mathbf{x} \in K} |f(\mathbf{x}; \boldsymbol{\theta}) - g(\mathbf{x})| < \varepsilon, \tag{1.2}$$

where $g(\mathbf{x})$ is the target function, $K$ is a compact set and $f(\mathbf{x}; \boldsymbol{\theta})$ is the neural network parametrized by $\boldsymbol{\theta} = \{W^{(l)}, \mathbf{b}^{(l)}\}_{l=1}^{L}$ [21].

When an ANN is used to parametrize a wavefunction, $\Psi(\mathbf{x}) \to \Psi_{\boldsymbol{\theta}}(\mathbf{x})$, the resulting ansatz is known as an NQS [3]. In this work, the NQS is parametrized as

$$\log \Psi_{\boldsymbol{\theta}}(\mathbf{x}) = A_{\boldsymbol{\theta}}(\mathbf{x}) + i\Phi_{\boldsymbol{\theta}}(\mathbf{x}), \tag{1.3}$$

where $A_{\boldsymbol{\theta}}(\mathbf{x})$ and $\Phi_{\boldsymbol{\theta}}(\mathbf{x})$ are real-valued neural networks representing the logarithm of the amplitude and the phase of the wavefunction, respectively[1]. This parametrization makes explicit that the sampling probability distribution depends only on the amplitude network, while the phase network remains relevant for certain observables and for real-time evolution.

[1] Working with $\log \Psi_{\boldsymbol{\theta}}$ is numerically convenient because, for instance, ratios of wavefunctions become differences of logarithms.

Throughout all the simulations presented in this work, $A_{\boldsymbol{\theta}}$ and $\Phi_{\boldsymbol{\theta}}$ are independent FFNNs with the same architecture. Each network takes the spatial coordinates $\mathbf{x} \in \mathbb{R}^3$ as input and returns a single real number. The architecture, shown in Fig. 1, consists of 2 hidden layers with 20 neurons per hidden layer, sigmoid activation functions[2] between consecutive layers, and a final linear output layer without bias[3], giving 520 trainable parameters per network. These choices are further discussed in Sec. 2.2.

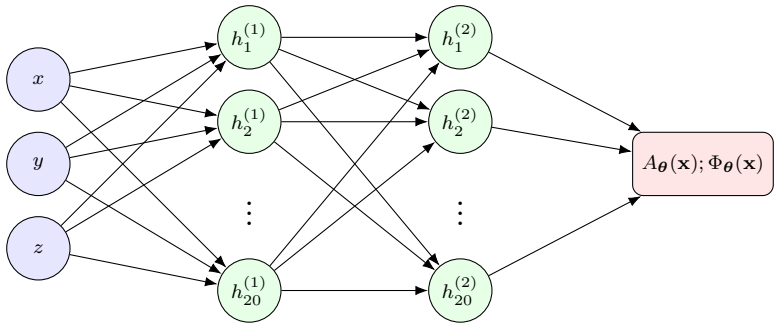


Figure 1: FFNN architecture used for both the amplitude network $A_{\boldsymbol{\theta}}(\mathbf{x})$ and the phase network $\Phi_{\boldsymbol{\theta}}(\mathbf{x})$. Each network takes the three spatial coordinates $(x, y, z)$ as input, has 2 hidden layers with 20 neurons each and sigmoid activation, and returns a single real-valued output.

The flexibility of ANNs makes them particularly appealing as variational ansätze, as they can represent a broad class of states while remaining computationally tractable. In the context of ground state optimization, the network parameters are treated as variational degrees of freedom and optimized according to the variational principle,

$$E_{\boldsymbol{\theta}} = \frac{\langle \Psi_{\boldsymbol{\theta}} | \hat{H} | \Psi_{\boldsymbol{\theta}} \rangle}{\langle \Psi_{\boldsymbol{\theta}} | \Psi_{\boldsymbol{\theta}} \rangle} \geq E_0, \tag{1.4}$$

where $E_0$ is the exact ground state energy. Consequently, minimizing $E_{\boldsymbol{\theta}}$ with respect to the network parameters $\boldsymbol{\theta}$ yields the best approximation to the ground state within the variational family [21]. Moreover, since the optimization does not rely on a pre-existing dataset of exact quantum states, but rather on minimizing the energy, this procedure can be formally understood as a type of reinforcement learning.

While the Universal Approximation Theorem guarantees the existence of parameters that approximate the target function arbitrarily well, it does not provide a practical procedure to find them. The optimization problem can be formulated in terms of a loss function $\mathcal{L}(\boldsymbol{\theta})$, which in the present context is the variational energy itself (Eq. (1.4)). The most common approach to optimization is Gradient Descent (GD), which iteratively updates the parameters in the direction of the steepest descent,

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}_t), \tag{1.5}$$

where $\eta$ is the learning rate, controlling the step size of each update [11]. While GD provides a baseline optimization strategy, it computes the direction of steepest descent with respect to the Euclidean metric in parameter space, effectively treating this space as flat. More sophisticated optimization algorithms, such as Adam [22], are frequently employed to improve convergence by adapting the learning rate based on gradient moments [11], but they still assume a Euclidean metric. In our case, we instead navigate the complex landscape of a parametrized quantum wavefunction by making use of SR, an optimization method that takes into account the geometry of the variational manifold.

[2]The sigmoid function is given by $\sigma(x) = \frac{1}{1+e^{-x}}$.

[3]Final biases would only multiply the wavefunction by a normalization factor and a global phase.

## 1.2 Variational Monte Carlo and Energy Gradient

Evaluating expectation values generally requires computing integrals over the configuration space of the system. A direct grid-based numerical evaluation would require discretizing each coordinate into a finite mesh. For a system of $A$ particles in three spatial dimensions, if $n$ grid points are used per coordinate, the total number of grid points scales as $n^{3A}$. This exponential growth with the number of degrees of freedom makes direct grid-based integration impractical for generic many-body systems, causing this approach to quickly fall victim to the *curse of dimensionality*.

In order to avoid this issue, VMC replaces deterministic integration with statistical estimates. Given a probability distribution $p_{\boldsymbol{\theta}}$, the expectation value of a function $g(\mathbf{x})$ can be estimated as

$$\int d\mathbf{x}\; p_{\boldsymbol{\theta}}(\mathbf{x}) g(\mathbf{x}) \approx \frac{1}{N} \sum_{i=1}^{N} g(\mathbf{x}_i), \tag{1.6}$$

where the configurations $\{\mathbf{x}_i\}_{i=1}^{N}$ are sampled from $p_{\boldsymbol{\theta}}(\mathbf{x})$. We denote expectation values over this sampled distribution by $\langle \cdot \rangle$. When working with wavefunctions, the natural choice is the Born probability distribution,

$$p_{\boldsymbol{\theta}}(\mathbf{x}) = \frac{|\Psi_{\boldsymbol{\theta}}(\mathbf{x})|^2}{\langle \Psi_{\boldsymbol{\theta}} | \Psi_{\boldsymbol{\theta}} \rangle}, \tag{1.7}$$

which ensures that the configurations contributing most strongly to the state are sampled more frequently. In practice, we generate samples from this distribution using random walks through the Metropolis-Hastings algorithm [23–26], summarized in Appendix A.

Within the NQS framework, energy is estimated by sampling configurations from $p_{\boldsymbol{\theta}}(\mathbf{x})$ and averaging an energy estimator evaluated at each configuration. The variational energy can be rewritten as

$$E_{\boldsymbol{\theta}} = \frac{\langle \Psi_{\boldsymbol{\theta}} | \hat{H} | \Psi_{\boldsymbol{\theta}} \rangle}{\langle \Psi_{\boldsymbol{\theta}} | \Psi_{\boldsymbol{\theta}} \rangle} = \int d\mathbf{x}\; p_{\boldsymbol{\theta}}(\mathbf{x}) \frac{\hat{H}\Psi_{\boldsymbol{\theta}}(\mathbf{x})}{\Psi_{\boldsymbol{\theta}}(\mathbf{x})} = \int d\mathbf{x}\; p_{\boldsymbol{\theta}}(\mathbf{x}) E_L(\mathbf{x}) \approx \frac{1}{N} \sum_{i=1}^{N} E_L(\mathbf{x}_i), \tag{1.8}$$

where $E_L(\mathbf{x}) = \frac{\hat{H}\Psi_{\boldsymbol{\theta}}(\mathbf{x})}{\Psi_{\boldsymbol{\theta}}(\mathbf{x})}$ is defined as the local energy. Note that, if $\Psi_{\boldsymbol{\theta}}$ is an exact eigenstate of the Hamiltonian, then $E_L(\mathbf{x})$ becomes independent of $\mathbf{x}$ and equal to the corresponding eigenvalue. The variance of $E_L(\mathbf{x})$ is therefore zero in this limit, independently of the sample size $N$, which is known as the *zero-variance property*. Consequently, fluctuations of the local energy provide a direct measure of the quality of the variational ansatz.

Since the variational energy plays the role of the loss function, $\mathcal{L}(\boldsymbol{\theta}) = E_{\boldsymbol{\theta}}$, optimizing the network typically requires evaluating its gradient with respect to the network parameters. This derivative is slightly subtle in VMC because the energy is an average over a probability distribution that depends on the wavefunction itself. Thus, when the parameters change, both the local energy, $E_L(\mathbf{x})$, and the sampling distribution, $p_{\boldsymbol{\theta}}(\mathbf{x})$, change. The latter contribution would be missed if one simply sampled a fixed set of Metropolis walkers, computed the average local energy and applied automatic differentiation to that average [21]. Taking this dependence into account, the energy gradient for a complex NQS with real-valued variational parameters can be written as

$$\partial_{\theta_k} E_{\boldsymbol{\theta}} = 2\,\mathrm{Re}\left[\langle (E_L(\mathbf{x}) - E_{\boldsymbol{\theta}})\, \partial_{\theta_k} \log \Psi^*_{\boldsymbol{\theta}}(\mathbf{x}) \rangle\right], \tag{1.9}$$

as derived in Appendix B. Importantly, evaluating this gradient only requires the parameter derivatives of the logarithm of the wavefunction, and does not require differentiating the local energy.

If a standard optimizer such as GD or Adam were used, the above expression would provide the gradient needed to update the network parameters. It is important to note that, in practice, when combining VMC with machine learning, the expectation value in Eq. (1.9) is evaluated using a finite set of Monte Carlo samples. The resulting gradient is therefore a stochastic estimate rather than the exact gradient of the variational energy. If inserted directly into the GD update of Eq. (1.5), this leads to the method known as Stochastic Gradient Descent (SGD).

However, SGD still corresponds to a Euclidean gradient descent in parameter space. As we shall see next, in the present work the parameters are updated using SR, which replaces the Euclidean gradient update with a geometry-aware update involving the metric of the variational manifold.

## 1.3 Stochastic Reconfiguration for Ground State Optimization

The energy gradient of Eq. (1.9) gives the direction of the steepest descent in ordinary Euclidean parameter space. However, this notion of steepest descent depends on measuring distances directly between parameter vectors. Since equal Euclidean displacements in parameter space may correspond to very different changes in the physical quantum state, this is not the most natural choice for a variational wavefunction.

Introduced by Sorella in the context of Green's Function Monte Carlo, and later reformulated for VMC, SR is an optimization method that updates the variational parameters according to the geometry of the variational manifold rather than following the ordinary Euclidean gradient [27, 28]. It is closely related to Natural Gradient Descent, introduced by Amari, in which the direction of steepest descent is defined using the Fisher Information, a metric tensor that measures distances between probability distributions [29]. SR applies the same geometric idea to variational quantum states, where the corresponding metric is given by the real part of the Quantum Geometric Tensor (QGT).

We now provide one way of obtaining the QGT. Since wavefunctions that differ only by a global phase or by a normalization factor represent the same state, a physically meaningful distance should be insensitive to these redundancies. The Fubini–Study distance provides such a measure [30, 31]. For two nearby variational states, it is given by

$$ds^2 = 1 - |\langle \Psi_{\boldsymbol{\theta}} | \Psi_{\boldsymbol{\theta}+d\boldsymbol{\theta}} \rangle|^2 , \tag{1.10}$$

where the states $|\Psi_{\boldsymbol{\theta}}\rangle$ are assumed to be normalized. Expanding the above expression to second order in $d\boldsymbol{\theta}$ yields

$$ds^2 = \sum_{ij} d\theta_i d\theta_j \,\mathrm{Re}\left[\langle \partial_{\theta_i}\Psi | \partial_{\theta_j}\Psi\rangle - \langle \partial_{\theta_i}\Psi|\Psi\rangle\langle\Psi|\partial_{\theta_j}\Psi\rangle\right] = \sum_{ij} d\theta_i d\theta_j \,\mathrm{Re}\, Q_{ij}, \tag{1.11}$$

where $Q_{ij}$ is the $(i,j)$ element of the QGT. The real part of $Q_{ij}$ defines the Fubini–Study metric, which measures distances between nearby physical states, while the imaginary part is associated with the Berry curvature [32]. Using the logarithmic derivative identity $\partial_{\theta_k}\Psi_{\boldsymbol{\theta}}(\mathbf{x}) = \Psi_{\boldsymbol{\theta}}(\mathbf{x})\partial_{\theta_k}\log\Psi_{\boldsymbol{\theta}}(\mathbf{x})$ and defining $O_i(\mathbf{x}) = \partial_{\theta_i}\log\Psi_{\boldsymbol{\theta}}(\mathbf{x})$, the QGT can be rewritten as the covariance of the logarithmic derivatives,

$$Q_{ij} = \langle O_i^*(\mathbf{x}) O_j(\mathbf{x})\rangle - \langle O_i^*(\mathbf{x})\rangle \langle O_j(\mathbf{x})\rangle . \tag{1.12}$$

The SR matrix is the real part of the QGT, which can be rewritten in centered form as

$$S_{ij} = \mathrm{Re}\, Q_{ij} = \mathrm{Re}\left[\langle \Delta O_i^*(\mathbf{x}) \Delta O_j(\mathbf{x})\rangle\right], \tag{1.13}$$

where $\Delta O_i(\mathbf{x}) = O_i(\mathbf{x}) - \langle O_i\rangle$. Suppose now that the parameters are updated by a small displacement $\delta\boldsymbol{\theta}$. To first order, the energy changes as

$$E_{\boldsymbol{\theta}+\delta\boldsymbol{\theta}} \approx E_{\boldsymbol{\theta}} + \sum_i \partial_{\theta_i} E_{\boldsymbol{\theta}} \delta\theta_i. \tag{1.14}$$

The SR direction is obtained by comparing infinitesimal updates with the same physical size, imposed through the constraint of fixed $ds^2 = \sum_{ij} \delta\theta_i \delta\theta_j S_{ij}$ [26]. This constrained minimization problem can be written in terms of the Lagrange function $\mathcal{M}$,

$$\mathcal{M} = \sum_i \partial_{\theta_i} E_{\boldsymbol{\theta}} \delta\theta_i + \lambda \left( \sum_{ij} \delta\theta_i \delta\theta_j S_{ij} - ds^2 \right), \tag{1.15}$$

where $\lambda$ is a Lagrange multiplier. This procedure leads to the SR update equation

$$\sum_j S_{kj}\, \delta\theta_j = (\nabla_{\boldsymbol{\theta}} E_{\boldsymbol{\theta}})_k\,, \tag{1.16}$$

where the energy gradient $(\nabla_{\boldsymbol{\theta}} E_{\boldsymbol{\theta}})_k = 2\mathrm{Re}\left[\langle \Delta E_L(\mathbf{x}) \Delta O_k^*(\mathbf{x}) \rangle\right]$ is the same as Eq. (1.9), written in centered form with $\Delta E_L(\mathbf{x}) = E_L(\mathbf{x}) - E_{\boldsymbol{\theta}}$. This way, the parameter update is $\boldsymbol{\theta} \longrightarrow \boldsymbol{\theta} - \eta\, \delta\boldsymbol{\theta}$ and SR can be interpreted as a preconditioned gradient method through

$$\delta\boldsymbol{\theta} = S^{-1} \nabla_{\boldsymbol{\theta}} E_{\boldsymbol{\theta}}. \tag{1.17}$$

In practice, however, the inverse of the matrix $S$ is not computed explicitly. Instead, the linear system of Eq. (1.16) is solved numerically through a linear solver algorithm. In our case, we make use of the Cholesky decomposition[4] [33].

The SR method is particularly useful for ground state optimization, as it can be interpreted as an imaginary time evolution projected onto the variational manifold [34]. The exact imaginary time evolution satisfies

$$\frac{d}{d\tau} |\Psi(\tau)\rangle = -\frac{1}{\hbar} \left( \hat{H} - E \right) |\Psi(\tau)\rangle, \tag{1.18}$$

which suppresses excited state components and projects a generic initial state towards the ground state. However, the exact imaginary time evolved state does not generally remain within the chosen variational family. SR addresses this by projecting the imaginary time evolution onto the tangent space of the variational manifold, providing the closest approximation to the exact imaginary time evolution.

## 1.4 Time-Dependent Variational Principle

Having discussed ground state optimization, we now turn to real-time dynamics, governed by the time-dependent Schrödinger equation

$$i\hbar \frac{d}{dt} |\Psi(t)\rangle = \hat{H} |\Psi(t)\rangle. \tag{1.19}$$

Once again, for a general initial state, the exact evolved state will not remain inside the variational family defined by the NQS. Therefore, the goal is to find the trajectory of parameters $\boldsymbol{\theta}(t)$ such that the parametrized state $|\Psi_{\boldsymbol{\theta}}(t)\rangle$[5] best approximates the exact solution of Eq. (1.19) within the variational manifold.

[4] The matrix $S$ is decomposed as $LL^T$, where $L$ is a lower triangular matrix. Cholesky decomposition requires $S$ to be positive definite. In numerical applications, $S$ is likely to become ill-conditioned, so we add a small diagonal regularization $S \to S + \varepsilon_{\mathrm{SR}} \mathbb{1}$ with $\varepsilon_{\mathrm{SR}} > 0$.

[5] Although we write $|\Psi_{\boldsymbol{\theta}}(t)\rangle$, the time dependence of the variational state is understood to enter exclusively through the parameters $\boldsymbol{\theta}(t)$, i.e. $|\Psi_{\boldsymbol{\theta}}(t)\rangle = |\Psi_{\boldsymbol{\theta}(t)}\rangle$.

For real variational parameters, the equations of motion can be obtained from the McLachlan variational principle [35]. Since the exact solution of Eq. (1.19) cannot be assumed known, this principle minimizes the local error in state norm between the Schrödinger time derivative and the time derivative allowed by the variational manifold,

$$\left\| \frac{d}{dt}|\Psi_{\boldsymbol{\theta}}\rangle - (-\frac{i}{\hbar}\hat{H})|\Psi_{\boldsymbol{\theta}}\rangle \right\|. \tag{1.20}$$

The tangent vectors of the variational manifold, corresponding to the directions in which the wavefunction changes under infinitesimal variations of the parameters, are

$$|v_k\rangle = \frac{\partial}{\partial\theta_k}|\Psi_{\boldsymbol{\theta}}\rangle. \tag{1.21}$$

Since the parameters of the networks are real, the time derivative of the parametrized wavefunction is a real linear combination of these tangent vectors,

$$|\dot{\Psi}_{\boldsymbol{\theta}}\rangle = \sum_j \dot{\theta}_j |v_j\rangle, \qquad \dot{\theta}_j \in \mathbb{R}. \tag{1.22}$$

Defining

$$|\varphi\rangle = |\dot{\Psi}_{\boldsymbol{\theta}}\rangle + \frac{i}{\hbar}\hat{H}|\Psi_{\boldsymbol{\theta}}\rangle, \tag{1.23}$$

McLachlan's principle corresponds to minimizing $\langle\varphi|\varphi\rangle$ with respect to the time derivatives of the parameters $\dot{\theta}_k$. Since $\frac{\partial|\varphi\rangle}{\partial\dot{\theta}_k} = |v_k\rangle$, this principle leads to

$$\hbar\sum_j \dot{\theta}_j \mathrm{Re}\left[\langle v_k|v_j\rangle\right] = \mathrm{Im}\left[\langle v_k|\hat{H}|\Psi_{\boldsymbol{\theta}}\rangle\right], \tag{1.24}$$

where we have used that the parameters are real. In order to express the above equation in a form analogous to Eq. (1.16), in which the $S$ matrix appears, the component of each tangent vector parallel to the wavefunction is projected out:

$$\widetilde{v}_k(\mathbf{x}) = \langle\mathbf{x}|v_k\rangle - \langle\mathbf{x}|\Psi_{\boldsymbol{\theta}}\rangle\frac{\langle\Psi_{\boldsymbol{\theta}}|v_k\rangle}{\langle\Psi_{\boldsymbol{\theta}}|\Psi_{\boldsymbol{\theta}}\rangle} = \Delta O_k(\mathbf{x})\Psi_{\boldsymbol{\theta}}(\mathbf{x}), \tag{1.25}$$

where $v_k(\mathbf{x}) = O_k(\mathbf{x})\Psi_{\boldsymbol{\theta}}(\mathbf{x})$ has been used. Employing this projected tangent basis, Eq. (1.24) can be rewritten in terms of covariances and the logarithmic derivatives of the wavefunction already introduced in Eq. (1.12):

$$\sum_j S_{kj}\dot{\theta}_j = \frac{1}{\hbar}\mathrm{Im}\left[\langle\Delta O_k^*(\mathbf{x})\Delta E_L(\mathbf{x})\rangle\right]. \tag{1.26}$$

This matrix $S$ is the same real part of the QGT that appears in the SR update of Eq. (1.16). We use Cholesky decomposition again to solve the above equation numerically[6]. The difference between SR and TDVP time evolution lies in the right-hand side of the corresponding linear system. The real-time TDVP involves the imaginary part of the local energy covariance, while SR uses the real part of the same covariance to follow the energy gradient to its minimum.

For complex holomorphic ansätze, the same equations of motion can be equivalently derived from the McLachlan, Lagrangian and Dirac–Frenkel variational principles [34, 36].

[6] TDVP time evolution is more sensitive to numerical instabilities than ground state SR. For this reason, we now use an adaptive diagonal regularization for the $S$ matrix, $S \to S + \varepsilon_{\mathrm{TDVP}}\,\mathrm{diag}(1 + |\mathbf{F}|)$, with $F_k = \frac{1}{\hbar}\mathrm{Im}\left[\langle\Delta O_k^*(\mathbf{x})\Delta E_L(\mathbf{x})\rangle\right]$.

In particular, the Dirac–Frenkel principle is especially powerful as it directly projects the Schrödinger equation onto the tangent space of the variational manifold. However, for our ansatz, this principle leads to a complex set of equations for the real-valued parameters, with the real and imaginary parts corresponding to the McLachlan and Lagrangian formulations, respectively. For this reason, we have used McLachlan's principle to derive the real system of equations used for the time evolution.

## 2 Quantum Harmonic Oscillator

The QHO serves as a benchmark to validate the methodology introduced in Sec. 1. Since both the ground state and the dynamics of the QHO can be obtained analytically, it provides a controlled setting in which the SR ground state optimization and the TDVP time evolution can be tested before considering more complex systems.

### 2.1 System and Exact Solution

In harmonic oscillator units, defined by $\hbar = m = 1$, the Hamiltonian is given by

$$\hat{H}(\mathbf{x}_0) = -\frac{1}{2}\nabla^2 + \frac{1}{2}\sum_{k=x,y,z} \omega_k^2 \left(x_k - x_{0,k}\right)^2, \tag{2.1}$$

where the dimensionless oscillation frequency in the $k$ direction is $\omega_k$ and the position of the potential center is $\mathbf{x}_0 = (x_{0,x}, x_{0,y}, x_{0,z})$. This Hamiltonian is separable, and therefore its eigenstates factorize into products of one-dimensional QHO eigenstates. The energy eigenvalues are

$$E_{n_x n_y n_z} = \omega_x\left(n_x + \frac{1}{2}\right) + \omega_y\left(n_y + \frac{1}{2}\right) + \omega_z\left(n_z + \frac{1}{2}\right), \tag{2.2}$$

and the ground state energy corresponds to the $n_x = n_y = n_z = 0$ case. Moreover, the exact ground state wavefunction is Gaussian and given by

$$\Psi_0(\mathbf{x}) = \left(\frac{\omega_x\omega_y\omega_z}{\pi^3}\right)^{1/4} e^{-\frac{1}{2}\sum_k \omega_k\left(x_k - x_{0,k}\right)^2}. \tag{2.3}$$

While the above expression serves as a reference for the optimization problem, coherent states provide a useful analytical benchmark for dynamics [37]. These are the eigenstates of the lowering operator, and can be triggered by displacing the potential origin and/or applying a momentum kick

$$\Psi(\mathbf{x}, t=0) = \Psi_0(\mathbf{x}) e^{i\sum_k p_{0,k}(x_k - x_{0,k})}, \tag{2.4}$$

where $\mathbf{p}_0$ is the applied momentum. When the potential origin is shifted, the initial state is displaced by $-\boldsymbol{\Delta} = -(\mathbf{x}_0' - \mathbf{x}_0)$ from the point of view of the Hamiltonian, where $\mathbf{x}_0' = (x_{0,x}', x_{0,y}', x_{0,z}')$ is the new potential center. The dynamics are then given by

$$\begin{cases} \langle x_k(t)\rangle = x_{0,k}' - \Delta_k \cos(\omega_k t) + \frac{p_{0,k}}{\omega_k}\sin(\omega_k t) \\ \langle p_k(t)\rangle = p_{0,k}\cos(\omega_k t) + \omega_k \Delta_k \sin(\omega_k t) \end{cases} \qquad k = x, y, z, \tag{2.5}$$

while the wavefunction remains Gaussian and evolves as

$$\Psi(\mathbf{x}, t) = \left(\frac{\omega_x\omega_y\omega_z}{\pi^3}\right)^{1/4} e^{-\frac{1}{2}\sum_k \omega_k\left(x_k - \langle x_k(t)\rangle\right)^2 + i\Phi(\mathbf{x},t)}, \tag{2.6}$$

where

$$\Phi(\mathbf{x},t) = \sum_{k=x,y,z} \langle p_k(t)\rangle \left[x_k - \frac{\langle x_k(t)\rangle + x'_{0,k}}{2}\right] + \frac{1}{2}\sum_{k=x,y,z} p_{0,k}\Delta_k - \frac{1}{2}\sum_{k=x,y,z}\omega_k t. \tag{2.7}$$

Moreover, after the dynamics are triggered, the new time-independent value of the energy is given by

$$E = E_0 + \frac{1}{2}\sum_{k=x,y,z}\left[p_{0,k}^2 + \omega_k^2\Delta_k^2\right]. \tag{2.8}$$

## 2.2 Numerical Setup and NQS Ansatz

All the simulations presented in this work are run using $N = 50,000$ walkers. A study on how the number of walkers affects the results is included in Appendix C. As stated in Sec. 1.1, both the amplitude and the phase networks are fixed to FFNNs with 2 hidden layers and 20 neurons per hidden layer.

The use of 2 hidden layers is motivated by the additional expressivity gained by composing nonlinear sigmoid functions. In a network with a single hidden layer, each neuron receives a linear combination of the input coordinates and applies the activation function, giving $h_j(\mathbf{x}) = \sigma(\mathbf{w}_j\mathbf{x} + b_j)$. The output is then built as a linear combination of these for all neurons in the layer. For the 1D QHO, this architecture yields good results, as the target wavefunction is a Gaussian that only depends on a single coordinate. However, in the 3D QHO, the wavefunction depends simultaneously on the 3D spatial coordinates through combinations such as $\omega_x^2x^2 + \omega_y^2y^2 + \omega_z^2z^2$, and a single hidden layer is not expressive enough to represent the corresponding 3D Gaussian. With 2 hidden layers, the neurons in the second layer receive as input the nonlinear features constructed by the first hidden layer, allowing the network to represent more complex spatial dependencies with a moderate number of neurons.

The number of neurons per hidden layer is fixed to 20 as a compromise between expressivity and computational cost. Although simpler cases, such as the isotropic QHO, can be accurately described with fewer neurons, more demanding cases, such as anisotropic oscillators and nuclear systems, require a more flexible ansatz.

In addition to the already detailed architecture, an extra element is added to the variational ansatz. For the QHO simulations, the wavefunction is multiplied by a Gaussian envelope according to

$$\Psi_{\boldsymbol{\theta}}(\mathbf{x}) \leftarrow e^{-\alpha_{\text{env}}|\mathbf{x}-\mathbf{x}_{\text{env}}|^2}\Psi_{\boldsymbol{\theta}}(\mathbf{x}). \tag{2.9}$$

The exponent factor is fixed to $\alpha_{\text{env}} = 0.04$, while the envelope variable $\mathbf{x}_{\text{env}}$ is initialized at the center of the harmonic potential and treated as a trainable parameter. This envelope imposes a large-distance decay on the wavefunction, allowing simple initial walker distributions, such as a uniform one. Without it, during the early stages of the optimization, some walkers may diffuse far from the main support of the wavefunction, producing unstable energy estimates and requiring a very large number of Metropolis steps to return to the physically relevant region. This way, the Gaussian envelope helps keep the walker ensemble localized, improving the stability of the sampling. Moreover, treating $\mathbf{x}_{\text{env}}$ as a trainable parameter is especially useful for the time-evolution simulations, as the dynamics may displace the probability distribution from the initial envelope position. Appendix D includes a detailed comparison between fixing $\mathbf{x}_{\text{env}}$ and optimizing it as a trainable parameter.

Regarding the time evolution, the TDVP linear system is solved at each time step using the Cholesky decomposition of the regularized $S$ matrix, and the resulting equations are integrated using Runge–Kutta 2 [38]. A more detailed comparison between time integration schemes and linear solver choices is found in Appendix E.

## 2.3 Ground State Optimization

Two different configurations are considered. The first one corresponds to the standard isotropic oscillator,

$$\omega_x = \omega_y = \omega_z = 1, \qquad \mathbf{x}_0 = (0, 0, 0),$$

for which the exact ground state energy is $E_0 = 3/2$. The second one is an anisotropic oscillator with

$$\omega_x = 1/3, \qquad \omega_y = 3, \qquad \omega_z = 1, \qquad \mathbf{x}_0 = (0, 1, -1),$$

for which the exact ground state energy is $E_0 = 13/6$. Although moving the potential center does not necessarily increase the intrinsic difficulty of the problem for NQS, as it only corresponds to a translation, the second configuration provides a more demanding test because the wavefunction is no longer rotationally invariant.

For both configurations, the NQS was optimized using an SR regularization parameter $\varepsilon_{\text{SR}} = 8 \cdot 10^{-5}$ and a fixed learning rate $\eta = 10^{-2}$. During the optimization, the total variational energy is monitored together with the kinetic and potential contributions. For the ground state of a QHO, the virial theorem $\langle K \rangle = \langle U \rangle = E_0/2$ is expected to be satisfied [37], providing an additional check beyond the convergence of the total energy.

Figures 2 and 3 show the result of the ground state optimization for the isotropic and anisotropic configurations, respectively. In both cases, the optimized NQS is compared with the exact ground state wavefunction through one-dimensional cuts along the Cartesian directions, while the convergence of the variational energy and its kinetic and potential contributions is monitored during training.

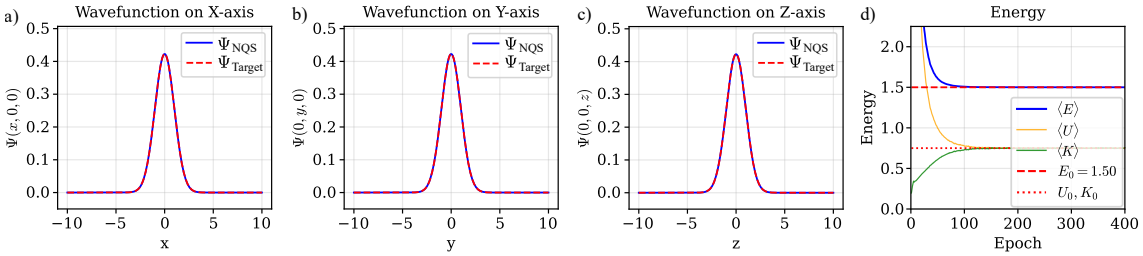


Figure 2: Ground state optimization of the 3D isotropic QHO with $\omega_x = \omega_y = \omega_z = 1$ using SR. The exact energy is $E_0 = 3/2$. Panels a)–c) compare the one-dimensional cuts of the NQS along the Cartesian axes with the exact ground state wavefunction. The solid blue curves correspond to the NQS, while the dashed red curves correspond to the exact wavefunction of Eq. (2.3). Panel d) shows the convergence of the total energy as a function of epochs, together with the kinetic and potential contributions. The blue, green and yellow curves correspond to the variational total, kinetic and potential energies, respectively. The dashed and dotted red lines indicate the exact energies.

For the isotropic oscillator, Fig. 2 a)–c) shows that the NQS accurately reproduces the Gaussian profile along the three Cartesian axes. The variational energy of panel d) converges to the target value, reaching $E_{\boldsymbol{\theta}} = 1.5001 \pm 0.0001$[7] within the first 400 epochs. The kinetic and potential energies also approach their reference value, consistent with the virial theorem prediction. This shows that the optimization does not only reproduce the correct total energy, but also learns the correct balance between the kinetic and potential contributions.

In the anisotropic case, the exact ground state is displaced from the origin and, due to the different oscillator frequencies, the Gaussian decay depends on the spatial direction.

[7]The uncertainties quoted for Monte Carlo estimates correspond to the $1\sigma$ statistical error of the mean, $\sigma/\sqrt{N}$, where $\sigma^2$ is the variance of the sampled observable.

Nevertheless, the NQS remains in excellent agreement with the exact wavefunction cuts in Fig. 3 a)–c), and panel d) shows that the energy converges to $E_{\boldsymbol{\theta}} = 2.1668 \pm 0.0001$ within the first 400 epochs. The virial theorem is also fulfilled in this case, although the kinetic energy converges to its expected value much earlier during the optimization. This behaviour is specific to the initialization and parameters adopted in this simulation.

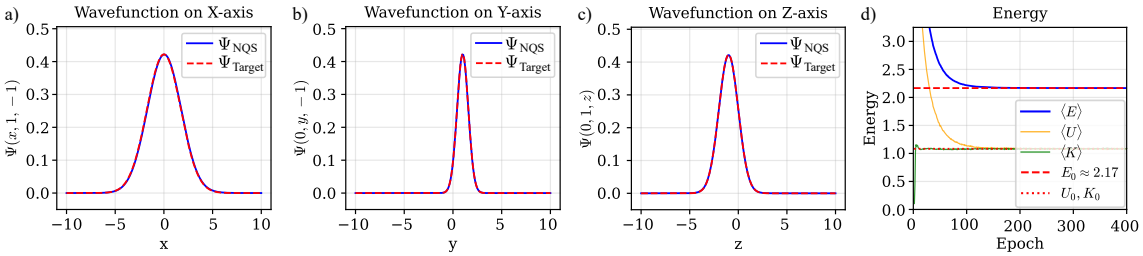


Figure 3: Same as Fig. 2, but for the anisotropic oscillator with $\omega_x = 1/3$, $\omega_y = 3$, $\omega_z = 1$, $\mathbf{x}_0 = (0, 1, -1)$ and exact energy $E_0 = 13/6$.

Figure 4 displays the histograms for the final Metropolis walker distributions after ground state optimization and the exact one-dimensional marginal probability densities, $p_{x_k}(x_k)$, for the anisotropic QHO. The latter are obtained by integrating $|\Psi_0(\mathbf{x})|^2$ of Eq. (2.3) over the remaining two spatial coordinates for each Cartesian direction.

The agreement between the Metropolis walker histograms and the exact probability densities confirms that the sampling procedure correctly reproduces the distribution $p_{\boldsymbol{\theta}}$. Having validated that the SR optimization provides a reliable variational representation of the QHO ground state, the optimized NQS can now be reused as the starting point for the TDVP time evolution.

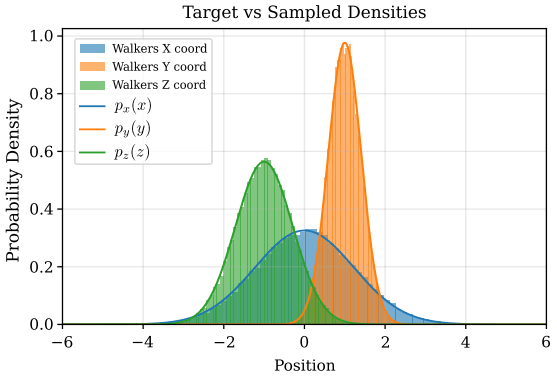


Figure 4: One-dimensional marginal probability densities of the anisotropic QHO. The shaded histograms show the walker distribution for the $x$, $y$ and $z$ coordinates in blue, orange and green, respectively. The solid curves of the same colors correspond to the exact marginal densities $p_{x_k}(x_k)$ with $k = x, y, z$.

## 2.4 Time Evolution

Time evolution is performed for the same two QHO configurations we considered in the previous subsection, using a regularization parameter $\varepsilon_{\mathrm{TDVP}} = 7 \cdot 10^{-6}$. On the one hand, the harmonic potential center of the isotropic oscillator is shifted from the origin to $\mathbf{x}_0' = (1, 1, 1)$, and a momentum kick $\mathbf{p}_0 = (1, 1, 1)$ is applied to the initial wavefunction, defining symmetric dynamics in which the wavefunction is expected to behave the same way along the Cartesian directions. On the other hand, the potential center of the anisotropic oscillator is shifted from $\mathbf{x}_0 = (0, 1, -1)$ to $\mathbf{x}_0' = (-1, 0, 1)$, and a momentum kick $\mathbf{p}_0 = (1, \frac{1}{2}, -1)$ is applied. Since the oscillator frequencies are different along the three Cartesian directions and the displacement and momentum kick are no longer symmetric, this second setup provides a more challenging dynamical test.

The quality of the TDVP time evolution is assessed using three complementary diagnostics. First, the expectation value of the position is compared with the exact analytical trajectory of Eq.(2.5). Second, the energy expectation value is monitored. Since the Hamiltonian is time-independent after the initial shift and kick, the exact energy should remain constant during the evolution. Finally, beyond comparing individual observables, the fi-

delity between the NQS and the exact coherent state of Eq.(2.6) serves as a global measure of the agreement with the TDVP dynamics.

Figure 5 shows the time evolution of the position expectation values for the two QHO configurations, together with the corresponding exact coherent state trajectories. Panel a) corresponds to the isotropic case, in which the three Cartesian components follow the same trajectory due to the symmetry of the displacement, kick and oscillator frequencies. Panel b) shows the anisotropic case, where the motion differs along each Cartesian direction, leading to distinct oscillation periods and amplitudes.

The agreement between the TDVP evolution and the exact trajectories confirms that the variational dynamics successfully reproduce the coherent state motion of the QHO. The only noticeable deviation is found in the $\langle x \rangle(t)$ component of the anisotropic case, where the absolute difference reaches at most $10^{-2}$.

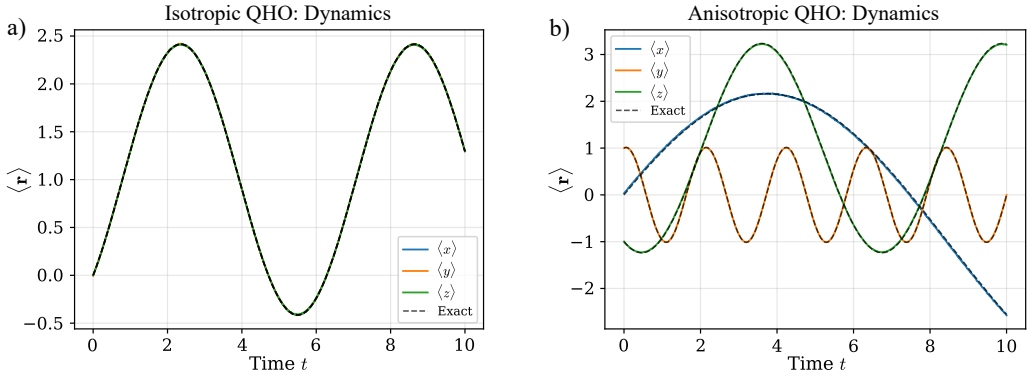


Figure 5: Time evolution of the position expectation values for the isotropic (panel a) and anisotropic (panel b) QHO configurations. Solid lines show the TDVP results, while dashed lines indicate the analytical coherent state trajectories of Eq. (2.5).

Figure 6 displays the energy expectation value during the TDVP time evolution. The Monte Carlo estimate is shown together with its statistical error, while the analytical statistical uncertainty, which can be derived using Eq. (2.6), is plotted around a linear fit to the numerical data.

In both configurations, energy remains approximately conserved over the full time interval, with fluctuations of the same order as the expected statistical uncertainty. The energy drift is slightly larger for the anisotropic case, where the linear fit gives a total variation of $-4.8 \cdot 10^{-2}$ over the simulated time interval, corresponding to a relative drift of 0.49%. This remains small compared with the energy scale of the system, indicating that the TDVP evolution is stable despite not being exactly energy conserving.

If an exact solution $\Psi_{\text{exact}}(\mathbf{x})$ is known, fidelity can also be written in a form suitable for estimation with the walker ensemble. The fidelity between the NQS and the exact wavefunction is defined as

$$\mathcal{F} = \frac{|\int d\mathbf{x} \Psi^*_{\boldsymbol{\theta}}(\mathbf{x}) \Psi_{\text{exact}}(\mathbf{x})|^2}{\int d\mathbf{x} |\Psi_{\boldsymbol{\theta}}(\mathbf{x})|^2 \cdot \int d\mathbf{x} |\Psi_{\text{exact}}(\mathbf{x})|^2}. \tag{2.10}$$

Using the fact that the walkers are distributed following $p_{\boldsymbol{\theta}}(\mathbf{x})$, as verified in Fig. 4, the fidelity can be rewritten in terms of expectation values over the NQS probability distribution:

$$\mathcal{F} = \frac{|\int d\mathbf{x}\, p_{\boldsymbol{\theta}}(\mathbf{x}) f(\mathbf{x})|^2}{\int d\mathbf{x}\, p_{\boldsymbol{\theta}}(\mathbf{x}) |f(\mathbf{x})|^2} \approx \frac{\left|\frac{1}{N}\sum_{i=1}^{N} f(\mathbf{x}_i)\right|^2}{\frac{1}{N}\sum_{i=1}^{N} |f(\mathbf{x}_i)|^2} = \frac{|\langle f \rangle|^2}{\langle |f|^2 \rangle}, \tag{2.11}$$

where $f(\mathbf{x}) = \frac{\Psi_{\text{exact}}(\mathbf{x})}{\Psi_{\boldsymbol{\theta}}(\mathbf{x})}$ and $N$ is the number of walkers.

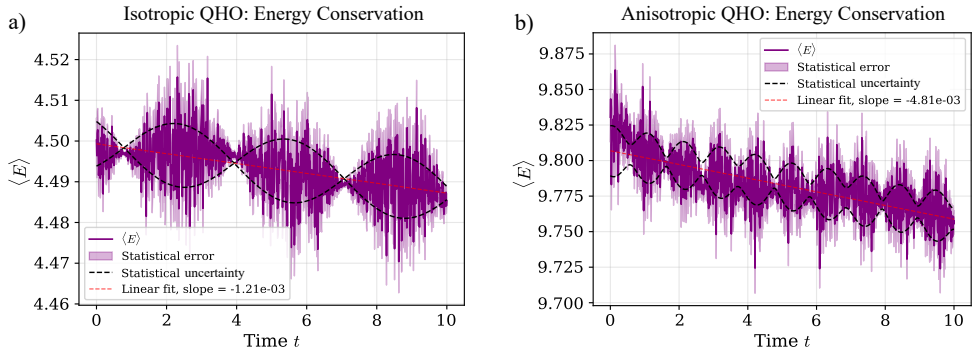


Figure 6: Energy expectation value during the TDVP time evolution for the isotropic (panel a) and anisotropic (panel b) QHO configurations. The shaded region indicates the Monte Carlo statistical error. The dashed curves show the analytical statistical uncertainty obtained from the exact energy variance, plotted around a linear fit to the numerical data.

Figure 7 shows the fidelity between the evolved NQS and the exact coherent state wavefunction. The walker-based estimate is validated by independently evaluating the wavefunction and computing the overlap on a spatial grid. In both configurations, the fidelity stays close to unity throughout the simulated time interval, exhibiting a gradual decrease as small errors accumulate during the time evolution. For the isotropic oscillator, panel a) shows that the fidelity stays above 0.99958, corresponding to a loss of $4.2 \cdot 10^{-4}$. For the anisotropic case, shown in panel b), the fidelity reaches 0.9986 at the end of the evolution, corresponding to a fidelity loss of order $10^{-3}$. As expected, the anisotropic case is more demanding and therefore leads to a lower fidelity. Nevertheless, the results obtained for both configurations demonstrate that the evolved NQS remains a reliable variational representation of the exact time-dependent wavefunction.

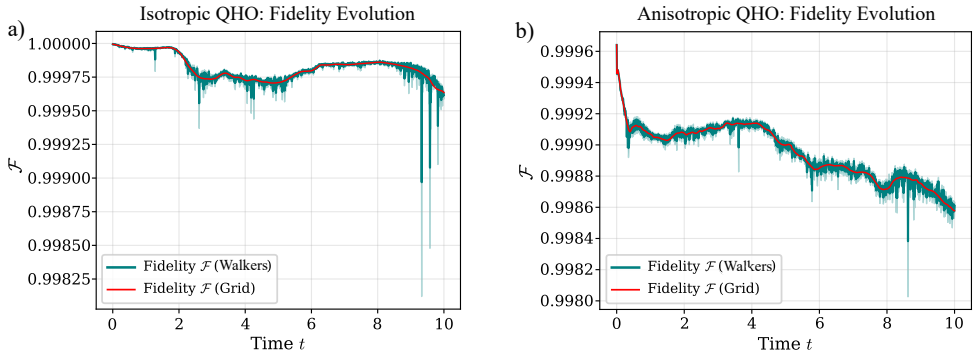


Figure 7: Fidelity between the TDVP-evolved NQS and the exact coherent state dynamics for the isotropic and anisotropic QHO configurations, shown in panels a) and b), respectively. The teal curve corresponds to the fidelity estimated from the walker ensemble, and the shaded teal region indicates its statistical error. The red curve shows the fidelity obtained by direct integration on a spatial grid. The two panels use different vertical scales.

## 3 Deuteron

After validating the NQS framework on the analytically solvable QHO, we now turn our attention to the deuteron, the simplest bound nuclear system. Formed by a single proton and a single neutron, the deuteron provides a setting in which the dynamics are governed by the nuclear interaction. Despite the apparent simplicity, it presents a notable challenge, since its weak binding and spatially extended wavefunction make it particularly sensitive to the long-distance behaviour of the variational ansatz.

## 3.1 Effective Field Theory Description of the Deuteron

Since the deuteron only contains two nucleons, the system can be formulated in terms of a single relative coordinate:

$$\mathbf{r} = \mathbf{r}_p - \mathbf{r}_n, \qquad r = |\mathbf{r}|, \tag{3.1}$$

where $\mathbf{r}_p$ and $\mathbf{r}_n$ are the proton and neutron coordinates, respectively. We choose a leading-order pionless effective field theory description of the proton-neutron interaction [14]. This model is suitable when the typical momenta involved are small compared to the pion mass scale, $m_\pi \simeq 140$ MeV/$c^2$, so that pion exchange is not resolved explicitly and the nuclear interaction is represented through contact terms. The singularities of the contact terms are regularized by introducing a Gaussian regulator, which leads to a Gaussian radial dependence of the potential [10].

The Leading-Order (LO) Hamiltonian in the center-of-mass frame takes the form

$$\hat{H}_{\mathrm{LO}} = -\frac{\hbar^2}{2\mu}\nabla^2_{\mathbf{r}} + V_{\mathrm{LO}}(r), \tag{3.2}$$

where $\mu = \frac{m_p m_n}{m_p + m_n}$ is the reduced mass. The LO interaction is given by

$$V_{\mathrm{LO}}(r) = C_{01} g_{01}(r) P_0^\sigma P_1^\tau + C_{10} g_{10}(r) P_1^\sigma P_0^\tau, \tag{3.3}$$

where $P_S^\sigma$ projects onto total spin $S$, $P_T^\tau$ projects onto total isospin $T$, $g_{ST}(r)$ denotes the regulated radial dependence of the contact interaction and $C_{ST}$ controls its strength. Although the deuteron, $|d\rangle$, has orbital angular momentum $L = 0, 2$, we assume that the $S$-wave component is dominant. It is experimentally known to belong to the spin-triplet ($S = 1$) channel. Together with its isospin-singlet ($T = 0$) character, this leads to $P_0^\sigma P_1^\tau |d\rangle = 0$ and $P_1^\sigma P_0^\tau |d\rangle = |d\rangle$. Introducing the Gaussian radial dependence of the potential, Eq. (3.3) simplifies to

$$V_{\mathrm{LO}}(r) = C_{10} e^{-\frac{r^2}{R_0^2}}, \tag{3.4}$$

where $R_0$ sets the potential range. The values of $C_{10} = -142$ MeV and $R_0 = 1$ fm are taken from Ref. [39].

In contrast with the QHO, the ground state and dynamics of this Hamiltonian are not known analytically. Consequently, in this case the NQS results are benchmarked against an independent numerical solution. In this reference calculation, the ground state is obtained by solving the Schrödinger equation using a finite-difference discretization, which leads to a tridiagonal Hamiltonian matrix. The resulting ground state solution is evolved using a split-operator method [40, 41]. In this approach, the potential part of the evolution is applied in coordinate space, while the kinetic part is applied in momentum space using Fourier transforms.

## 3.2 Adapting the NQS Ansatz to the Deuteron

We employ the same neural-network architecture, number of walkers, TDVP linear solver and time integrator as for the QHO. However, we change the envelope of the amplitude network. Since the system is now a nuclear bound state, an exponential decay at large distances is expected. The exponential $e^{-\alpha r}$ may seem like a natural choice. However, this is not suitable because $r = |\mathbf{r}|$ is not smooth at the origin and it would introduce singular contributions to the local kinetic energy. For this reason, the wavefunction is instead multiplied by an envelope according to

$$\Psi_{\boldsymbol{\theta}}(\mathbf{r}) \leftarrow e^{-\alpha_{\mathrm{env}} r_{\mathrm{aux}}} \Psi_{\boldsymbol{\theta}}(\mathbf{r}), \tag{3.5}$$

where $r_{\text{aux}} = \sqrt{r^2 + r_c^2} - r_c$ with $r_c = 0.5\,\text{fm}$ and $\alpha_{\text{env}}$ is a trainable parameter. Equivalently, this amounts to subtracting the exponent $\alpha_{\text{env}} r_{\text{aux}}$ from the raw output of the amplitude network. At large distances, this envelope produces the desired exponential decay, while close to the origin $r_{\text{aux}}$ behaves quadratically in $r$, avoiding divergences in the local kinetic energy. The length $r_c$ sets the transition scale between these two regimes.

The dynamics are initialized by inducing a small monopolar expansion of the optimized ground state. This is implemented by adding a quadratic radial term to the output of the phase network, so that the wavefunction is effectively transformed as

$$\Psi_{\boldsymbol{\theta}}(\mathbf{r}) \leftarrow e^{i\beta r^2} \Psi_{\boldsymbol{\theta}}(\mathbf{r}), \tag{3.6}$$

where $\beta$ controls the strength of the perturbation. Since the phase depends only on the radial distance, this kick does not displace the center of the wavefunction. Instead, it induces a breathing motion, in which the wavefunction periodically expands and contracts around the origin [42].

A major challenge in applying the NQS method to the physical deuteron arises from its weakly bound nature. As a consequence, the wavefunction is spatially extended and very sensitive to the long-distance behaviour of the variational ansatz. Direct NQS simulations of the TDVP time evolution of this weakly bound system are therefore particularly challenging. For this reason, two complementary setups are considered for the dynamics. Firstly, since a larger mass produces a more compact bound state, a mass-rescaled deuteron is studied in Sec. 3.4. Secondly, the physical-mass deuteron is studied in a finite box by enforcing periodic boundary conditions in the neural network. This reduces the difficulty associated with representing the extended tail of the weakly bound wavefunction, as explained in Sec. 3.5.

### 3.3 Ground State Optimization

For the ground state optimization of the physical-mass deuteron, the NQS is trained using an SR regularization parameter $\varepsilon_{\text{SR}} = 4 \cdot 10^{-5}$ and a fixed learning rate $\eta = 3 \cdot 10^{-3}$.

Figure 8 shows the SR ground state optimization of the physical-mass deuteron. The optimized NQS wavefunction is compared with the independent numerical solution through one-dimensional cuts along the three Cartesian axes in panels a)–c), and the convergence of the total variational energy, together with its kinetic and potential contributions, is monitored in panels d) and e).

The agreement of the wavefunction profiles along the three Cartesian axes demonstrates that the NQS successfully learns the spherical symmetry of the deuteron ground state without imposing this symmetry in the neural network architecture. The variational energies converge to

$$\begin{aligned} E_0 &= -2.171 \pm 0.006 \text{ MeV}, \\ K_0 &= 16.8 \pm 0.1 \text{ MeV}, \\ U_0 &= -18.9 \pm 0.1 \text{ MeV}, \end{aligned}$$

in agreement with the independent numerical result $E_{\text{num}} = -2.1755$ MeV. These values illustrate a characteristic feature of the deuteron. The kinetic and potential energies are large in magnitude, but they nearly cancel each other to yield a small binding energy [43].

To characterize the spatial extension of wavefunction, we show the convergence of two radial observables during the SR optimization and their numerical values in Fig. 9. The mean proton-neutron separation $\langle r \rangle$ (panel a) and the quantity $\sqrt{\langle r^2 \rangle}/2$ (panel b) are considered. Since the two nucleons have approximately equal masses, the center of mass lies approximately halfway between them, so the latter observable corresponds to the point-nucleon radius.

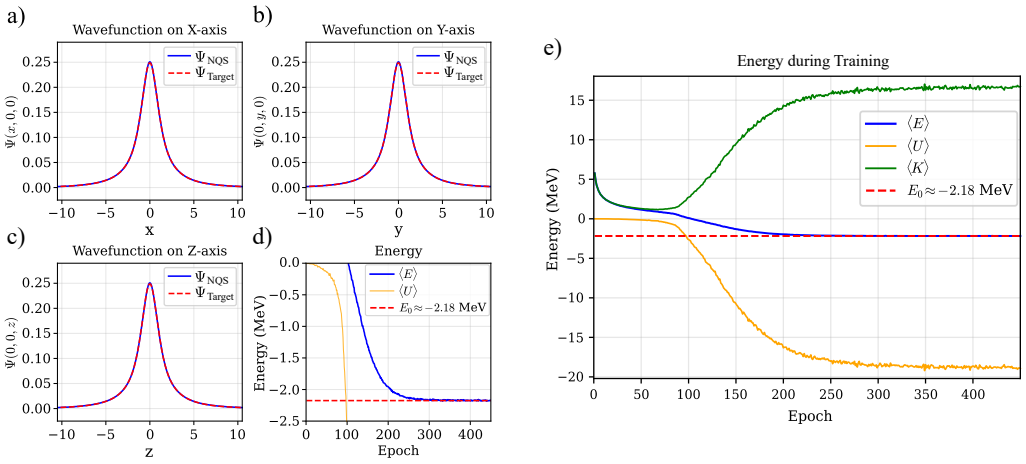


Figure 8: Ground state optimization of the physical-mass deuteron using SR. The wavefunction cuts in panel a)–c) compare the NQS along the Cartesian axes with the independent numerical solution. The solid blue curves correspond to the NQS, while the dashed red curves correspond to the independent numerical solution. The energy convergence plot in the panel d) shows the total variational energy and the potential contribution, represented by the blue and yellow curves, respectively, while the dashed horizontal line indicates the numerical ground state energy. Panel e) shows the same energy convergence figure on an expanded energy scale, making the kinetic contribution, represented by a green curve, visible.

The results, obtained averaging the corresponding radial quantities over the walker distribution, are in good agreement with the independent numerical measures:

$$\begin{cases} \langle r \rangle = 2.86 \pm 0.01 \text{ fm} \\ \dfrac{\sqrt{\langle r^2 \rangle}}{2} = 1.806 \pm 0.007 \text{ fm} \end{cases} \qquad \begin{cases} \langle r \rangle_{\text{num}} = 2.867 \text{ fm} \\ \dfrac{\sqrt{\langle r^2 \rangle_{\text{num}}}}{2} = 1.811 \text{ fm} \end{cases}$$

The obtained value of $\frac{\sqrt{\langle r^2 \rangle}}{2}$ is slightly below the literature value for the deuteron matter radius, 1.975 fm [44, 45], which reflects the limitations of the simplified LO interaction and of the pure $S$-wave description.

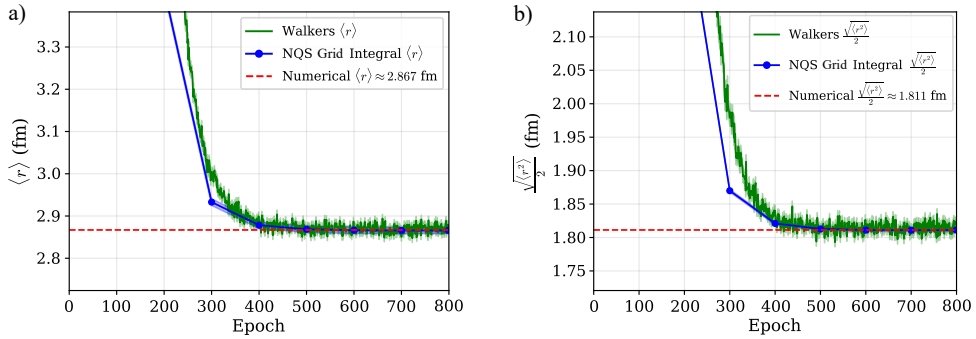


Figure 9: Convergence of the radial observables during the SR ground state optimization of the physical-mass deuteron. Panel a) shows the mean proton-neutron separation $\langle r \rangle$, while panel b) shows the point-nucleon radius, $\sqrt{\langle r^2 \rangle}/2$. The green and blue solid curves correspond to the walker and grid integration estimates of the observables, respectively, while the dashed red lines indicate the corresponding numerical values.

## 3.4 Heavy Deuteron Dynamics

As a first step towards the dynamical simulation of the physical deuteron, we consider a mass-rescaled system in which both the proton and neutron masses are multiplied by a

factor of 10. The interaction is kept unchanged, but the larger reduced mass decreases the kinetic energy prefactor of the Hamiltonian in Eq. (3.2). The same attractive potential produces a deeper and more spatially compact bound state. As a result, the characteristic asymptotic exponential tail of this mass-rescaled deuteron decays more rapidly, reducing the sensitivity of the simulation to the long-distance behaviour of the ansatz.

Since the mass-rescaled system corresponds to a different Hamiltonian, its ground state must first be obtained again. The ground state energy, obtained through SR, is in good agreement with the independent numerical solution,

$$E_0 = -77.2872 \pm 0.0008 \text{ MeV}, \qquad E_{\text{num}} = -77.2875 \text{ MeV},$$

which confirms that the NQS accurately represents the heavy deuteron before the dynamics are triggered with a regularization parameter $\varepsilon_{\text{TDVP}} = 5 \cdot 10^{-7}$.

Figure 10 displays the time evolution of the expectation value of the proton-neutron separation after applying a quadratic phase kick with $\beta = 0.1$ fm$^{-2}$ (see Eq. (3.6)). The observable is estimated both by averaging the radial positions of the walkers and by evaluating the wavefunction on a mesh and performing the corresponding grid integration. We also compare these results with the independent numerical solution.

The phase kick excites a breathing mode, in which the wavefunction periodically expands and contracts. This is reflected in the oscillatory behaviour of $\langle r \rangle$, whose relative variation remains below 5%. The grid integration estimate follows the numerical solution very closely, with a maximum difference of $1.5 \cdot 10^{-3}$ fm over the considered time interval. The walker estimate also follows the oscillations, with the expected Monte Carlo fluctuations.

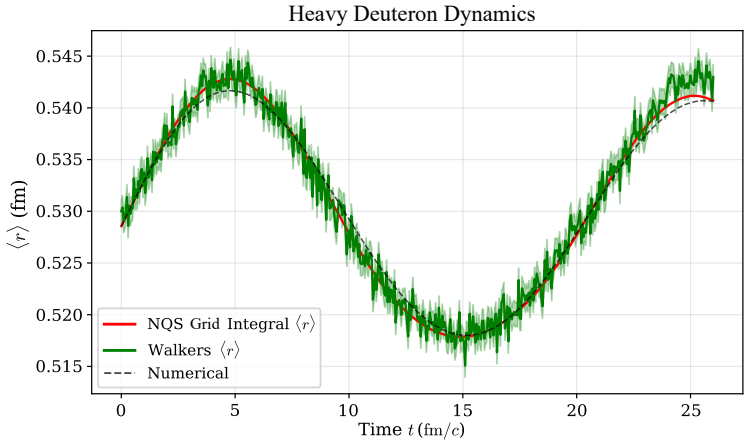


Figure 10: Time evolution of the expectation value of the proton-neutron separation for the heavy deuteron after a phase kick with $\beta = 0.1$ fm$^{-2}$. The green curve shows the NQS estimate obtained from the walkers, the red curve shows the NQS estimate obtained from the grid integral, and the black dashed curve corresponds to the independent numerical solution.

Figure 11 illustrates the quality of the TDVP time evolution through two diagnostics. The energy expectation value is monitored in the left panel, and a linear fit is used to quantify the drift over the simulated time interval. The fidelity with respect to the numerical solution, shown on the right panel, is computed both from the walker distribution and from the grid representation of the NQS.

The energy remains close to its initial value during most of the evolution, although there is a small upward drift that becomes more pronounced near the end of the simulation. The linear fit gives a total drift of $6.4 \cdot 10^{-2}$ MeV. Similarly, the fidelity remains close to unity throughout the simulation, with a more pronounced decrease at late times, reaching a final

value of 0.9994. Together with the agreement observed in Fig. 10, these results show that the TDVP evolution gives a reliable description of the mass-rescaled system.

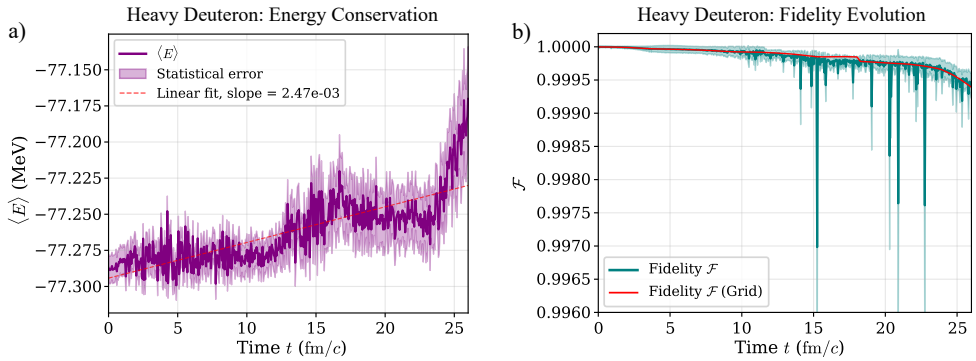


Figure 11: Energy (left panel) and fidelity (right panel) with respect to the independent numerical solution during the TDVP evolution of the heavy deuteron. In panel a), the purple curve and the shaded region show the energy expectation value and its statistical error, respectively, while the dashed red line shows the linear fit that quantifies the energy drift. In panel b), the teal curve and the shaded region show the fidelity estimated from the walker ensemble and its statistical error, respectively, while the red curve shows the grid integration results for the fidelity.

## 3.5 Deuteron Dynamics in a Periodic Box

As a final dynamical test, we consider the physical-mass deuteron inside a finite cubic box with periodic boundary conditions. The dynamics are initialized from the optimized NQS ground state obtained in Sec. 3.3. The use of a finite box is motivated by the extended tail of the deuteron wavefunction, which makes the TDVP evolution particularly sensitive to the long-distance behaviour of the variational ansatz. Periodic boundary conditions must be imposed both on the walkers and on the neural network. After each Metropolis proposal, the walker positions are mapped back into the interval $[-L/2, L/2]$ through periodic wrapping. However, since the network has no built-in reason to satisfy matching values and derivatives at the boundaries, feeding the wrapped Cartesian coordinates directly to the neural network would not guarantee a smooth periodic function. To enforce this structure, each coordinate is instead mapped to

$$x_k \longrightarrow \left(\cos\frac{2\pi x_k}{L}, \sin\frac{2\pi x_k}{L}\right), \qquad k = x, y, z, \tag{3.7}$$

so that the input layer receives six periodic variables instead of the three Cartesian coordinates. This construction ensures that the NQS is periodic and smooth across the boundaries of the box by making the input itself periodic.

The envelope and the radial phase kick are also adapted to the periodic geometry. We use the periodic radial distance

$$r_{\sin}^2 = \sum_{k=x,y,z} \left(\frac{L}{\pi}\sin\frac{\pi x_k}{L}\right)^2, \tag{3.8}$$

which ensures that points whose coordinates differ by one box length $L$ are assigned the same distance from the origin. Moreover, close to the origin, this quantity reduces to the usual radial distance, $r_{\sin} \simeq r$, while remaining smooth at the boundaries. In contrast, since the Gaussian potential is centered at the origin, it is still evaluated using the minimum image distance.

Figure 12 shows the time evolution of the mean proton-neutron separation $\langle r \rangle$ after applying a radial phase kick $\beta = 7.5{\cdot}10^{-4}$ fm$^{-2}$ in a periodic box of $L = 60$. The observable

is estimated by averaging over the radial position of the walkers and by evaluating the NQS on a spatial mesh and performing the corresponding grid integral, and both estimates are compared with the independent numerical solution.

The phase kick excites a radial breathing mode of the deuteron. Both the grid integration and the walker estimates closely follow the numerical result over the full simulated time interval, reproducing the initial expansion and the subsequent contraction of the wavefunction. The statistical uncertainty of the Monte Carlo estimate is visibly larger in this case, consistent with Fig. 9 and expected when sampling such a wide spatial distribution. Compared to the heavy deuteron dynamics of Fig. 10, the physical-mass deuteron exhibits a much longer oscillation period and a larger radial amplitude, reflecting its more extended and loosely bound nature. As the simulation progresses, numerical errors accumulate over time, which manifests as a slight damping in the amplitude of the oscillations. The largest differences between the estimates and the numerical solution are found at the antinodes of the curve, reaching a maximum discrepancy of $1.2 \cdot 10^{-2}$ fm with the grid-integration estimate.

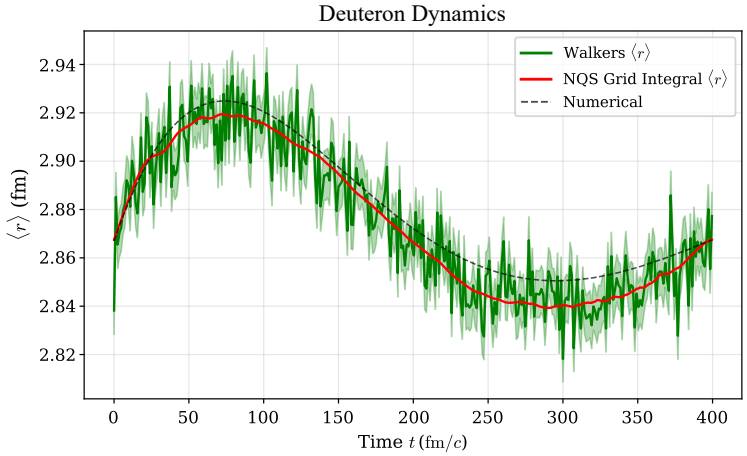


Figure 12: Same as Fig. 10, but for the physical-mass deuteron in a periodic box of $L = 60$ fm after a phase kick with $\beta = 7.5 \cdot 10^{-4}$ fm$^{-2}$.

Fig. 13 displays two complementary diagnostics for the TDVP evolution. The energy expectation value is plotted together with its Monte Carlo statistical error, and a linear fit is included to quantify the drift. The fidelity is estimated both from the walker ensemble and from the grid representation of the NQS.

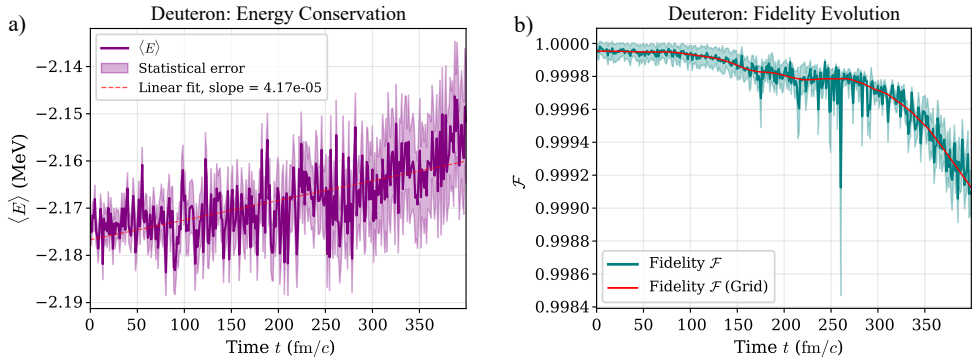


Figure 13: Same as Fig. 11, but for the physical-mass deuteron in a periodic box of $L = 60$ fm after a phase kick with $\beta = 7.5 \cdot 10^{-4}$ fm$^{-2}$

The energy remains close to the optimized ground state value throughout the evolution,

although a small upward drift is visible. Unlike the drift observed for the QHO, the energy increase here is not linear and visibly accelerates past 250 fm/$c$. The linear fit yields an approximate energy variation of $1.7 \cdot 10^{-2}$ MeV over the considered time interval. The fidelity also remains very close to unity during the evolution. Its decrease becomes more pronounced at late times, although it stays above 0.999 throughout the full simulation. Ultimately, by mitigating the boundary issues through a periodic NQS architecture, this approach successfully extends the TDVP framework to capture the real-time dynamics of the spatially extended, weakly bound deuteron.

# 4 Conclusions and Outlook

This thesis has explored the use of NQS for ground state optimization and real-time dynamics by studying the three-dimensional QHO and the deuteron. The wavefunction was parametrized through two real-valued FFNNs, representing separately the logarithm of the amplitude and the phase. Throughout the simulations, expectation values and variational quantities were estimated using VMC sampling. Ground states were obtained using SR, while time evolution was performed through the McLachlan TDVP. Both procedures act by updating the network parameters of the variational representation of the quantum state, which is sampled through Monte Carlo configurations rather than represented on a fixed grid.

The method was first validated with the QHO, for which both the ground state and dynamics are known analytically. The SR optimization accurately described the ground states of the isotropic and anisotropic configurations considered, with relative energy errors below 0.007%. The subsequent TDVP evolution reproduced the expected coherent state motion, with fidelity remaining above 0.9986 for the more demanding anisotropic case. The same framework was then applied to the deuteron described by an LO pionless effective field theory interaction. Although the network takes Cartesian coordinates as input, the NQS learned the spherical symmetry of the system without imposing it explicitly in the network architecture. For the physical-mass system, the optimized ground state was found with a relative energy error of 0.2 % with respect to the independent numerical solution, while for the mass-rescaled deuteron the optimized ground state energy had a relative error of $3.9 \cdot 10^{-4}$ %. The TDVP evolution of both nuclear systems was performed successfully, with fidelity remaining above the 0.999 level.

The successful application of the NQS framework to both the QHO and the deuteron demonstrates that NQS provide flexible variational ansätze capable not only of characterizing ground states, but also of accurately describing real-time dynamics. The agreement obtained in these two different systems, from an analytically solvable benchmark to an interacting nuclear problem, supports the use of this framework as a promising variational approach to quantum dynamics in real space formulations.

There are, nevertheless, limitations to the present work. Although stable and accurate dynamics were obtained over the simulated time intervals, the method cannot yet be extended reliably to longer times. During real-time evolution, statistical noise and numerical errors in the solution of the TDVP equations accumulate, eventually leading to a complete breakdown of the simulation. In addition, the systems considered here remain relatively simple. A natural continuation of this work would therefore be to apply the same methodology to more realistic nuclear Hamiltonians and, ultimately, to systems with more than two particles. This would also require exploring more expressive neural-network architectures and improved ways of imposing boundary conditions and asymptotic behaviour, as well as refining numerical aspects such as the regularization strategy and the linear solver for the TDVP equations.

## Bibliography


[1] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. ISBN 9780262337373.

[2] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. DOI: 10.1038/nature14539.

[3] G. Carleo and M. Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017. DOI: 10.1126/science.aag2302.

[4] S.-J. Ran, E. Tirrito, C. Peng, X. Chen, et al. *Tensor Network Contractions: Methods and Applications to Quantum Many-Body Systems*, volume 964 of *Lecture Notes in Physics*. Springer Cham, 2020. ISBN 978-3-030-34488-7. DOI: 10.1007/978-3-030-34489-4.

[5] M. C. Bañuls. Tensor network algorithms: A route map. *Annual Review of Condensed Matter Physics*, 14(1):173–191, 2023. DOI: 10.1146/annurev-conmatphys-040721-022705.

[6] J. I. Cirac, D. Pérez-García, N. Schuch, and F. Verstraete. Matrix product states and projected entangled pair states: Concepts, symmetries, theorems. *Reviews of Modern Physics*, 93(4):045003, 2021. DOI: 10.1103/RevModPhys.93.045003.

[7] S. Carignano, G. Lami, J. De Nardis, and L. Tagliacozzo. Overcoming the entanglement barrier with sampled tensor networks, 2025. arXiv:2505.09714.

[8] M. Grundner, P. Westhoff, F. B. Kugler, O. Parcollet, et al. Complex time evolution in tensor networks and time-dependent green's functions. *Physical Review B*, 109(15): 155124, 2024. DOI: 10.1103/PhysRevB.109.155124.

[9] M. Medvidović and J. R. Moreno. Neural-network quantum states for many-body physics. *The European Physical Journal Plus*, 139(7):631, 2024. DOI: 10.1140/epjp/s13360-024-05311-y.

[10] A. Lovato, G. Carleo, B. Fore, H.-J. Morten, et al. Neural-network quantum states for the nuclear many-body problem, 2026. arXiv:2602.13826.

[11] A. Dawid, J. Arnold, B. Requena, A. Gresch, et al. *Machine Learning in Quantum Sciences*. Cambridge University Press, 2025. ISBN 9781009504935. DOI: 10.1017/9781009504942.

[12] J. W. T. Keeble and A. Rios. Machine learning the deuteron. *Physics Letters B*, 809: 135743, 2020. DOI: 10.1016/j.physletb.2020.135743.

[13] J. Rozalén, J. W. T. Keeble, and A. Rios. Machine learning the deuteron: new architectures and uncertainty quantification. *The European Physical Journal Plus*, 139(2):153, 2024. DOI: 10.1140/epjp/s13360-024-04983-w.

[14] C. Adams, G. Carleo, A. Lovato, and N. Rocco. Variational monte carlo calculations of $A \leq 4$ nuclei with an artificial neural-network correlator ansatz. *Physical Review Letters*, 127(2):022502, 2021. DOI: 10.1103/PhysRevLett.127.022502.

[15] A. Lovato, C. Adams, G. Carleo, and N. Rocco. Hidden-nucleons neural-network quantum states for the nuclear many-body problem. *Physical Review Research*, 4(4): 043178, 2022. DOI: 10.1103/PhysRevResearch.4.043178.

[16] A. Romero-Ros, J. Rozalén, and A. Rios. Quantum dynamics with time-dependent neural quantum states, 2025. arXiv:2509.24865.

[17] I. López and C. B. Mendl. Real time evolution with neural-network quantum states. *Quantum*, 6:627, 2022. DOI: 10.22331/q-2022-01-20-627.

[18] M. Schmitt and M. Heyl. Quantum many-body dynamics in two dimensions with artificial neural networks. *Physical Review Letters*, 125(10):100503, 2020. DOI: 10.1103/PhysRevLett.125.100503.

[19] M. Fernández. NQS-Dynamics, 2026. GitHub repository. Accessed: 2026-07-07.
[20] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989. DOI: 10.1007/BF02551274.
[21] T. Heightman and M. Płodzień. Deep learning in classical and quantum physics, 2025. arXiv:2508.10666.
[22] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2017. arXiv:1412.6980.
[23] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, et al. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21 (6):1087–1092, 1953. DOI: 10.1063/1.1699114.
[24] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970. DOI: 10.1093/biomet/57.1.97.
[25] M. H. Kalos and P. A. Whitlock. *Monte Carlo Methods*. Wiley-VCH, 2nd edition, 2008. ISBN 9783527407606. DOI: 10.1002/9783527626212.
[26] F. Becca and S. Sorella. *Quantum Monte Carlo Approaches for Correlated Systems*. Cambridge University Press, 2017. ISBN 9781316417041. DOI: 10.1017/9781316417041.
[27] S. Sorella. Green function monte carlo with stochastic reconfiguration. *Physical Review Letters*, 80(20):4558–4561, 1998. DOI: 10.1103/PhysRevLett.80.4558.
[28] S. Sorella. Generalized lanczos algorithm for variational quantum monte carlo. *Physical Review B*, 64(2):024512, 2001. DOI: 10.1103/PhysRevB.64.024512.
[29] S. Amari. Natural gradient works efficiently in learning. *Neural Computation*, 10(2): 251–276, 1998. DOI: 10.1162/089976698300017746.
[30] J. Stokes, J. Izaac, N. Killoran, and G. Carleo. Quantum natural gradient. *Quantum*, 4:269, 2020. DOI: 10.22331/q-2020-05-25-269.
[31] C.-Y. Park and M. J. Kastoryano. Geometry of learning neural quantum states. *Physical Review Research*, 2(2):023232, 2020. DOI: 10.1103/PhysRevResearch.2.023232.
[32] M. Kolodrubetz, D. Sels, P. Mehta, and A. Polkovnikov. Geometry and non-adiabatic response in quantum and classical systems. *Physics Reports*, 697:1–87, 2017. DOI: 10.1016/j.physrep.2017.07.001.
[33] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 4th edition, 2013. ISBN 9781421407944.
[34] L. Hackl, T. Guaita, T. Shi, J. Haegeman, et al. Geometry of variational methods: dynamics of closed quantum systems. *SciPost Physics*, 9:048, 2020. DOI: 10.21468/SciPostPhys.9.4.048.
[35] A. D. McLachlan. A variational solution of the time-dependent schrödinger equation. *Molecular Physics*, 8(1):39–44, 1964. DOI: 10.1080/00268976400100041.
[36] A. Raab. On the dirac–frenkel/mclachlan variational principle. *Chemical Physics Letters*, 319(5):674–678, 2000. DOI: 10.1016/S0009-2614(00)00200-1.
[37] C. Cohen-Tannoudji, B. Diu, and F. Laloë. *Quantum Mechanics, Volume 1: Basic Concepts, Tools, and Applications*. Wiley-VCH, 2nd edition, 2020. ISBN 9783527345533.
[38] J.C. Butcher. A history of runge-kutta methods. *Applied Numerical Mathematics*, 20 (3):247–260, 1996. DOI: https://doi.org/10.1016/0168-9274(95)00108-5.
[39] J. Kirscher, N. Barnea, D. Gazit, F. Pederiva, et al. Spectra and scattering of light lattice nuclei from effective field theory. *Physical Review C*, 92(5):054002, 2015. DOI: 10.1103/PhysRevC.92.054002.

[40] M. D. Feit, J. A. Fleck, and A. Steiger. Solution of the schrödinger equation by a spectral method. *Journal of Computational Physics*, 47(3):412–433, 1982. DOI: 10.1016/0021-9991(82)90091-2.

[41] M. R. Hermann and J. A. Fleck. Split-operator spectral method for solving the time-dependent schrödinger equation in spherical coordinates. *Physical Review A*, 38(12): 6000–6012, 1988. DOI: 10.1103/PhysRevA.38.6000.

[42] P. D. Stevenson and S. Fracasso. Extracting structure information from tdhf. *Journal of Physics G: Nuclear and Particle Physics*, 37(6):064030, 2010. DOI: 10.1088/0954-3899/37/6/064030.

[43] H. Müther and A. Polls. Two-body correlations in nuclear systems. *Progress in Particle and Nuclear Physics*, 45(1):243–334, 2000. DOI: 10.1016/s0146-6410(00)00105-8.

[44] M. Garçon and J. W. Van Orden. The deuteron: Structure and form factors. In *Advances in Nuclear Physics*, pages 293–378. Springer US, 2001. ISBN 9780306479151. DOI: 10.1007/0-306-47915-X_4.

[45] U. D. Jentschura, A. Matveev, C. G. Parthey, J. Alnis, et al. Hydrogen-deuterium isotope shift: From the $1S$–$2S$ transition frequency to the proton-deuteron charge-radius difference. *Physical Review A*, 83(4):042505, 2011. DOI: 10.1103/PhysRevA.83.042505.

# A Metropolis-Hastings Algorithm

The Metropolis-Hastings algorithm is a Markov Chain Monte Carlo method used to generate representative samples from a given probability distribution (Alg. 1). As stated in Eq. (1.6), the expectation value of an observable $g(\mathbf{x})$ can be estimated as

$$\int d\mathbf{x}\; p_{\boldsymbol{\theta}}(\mathbf{x}) g(\mathbf{x}) \approx \frac{1}{N} \sum_{i=1}^{N} g(\mathbf{x}_i), \tag{A.1}$$

where the configurations $\{\mathbf{x}_i\}_{i=1}^{N}$ are sampled from the probability distribution $p_{\boldsymbol{\theta}}(\mathbf{x})$. For a parametrized quantum state, these configurations are sampled using the Born probability distribution,

$$p_{\boldsymbol{\theta}}(\mathbf{x}) = \frac{|\Psi_{\boldsymbol{\theta}}(\mathbf{x})|^2}{\langle \Psi_{\boldsymbol{\theta}} | \Psi_{\boldsymbol{\theta}} \rangle}, \tag{A.2}$$

which ensures that the configurations contributing most strongly to the expectation value are the regions where $|\Psi_{\boldsymbol{\theta}}(\mathbf{x})|^2$ is large. The sampling process begins by choosing an arbitrary initial configuration $\mathbf{x}$. A small random displacement $\boldsymbol{\xi}$ is then proposed to generate a new configuration $\mathbf{x}'$. The proposed move is accepted or rejected according to the probability ratio between the new and old configurations. Satisfying the detailed balance condition

$$p_{\boldsymbol{\theta}}(\mathbf{x}) T(\mathbf{x} \to \mathbf{x}') = p_{\boldsymbol{\theta}}(\mathbf{x}') T(\mathbf{x}' \to \mathbf{x}) \tag{A.3}$$

is enough to ensure that the random walk correctly samples the desired probability distribution, where $T(\mathbf{x} \to \mathbf{x}')$ is the transition probability [25]. For symmetric proposal distributions, the standard choice is

$$T(\mathbf{x} \to \mathbf{x}') = \min\left[1, \frac{p_{\boldsymbol{\theta}}(\mathbf{x}')}{p_{\boldsymbol{\theta}}(\mathbf{x})}\right]. \tag{A.4}$$

Note that, because the acceptance ratio only depends on the relative probability,

$$\frac{p_{\boldsymbol{\theta}}(\mathbf{x}')}{p_{\boldsymbol{\theta}}(\mathbf{x})} = \frac{|\Psi_{\boldsymbol{\theta}}(\mathbf{x}')|^2}{|\Psi_{\boldsymbol{\theta}}(\mathbf{x})|^2}, \tag{A.5}$$

the computationally expensive normalization constant cancels out and therefore does not need to be computed explicitly.

**Algorithm 1** Metropolis-Hastings sampling

1: Initialize $N$ walkers $\{\mathbf{x}_k\}_{k=1}^{N}$.
2: Choose a proposal step size $\Delta$ and a number of sweeps $N_{\text{sweeps}}$.
3: **for** $s = 1$ to $N_{\text{sweeps}}$ **do**
4:  **for** $k = 1$ to $N$ **do**
5:  Propose a new configuration $\mathbf{x}'_k = \mathbf{x}_k + \Delta \cdot \boldsymbol{\xi}_k$.
6:  Compute the acceptance probability

$$a_k = \min\left[1, \frac{|\Psi_{\boldsymbol{\theta}}(\mathbf{x}'_k)|^2}{|\Psi_{\boldsymbol{\theta}}(\mathbf{x}_k)|^2}\right].$$

7:  **if** $u_k \sim \mathcal{U}(0,1) < a_k$ **then**
8:  Accept the proposal $\mathbf{x}_k \leftarrow \mathbf{x}'_k$.
9:  **else**
10:  Reject the proposal and keep $\mathbf{x}_k$ unchanged.
11: **return** The final walker ensemble $\{\mathbf{x}_k\}_{k=1}^{N}$.

In practice, when the walkers are initialized and before measurements are taken, the walker ensemble must be equilibrated with respect to the probability distribution of the wavefunction. For this reason, an initial thermalization stage is performed, in which the walkers are evolved for 500 Metropolis steps but the generated configurations are discarded. Moreover, each walker has a position in configuration space and evolves through a Markov chain of Metropolis proposals. As a result, two consecutive configurations are generally correlated. For this reason, observables are not evaluated after every single Metropolis proposal. Instead, the walker ensemble is propagated for 100 Metropolis steps before computing the relevant observables, reducing the correlation between configurations used in successive measurements.

## B Energy Gradient

The variational energy can be written as the expectation value of the local energy over the probability distribution $p_{\boldsymbol{\theta}}$:

$$E_{\boldsymbol{\theta}} = \langle E_L(\mathbf{x}) \rangle = \int d\mathbf{x}\, p_{\boldsymbol{\theta}}(\mathbf{x}) E_L(\mathbf{x}), \tag{B.1}$$

where

$$p_{\boldsymbol{\theta}}(\mathbf{x}) = \frac{|\Psi_{\boldsymbol{\theta}}(\mathbf{x})|^2}{\langle \Psi_{\boldsymbol{\theta}} | \Psi_{\boldsymbol{\theta}} \rangle}, \qquad E_L(\mathbf{x}) = \frac{\hat{H}\Psi_{\boldsymbol{\theta}}(\mathbf{x})}{\Psi_{\boldsymbol{\theta}}(\mathbf{x})}. \tag{B.2}$$

Using the logarithmic derivative identity $\partial_{\theta_k} p_{\boldsymbol{\theta}}(\mathbf{x}) = p_{\boldsymbol{\theta}}(\mathbf{x}) \partial_{\theta_k} \log p_{\boldsymbol{\theta}}(\mathbf{x})$, the derivative of the variational energy becomes

$$\partial_{\theta_k} E_{\boldsymbol{\theta}} = \int d\mathbf{x}\, p_{\boldsymbol{\theta}}(\mathbf{x}) \left[ E_L(\mathbf{x}) \partial_{\theta_k} \log p_{\boldsymbol{\theta}}(\mathbf{x}) + \partial_{\theta_k} E_L(\mathbf{x}) \right]. \tag{B.3}$$

With the definition $O_k(\mathbf{x}) = \partial_{\theta_k} \log \Psi_{\boldsymbol{\theta}}(\mathbf{x})$, the logarithmic derivative of the normalized probability distribution is rewritten as

$$\partial_{\theta_k} \log p_{\boldsymbol{\theta}}(\mathbf{x}) = O_k(\mathbf{x}) + O_k^*(\mathbf{x}) - \langle O_k + O_k^* \rangle, \tag{B.4}$$

where $\partial_{\theta_k} \log |\Psi_{\boldsymbol{\theta}}(\mathbf{x})|^2 = O_k(\mathbf{x}) + O_k^*(\mathbf{x})$ has been used. Substituting this expression into Eq. (B.3) gives

$$\partial_{\theta_k} E_{\boldsymbol{\theta}} = \langle E_L \left[ O_k + O_k^* - \langle O_k + O_k^* \rangle \right] \rangle + \langle \partial_{\theta_k} E_L \rangle. \tag{B.5}$$

In order to evaluate the last term, since $\partial_{\theta_k} \Psi_{\boldsymbol{\theta}}(\mathbf{x}) = O_k(\mathbf{x}) \Psi_{\boldsymbol{\theta}}(\mathbf{x})$, the derivative of the local energy is written as

$$\partial_{\theta_k} E_L(\mathbf{x}) = \frac{\hat{H} \partial_{\theta_k} \Psi_{\boldsymbol{\theta}}(\mathbf{x})}{\Psi_{\boldsymbol{\theta}}(\mathbf{x})} - E_L(\mathbf{x}) O_k(\mathbf{x}). \tag{B.6}$$

Averaging over $p_{\boldsymbol{\theta}}$ yields

$$\langle \partial_{\theta_k} E_L \rangle = \frac{\int d\mathbf{x}\, \Psi_{\boldsymbol{\theta}}^*(\mathbf{x}) \hat{H} \partial_{\theta_k} \Psi_{\boldsymbol{\theta}}(\mathbf{x})}{\langle \Psi_{\boldsymbol{\theta}} | \Psi_{\boldsymbol{\theta}} \rangle} - \langle E_L O_k \rangle, \tag{B.7}$$

which can be further simplified using the Hermiticity of $\hat{H}$

$$\int d\mathbf{x}\, \Psi_{\boldsymbol{\theta}}^*(\mathbf{x}) \hat{H} \partial_{\theta_k} \Psi_{\boldsymbol{\theta}}(\mathbf{x}) = \int d\mathbf{x}\, (\hat{H}\Psi_{\boldsymbol{\theta}})^*(\mathbf{x}) \partial_{\theta_k} \Psi_{\boldsymbol{\theta}}(\mathbf{x}), \tag{B.8}$$

to

$$\langle \partial_{\theta_k} E_L \rangle = \langle E_L^* O_k \rangle - \langle E_L O_k \rangle. \tag{B.9}$$

Substituting this result back into Eq. (B.5), the terms involving $\langle E_L O_k \rangle$ cancel, resulting in

$$\partial_{\theta_k} E_{\boldsymbol{\theta}} = \langle E_L O_k^* \rangle + \langle E_L^* O_k \rangle - E_{\boldsymbol{\theta}} \langle O_k + O_k^* \rangle . \tag{B.10}$$

Finally, since $E_{\boldsymbol{\theta}}$ is real, this can be written compactly as

$$\partial_{\theta_k} E_{\boldsymbol{\theta}} = 2 \operatorname{Re} \left[ \langle (E_L(\mathbf{x}) - E_{\boldsymbol{\theta}}) O_k^*(\mathbf{x}) \rangle \right] \tag{B.11}$$

or, equivalently,

$$\partial_{\theta_k} E_{\boldsymbol{\theta}} = 2 \operatorname{Re} \left[ \langle \Delta O_k^*(\mathbf{x}) \Delta E_L(\mathbf{x}) \rangle \right] , \tag{B.12}$$

where $\Delta O_k(\mathbf{x}) = O_k(\mathbf{x}) - \langle O_k \rangle$, $\Delta E_L(\mathbf{x}) = E_L(\mathbf{x}) - E_{\boldsymbol{\theta}}$ and $\langle E_L - E_{\boldsymbol{\theta}} \rangle = 0$ has been used.

# C Dependence on the Number of Walkers

The number of walkers controls the statistical accuracy of the Monte Carlo estimates used throughout the variational calculation. In order to study its effect, the isotropic QHO configuration introduced in Sec. 2.3 is simulated using different numbers of walkers. This configuration and its dynamics were characterized by

$$\omega_x = \omega_y = \omega_z = 1, \quad \mathbf{x}_0 = (0, 0, 0), \quad \mathbf{x}_0' = (1, 1, 1), \quad \mathbf{p}_0 = (1, 1, 1).$$

Although this comparison is performed only for the above parameters, the conclusions apply generally. The role of the number of walkers is to reduce the statistical fluctuations of the Monte Carlo estimators independently of the specific configuration or dynamics considered.

The ground state optimization is relatively insensitive to the number of walkers, and accurate ground states can already be obtained with fewer than the 50,000 walkers used in the main simulations. However, the dynamics are significantly more sensitive to the number of walkers. In the TDVP evolution, statistical noise affects the estimation of the terms in Eq. (1.26), and these errors can accumulate over time. Therefore, a larger walker population is required to obtain stable trajectories, good energy conservation and high fidelity with respect to the exact solution.

Figure 14 compares the TDVP dynamics of the isotropic QHO for different numbers of walkers. The energy expectation value and the fidelity with respect to the exact coherent state solution are considered, and the different curves correspond to simulations performed with different walker populations.

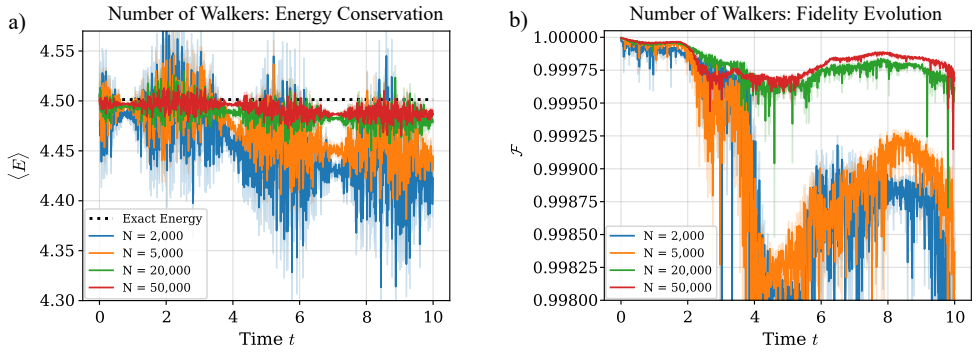


Figure 14: Effect of the number of walkers on the QHO dynamics. Panel a) shows the energy expectation value, while panel b) shows fidelity with respect to the exact coherent state solution. Each color corresponds to a different number of walkers.

For small walker populations, the energy (panel a) exhibits larger fluctuations and a more pronounced drift during the evolution. Increasing the number of walkers reduces the noise in the Monte Carlo estimate and improves energy conservation. This can be seen most clearly by comparing the simulations with 2,000 and 50,000 walkers. In the former case, the energy decreases by $8 \cdot 10^{-2}$ over the simulated time interval, while for 50,000 walkers the drift is reduced to $1 \cdot 10^{-2}$. Similarly, the cases with fewer walkers display larger losses of fidelity (panel b) with respect to the exact coherent state dynamics and stronger fluctuations, while the simulations with 20,000 and especially 50,000 walkers remain much more stable throughout the time interval considered. In particular, the final fidelity falls to 0.9986 for 2,000 walkers, whereas for 50,000 walkers it remains at 0.9997.

## D Role of the Envelope Center in QHO Dynamics

As discussed in Sec. 2, the amplitude part of the variational ansatz includes an envelope. In the QHO simulations, this envelope is chosen to be Gaussian, so the logarithm of the amplitude takes the form

$$A_{\boldsymbol{\theta}}(\mathbf{x}) = A^{\mathrm{NN}}_{\boldsymbol{\theta}}(\mathbf{x}) - \alpha_{\mathrm{env}} |\mathbf{x} - \mathbf{x}_{\mathrm{env}}|^2.$$

At first sight, treating $\mathbf{x}_{\mathrm{env}}$ as an additional trainable parameter may seem unnecessary, since it increases the number of variational degrees of freedom. During the ground state optimization, the natural choice is to initialize the envelope center at the center of the harmonic potential $\mathbf{x}_{\mathrm{env}} = \mathbf{x}_0$. However, this choice becomes less obvious during the time evolution. After the dynamics are triggered, the probability density no longer remains centered around its initial position. If the envelope center is fixed, the large-distance decay imposed by the ansatz remains anchored to the original ground state configuration, which can restrict the ability of the NQS to accurately follow the moving wavefunction. Allowing $\mathbf{x}_{\mathrm{env}}$ to be optimized gives the ansatz more flexibility, while keeping the stabilizing effect of the Gaussian envelope.

In order to test this effect, the isotropic QHO configuration introduced in Sec. 2.3 is simulated again using both choices. This configuration and its dynamics are characterized by

$$\omega_x = \omega_y = \omega_z = 1, \quad \mathbf{x}_0 = (0,0,0), \quad \mathbf{x}_0' = (1,1,1), \quad \mathbf{p}_0 = (1,1,1).$$

Figure 15 shows the resulting TDVP dynamics for a fixed envelope center and for a trainable envelope center. The energy expectation value and the fidelity with respect to the exact coherent state solution are considered.

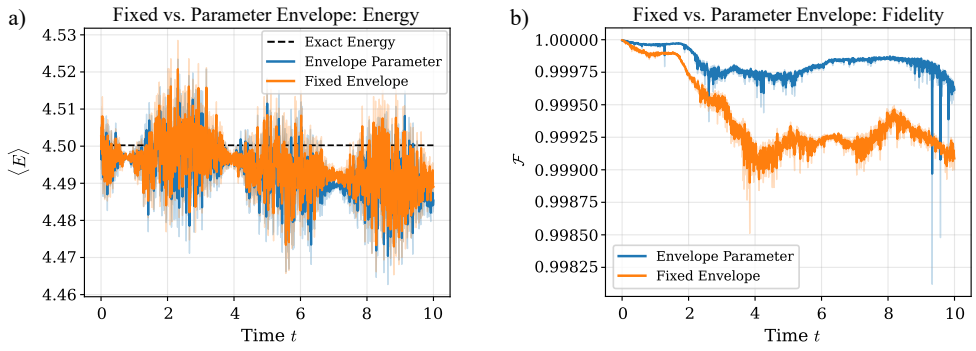


Figure 15: Comparison between a fixed envelope center and a trainable envelope center during the QHO dynamics with a shift $\boldsymbol{\Delta} = \mathbf{x}_0' - \mathbf{x}_0 = (1,1,1)$. Panel a) shows the energy expectation value, while panel b) shows the fidelity with respect to the exact coherent state solution. The blue and orange curves correspond to the trainable envelope and the fixed envelope, respectively.

Although both choices demonstrate similar energy conservation (panel a), the trainable-envelope simulation remains closer to unit fidelity throughout the whole simulation (panel b). With a trainable envelope center, the final fidelity remains at 0.99962, whereas for the fixed envelope it decreases to 0.99912.

This improvement becomes more notable for larger displacements, where a fixed envelope may artificially constrain or destroy the motion of the wavefunction. Fig. 16 shows the results for similar dynamics but for a larger displacement of the potential center $\boldsymbol{\Delta} = \mathbf{x}_0' - \mathbf{x}_0 = (2, 2, 2)$.

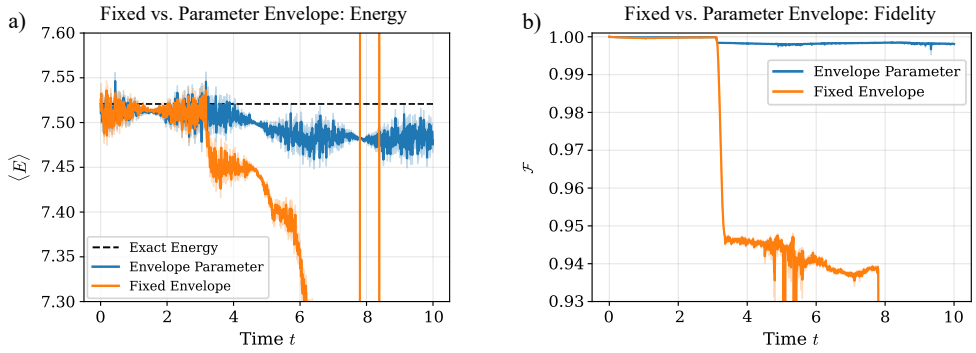


Figure 16: Comparison between a fixed envelope center and a trainable envelope center during the QHO dynamics, as in Fig. 15, but for a larger shift $\boldsymbol{\Delta} = \mathbf{x}_0' - \mathbf{x}_0 = (2, 2, 2)$.

In this case, the larger displacement makes the effect of the envelope center much more significant. The simulation with a trainable envelope remains stable, while the fixed-envelope simulation suffers a complete breakdown once the wavefunction moves sufficiently far from the initial envelope position, resulting in an energy spike of order $10^{10}$ (panel a) and a fidelity drop to values of order $10^{-3}$ (panel b).

An alternative possibility was also considered, namely to keep $\mathbf{x}_{\text{env}}$ as a non-trainable variable and update its value manually during the time evolution, following the mean position of the walker ensemble. However, this procedure led to significantly poorer results. A possible explanation for this is that the TDVP parameter update is computed for the variational state defined at the current value of the envelope center. If the envelope is then moved externally, this modification is not determined by the variational principle and can interfere with the parameter update. The network parameters may already be changing in order to displace the probability density, while the manually shifted envelope pushes the state in the same direction, effectively overcorrecting the motion.

## E Numerical Methods for TDVP Time Evolution

The TDVP time evolution requires solving the linear system of Eq. (1.26) at each time step in order to obtain the update of the network parameters. In simpler terms, this system takes the form

$$S\dot{\boldsymbol{\theta}} = \mathbf{F},$$

where $S$ is the real part of the QGT and $\mathbf{F}$ is the TDVP force vector. Since $S$ is estimated stochastically from the walker ensemble, it can be affected by Monte Carlo noise and may become ill-conditioned. For this reason, the diagonal regularization

$$S \to S + \varepsilon_{\text{TDVP}}\, \text{diag}(1 + |\mathbf{F}|)$$

is added before solving the system. Two different strategies are considered for solving this linear system. The first one is based on the Moore–Penrose pseudo-inverse, which can be applied even when $S$ is singular or not positive definite. This makes it relatively robust, but also computationally expensive. The second strategy uses a Cholesky decomposition, which is faster and numerically efficient, but requires $S$ to be positive definite [33].

In addition to the choice of linear solver, different Runge–Kutta integration schemes are also considered for parameter evolution. RK2 requires fewer TDVP evaluations per time step, while RK4 evaluates the TDVP update at more intermediate points. Normally, RK4 is expected to provide a more accurate integration. However, in our case each intermediate RK4 evaluation requires a new stochastic estimate of $S$ and $\mathbf{F}$. As a consequence, RK4 is more computationally expensive and is more exposed to noisy or ill-conditioned estimates of the TDVP linear system.

The comparison was performed considering all combinations of the introduced methods for the same isotropic QHO configuration considered in Sec. 2, characterized by

$$\omega_x = \omega_y = \omega_z = 1, \quad \mathbf{x}_0 = (0,0,0), \quad \mathbf{x}_0' = (1,1,1), \quad \mathbf{p}_0 = (1,1,1).$$

Fig. 17 compares the TDVP dynamics obtained with different combinations of time integrator and linear solver, using the same value of $\varepsilon_{\text{TDVP}} = 4 \cdot 10^{-5}$ in all cases. The energy expectation value and the fidelity with respect to the exact coherent state solution are considered, and the different curves correspond to the four combinations of RK2 or RK4 integration with either Cholesky or Moore–Penrose pseudo-inverse solution of the TDVP linear system.

When the same regularization is used for all methods, the RK4-based schemes yield slightly better overall time evolution. The Cholesky and pseudo-inverse solvers paired with RK4 lead to similar results, while the Cholesky solver paired with RK2 gives the lowest fidelity.

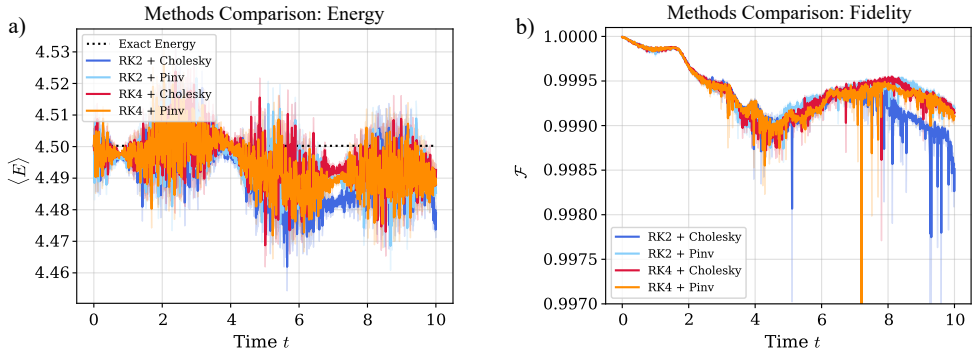


Figure 17: Comparison between RK2 and RK4 time integration, combined with either Moore–Penrose pseudo-inverse or Cholesky solution of the TDVP linear system, using the same value of $\varepsilon_{\text{TDVP}} = 4 \cdot 10^{-5}$ for all cases. Panel a) shows the energy expectation value, while panel b) shows the fidelity with respect to the exact coherent state dynamics. Each color corresponds to a different combination of time integrator and linear solver.

A different picture emerges when the value of $\varepsilon_{\text{TDVP}}$ is optimized separately for RK2 and RK4. Fig. 18 repeats the same comparison using $\varepsilon_{\text{TDVP}} = 7 \cdot 10^{-6}$ for the RK2 schemes and $\varepsilon_{\text{TDVP}} = 4 \cdot 10^{-5}$ for the RK4 schemes.

With separately optimized regularization parameters, the RK2 schemes become competitive and even provide higher fidelity than the RK4 schemes over the simulated time interval. In particular, the RK2+Cholesky evolution remains closest to unit fidelity, reaching a final fidelity of 0.9996, while the RK4 simulations end with fidelities around 0.9987.

This shows that the order of the Runge–Kutta method is not the only relevant factor in a stochastic TDVP simulation. Since RK2 requires fewer TDVP evaluations per time step, it is less sensitive to the accumulation of noise and errors in the estimates of $S$ and $\mathbf{F}$, and a smaller regularization can be used. Consequently, the quality of the Monte Carlo estimate of the TDVP linear system and the amount of regularization required to solve it also play an important role.

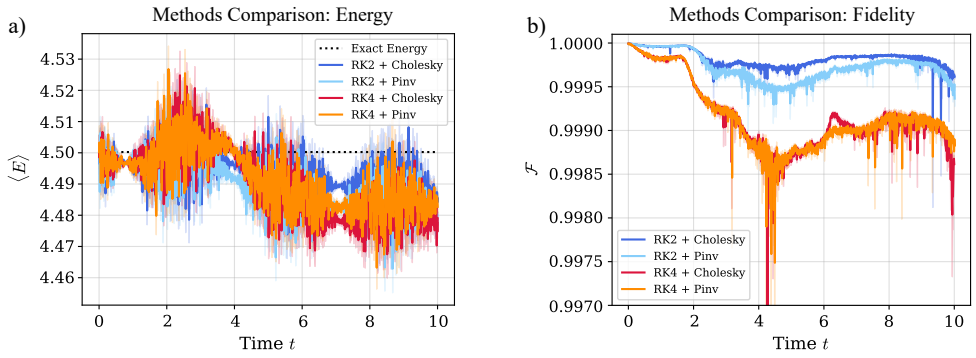


Figure 18: Same comparison as in Fig. 17, but using separately optimized values of $\varepsilon_{\mathrm{TDVP}}$ for the RK2 and RK4 methods. The RK2 schemes use $\varepsilon_{\mathrm{TDVP}} = 7 \cdot 10^{-6}$, while the RK4 schemes use $\varepsilon_{\mathrm{TDVP}} = 4 \cdot 10^{-5}$.