

UC Berkeley

UC Berkeley Previously Published Works

Title

A Kaczmaz-inspired approach to accelerate the optimization of neural network wavefunctions

Permalink

<https://escholarship.org/uc/item/2v45r6qz>

Authors

Goldshlager, Gil

Abrahamsen, Nilin

Lin, Lin

Publication Date

2024-11-01


DOI

10.1016/j.jcp.2024.113351

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed



A Kaczmarz-inspired approach to accelerate the optimization of neural network wavefunctions

Gil Goldshlager^{a,*}, Nilin Abrahamsen^b, Lin Lin^{a,c}

^a Department of Mathematics, University of California, Berkeley, CA 94720, USA

^b The Simons Institute for the Theory of Computing, Berkeley, CA 94720, USA

^c Applied Mathematics and Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

ABSTRACT

Neural network wavefunctions optimized using the variational Monte Carlo method have been shown to produce highly accurate results for the electronic structure of atoms and small molecules, but the high cost of optimizing such wavefunctions prevents their application to larger systems. We propose the Subsampled Projected-Increment Natural Gradient Descent (SPRING) optimizer to reduce this bottleneck. SPRING combines ideas from the recently introduced minimum-step stochastic reconfiguration optimizer (MinSR) and the classical randomized Kaczmarz method for solving linear least-squares problems. We demonstrate that SPRING outperforms both MinSR and the popular Kronecker-Factored Approximate Curvature method (KFAC) across a number of small atoms and molecules, given that the learning rates of all methods are optimally tuned. For example, on the oxygen atom, SPRING attains chemical accuracy after forty thousand training iterations, whereas both MinSR and KFAC fail to do so even after one hundred thousand iterations.

1. Introduction

Predicting the properties of molecules and materials from first principles has numerous applications. For many chemical properties, it suffices to work within the Born-Oppenheimer approximation, in which the nuclei are viewed as classical point charges and only the electrons exhibit quantum-mechanical behavior. The study of chemistry through this lens is known as *electronic structure theory*.

Within electronic structure theory, methods to model the many-body electron wavefunction include Hartree-Fock theory, configuration interaction methods, and coupled cluster theory. A typical ansatz for such methods is a sum of Slater determinants which represent antisymmetrized products of single-particle states. The benefit of such an ansatz is that the energy and other properties of the wavefunction can be evaluated analytically from pre-computed few-particle integrals.

Another approach to the electronic structure problem is the variational Monte Carlo method (VMC) [1,2]. In VMC, the properties of the wavefunction are calculated using Monte Carlo sampling rather than direct numerical integration, and the energy is variationally minimized through a stochastic optimization procedure. This increases the cost of the calculations, especially when high accuracy is required, but it enables the use of much more general ansatzes. Traditionally, these ansatzes included Slater-Jastrow wavefunctions and Slater-Jastrow-backflow wavefunctions [1].

* Corresponding author.

E-mail address: goldsh@berkeley.edu (G. Goldshlager).

<https://doi.org/10.1016/j.jcp.2024.113351>

Received 2 February 2024; Received in revised form 18 July 2024; Accepted 14 August 2024

Available online 20 August 2024

0021-9991/© 2024 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

In recent decades the modeling of very high-dimensional data and functions has seen impressive progress through the use of neural networks. For the quantum chemistry problem, the high-dimensional quantum wavefunction can be modeled by combining several neural network layers with a determinantal layer that enforces the Fermionic antisymmetry [3,4]. Such neural network wavefunctions, optimized using the variational Monte Carlo method, have enabled near-exact computation of the ground-state energy for small molecules [5–7] and certain condensed matter systems [8–12].

Due to the large number of parameters and the highly nonlinear nature of neural networks, the optimization of neural network wavefunctions poses a significant challenge. Prior to neural network wavefunctions, VMC simulations typically relied on powerful domain-specific optimizers such as stochastic reconfiguration (SR) [13–15] and the linear method [16–20]. These optimizers are able to converge in only tens or hundreds of iterations using highly accurate gradient estimates based on millions or more Monte Carlo samples. However, when applied to a wavefunction with N_p parameters, they require solving either a linear system or a generalized eigenvalue problem involving a dense $N_p \times N_p$ matrix known as the S matrix, with a cost that is $O(N_p^3)$ in general. The recently proposed Rayleigh-Gauss-Newton optimizer [21] presents a potentially favorable intermediate point between stochastic reconfiguration and the linear method, but still incurs a similar computational cost.

Since high-accuracy neural network architectures for molecules typically involve at least hundreds of thousands of parameters, it is not possible to directly apply SR or the linear method in this context. Instead, molecular neural network wavefunctions are typically optimized using methods from the machine learning community. These methods have a low per-iteration cost and require only a small number of Monte Carlo samples per iteration (or in machine learning parlance, a small minibatch size). However, they can require hundreds of thousands of optimization steps to converge fully.

For the FermiNet and related architectures [3,6,12,5], the most popular optimizer is the Kronecker-Factored Approximate Curvature method (KFAC) [22]. KFAC was originally designed as a tractable approximation to natural gradient descent (NGD) [23] for machine learning models with millions of parameters. In the quantum setting, natural gradient descent is equivalent to stochastic reconfiguration since the S matrix can be viewed as the Fisher information matrix of the normalized probability distribution defined by the wavefunction [3]. Taking the view of stochastic reconfiguration, the key ingredient of KFAC is an approximate factorization of the S matrix. This factorization enables the efficient inversion of the approximate S matrix and also makes it possible to approximate S based on multiple recent minibatches. This can significantly improve performance in highly stochastic settings where a single minibatch can provide only a very noisy approximation to the true S matrix. It is worth noting that the interpretation of KFAC as an approximate natural gradient method has recently been called into question due to some experiments which show that KFAC performs *better* than exact natural gradient updates [24]. Regardless of the underlying mechanism, KFAC represents the state-of-the-art in optimizing neural network wavefunctions for molecules and solids.

An alternative to KFAC is to take advantage of the fact that the estimate of the S matrix is always low-rank when it is based on only a small set of Monte Carlo samples. This idea was first introduced in the machine learning community in the form of efficient subsampled natural gradient descent [25]. In the physics community, several recent works have proposed similar methods for the VMC setting [26,27]. These methods make fewer heuristic assumptions than KFAC and are much simpler to describe and implement. However, unlike KFAC, they are limited to estimating the S matrix based on only a single minibatch at a time.

In this work, we improve upon these existing optimizers with a new method that we call Subsampled Projected-Increment Natural Gradient Descent (SPRING). Our method is most directly inspired by the minimum-step stochastic reconfiguration (MinSR) approach of Chen and Heyl [26]. MinSR is based on the observation that the SR parameter update can be formulated as the solution to an overdetermined linear least-squares problem. At each iteration, the Monte Carlo samples provide access to a small set of rows from this least-squares problem, yielding an underdetermined subproblem which has many possible solutions. To make the parameter update unique, Chen and Heyl propose to choose the minimal-norm solution to the sampled subproblem.

We improve upon this scheme by taking inspiration from the randomized Kaczmarz method for solving overdetermined linear least-squares problems [28]. The original randomized Kaczmarz method solves the problem by sampling a single row at a time. The subsequently developed block version of the method (see i.e. [29]) instead uses a small set of rows at each iteration. In either case, the key idea of the Kaczmarz method is to iteratively project the solution vector onto the hyperplane of solutions to each sampled subproblem. When the system is consistent and the rows are sampled from an appropriate probability distribution, this approach provably converges to the true solution with an expected error that decays by a constant factor at each iteration.

Recall that the SR parameter update can be formulated as the solution to an overdetermined linear least-squares problem, and that each minibatch of Monte Carlo samples provides access to a randomly sampled underdetermined subproblem. In direct analogy to the Kaczmarz method, the parameter update in SPRING is calculated by projecting the *previous* parameter update onto the hyperplane of solutions to the newly sampled subproblem. By leveraging the previous parameter update as the starting point for the projection, SPRING is able to use data from previous minibatches to obtain a more accurate approximation to the true SR update direction. Furthermore, this improvement is obtained at essentially no extra cost relative to MinSR.

To demonstrate the effectiveness of SPRING, we apply it to optimize a FermiNet wavefunction for several small atoms and molecules, namely the carbon, nitrogen, and oxygen atoms, two configurations of the N_2 molecule, and the CO molecule at equilibrium. We use the VMCNet code [30] for all numerical experiments. We find that SPRING consistently outperforms MinSR and KFAC when all three methods have their learning rates tuned for maximal performance. Remarkably, on the oxygen atom, SPRING is able to attain chemical accuracy after forty thousand training iterations, whereas both MinSR and KFAC fail to do so even after one hundred thousand iterations.

2. Background

2.1. Variational Monte Carlo

In this work we focus on the application of the variational Monte Carlo method to find ground-state wavefunctions for molecular systems. A molecule is defined by the positions R_I and charges Z_I of M atomic nuclei. Using the Born-Oppenheimer approximation [31], we model the nuclei as fixed point charges. The quantum system of interest then consists of N electrons interacting under the influence of these charges. The Hamiltonian for the system is given in atomic units by

$$H = -\frac{1}{2} \sum_{i=1}^N \Delta_{r_i} - \sum_{i=1}^N \sum_{I=1}^M \frac{Z_I}{|r_i - R_I|} + \sum_{i < j}^N \frac{1}{|r_i - r_j|} + \sum_{I < J}^M \frac{Z_I Z_J}{|R_I - R_J|}, \quad (1)$$

where $r_i \in \mathbb{R}^3$ represents the position of electron i and Δ_{r_i} represents the corresponding Laplacian operator. Because electrons are fermions, the wavefunction must be antisymmetric with respect to particle exchange. Within the space of antisymmetric wavefunctions, the ground-state energy E_0 and wavefunction $|\psi_0\rangle$ correspond to the smallest eigenvalue of H and its corresponding eigenfunction.

We make several simplifications to the space of wavefunctions before searching for the ground-state. First, since the Hamiltonian is Hermitian, its ground-state wavefunction can be chosen to be purely real. We thus limit our search to real normalizable many-body wavefunctions of the form $\psi : \mathbb{R}^{3N} \rightarrow \mathbb{R}$, usually written as $\psi(R) = \psi(r_1, \dots, r_N)$ with $R \in \mathbb{R}^{3N}$. Next, we fix the numbers N_\uparrow, N_\downarrow of up- and down-spin electrons *a priori* for each calculation. Given this assumption, together with the antisymmetry of the wavefunction and the spin-independence of the Hamiltonian, it is possible to assume without loss of generality that the first N_\uparrow electrons are always spin-up and the last N_\downarrow electrons are always spin-down [1].

The ground-state wavefunction can be found by minimizing the expectation value of the energy:

$$|\psi_0\rangle = \operatorname{argmin}_{\psi} \frac{\langle \psi | H | \psi \rangle}{\langle \psi | \psi \rangle}. \quad (2)$$

Given a variational ansatz $|\psi_\theta\rangle$, we can then define a loss function

$$L(\theta) = \frac{\langle \psi_\theta | H | \psi_\theta \rangle}{\langle \psi_\theta | \psi_\theta \rangle} \quad (3)$$

and represent the ground-state approximately via

$$\theta^* = \operatorname{argmin}_{\theta} L(\theta), \quad |\psi_0\rangle \approx |\psi_{\theta^*}\rangle, \quad E_0 \approx L(\theta^*). \quad (4)$$

In some electronic structure methods, such as the Hartree-Fock method and configuration interaction methods, the ansatz is chosen to enable the direct calculation of $L(\theta)$ based on precomputed one- and two-electron integrals. In variational Monte Carlo, the ansatz is chosen in a more flexible way, such as a Slater-Jastrow ansatz, a Slater-Jastrow-backflow ansatz, or a neural network wavefunction. For such ansatzes it is not possible to directly evaluate $L(\theta)$, so it is necessary to instead approximate it via Monte Carlo integration. To this end, $L(\theta)$ can be reformulated stochastically as

$$L(\theta) = \mathbb{E}_{R \sim p} [E_L(\theta)], \quad (5)$$

where $p(R) = \psi_\theta(R)^2 / \langle \psi_\theta | \psi_\theta \rangle$ and $E_L(R) = H\psi_\theta(R) / \psi_\theta(R)$ (see i.e. [1], Section III.C). This last term, $E_L(R)$, is known as the *local energy* of the wavefunction $|\psi_\theta\rangle$ at the position R .

To optimize a variational wavefunction using this formula, it is usually necessary to calculate the gradient $\nabla_\theta L(\theta)$. A convenient formula for this gradient can be derived taking advantage of the fact that H is Hermitian (see i.e. [32], Appendix E):

$$g = \nabla_\theta L(\theta) = 2\mathbb{E}_{R \sim p} [\nabla_\theta \log |\psi_\theta(R)| (E_L(R) - L(\theta))]. \quad (6)$$

In practice, both $L(\theta)$ and g are estimated stochastically using Markov-Chain Monte Carlo sampling to generate samples from $p(R)$. The parameters can then be updated by a variety of optimization schemes, several of which we discuss in the following sections.

2.2. Stochastic reconfiguration

One popular approach for optimizing variational wavefunctions is known as stochastic reconfiguration (SR) [13–15]. In recent years, SR has been widely adopted for optimizing neural quantum states in second quantization [33,34]. It also serves as the starting point for the MinSR method to be introduced later.

SR is based on the idea of imaginary time evolution, which states that the ground-state $|\psi_0\rangle$ can be found via the formula

$$|\psi_0\rangle = \lim_{\tau \rightarrow \infty} e^{-\tau H} |\psi\rangle \quad (7)$$

for any wavefunction $|\psi\rangle$ such that $\langle \psi_0 | \psi \rangle \neq 0$. In SR, we optimize a variational ansatz $|\psi_\theta\rangle$ by finding parameter updates that approximate small imaginary time steps. In particular, we choose a step-size $\delta\tau$ and use this to define the desired updated wavefunction

$|\psi'\rangle = e^{-\delta\tau H} |\psi_\theta\rangle$. Due to the constraints of the ansatz, we cannot update our wavefunction to $|\psi'\rangle$ directly. Instead we seek a parameter update $d\theta$ such that $|\psi_{\theta+d\theta}\rangle$ is close to $|\psi'\rangle$. To measure the closeness, we rely on the Fubini-Study distance defined by

$$D(|\psi\rangle, |\phi\rangle) = \arccos \frac{|\langle\psi|\phi\rangle|}{\|\psi\| \|\phi\|}. \quad (8)$$

In practice, we make a first order approximation to both $|\psi'\rangle$ and $|\psi_{\theta+d\theta}\rangle$ and then minimize a second-order approximation to $D^2(|\psi'\rangle, |\psi_{\theta+d\theta}\rangle)$. See [35] for a more thorough discussion of the geometric underpinnings of stochastic reconfiguration.

To present the formula for the update, we first define the functions $\bar{O}(R)$ and $\bar{\epsilon}(R)$. The first quantity, $\bar{O}(R)$, is the transposed gradient of the logarithm of the normalized wavefunction, evaluated at the point R . The second quantity, $\bar{\epsilon}(R)$, is the first-order change in the logarithm of the normalized wavefunction induced by the imaginary time-step $\delta\tau$, evaluated at the point R . These quantities are given by

$$\bar{O}(R) = (\nabla_\theta \log |\psi_\theta(R)| - \mathbb{E}_{R\sim p} [\nabla_\theta \log |\psi_\theta(R)|])^T, \quad (9)$$

$$\bar{\epsilon}(R) = -\delta\tau (E_L(R) - \mathbb{E}_{R\sim p} [E_L(R)]). \quad (10)$$

It can be shown that, to second order, the square of the Fubini-study distance is given by

$$D^2(|\psi'\rangle, |\psi_{\theta+d\theta}\rangle) \approx \mathbb{E}_{R\sim p} [|\bar{O}(R)d\theta - \bar{\epsilon}(R)|^2]. \quad (11)$$

This leads us to define the SR parameter update by

$$d\theta = \operatorname{argmin}_{d\theta'} \mathbb{E}_{R\sim p} [|\bar{O}(R)d\theta' - \bar{\epsilon}(R)|^2]. \quad (12)$$

A detailed derivation of this formulation of stochastic reconfiguration can be found in [26].

In practice, (11) must always be approximated using a finite collection of N_s samples R_1, \dots, R_{N_s} . To estimate $\bar{O}(R)$ and $\bar{\epsilon}(R)$ at the sampled points, we must empirically estimate the expectation values that appear in their definitions. We first define

$$O = \frac{1}{\sqrt{N_s}} \begin{bmatrix} \nabla_\theta \log |\psi_\theta(R_1)| \\ \vdots \\ \nabla_\theta \log |\psi_\theta(R_{N_s})| \end{bmatrix}, \quad \epsilon = -\frac{\delta\tau}{\sqrt{N_s}} \begin{bmatrix} E_L(R_1) \\ \vdots \\ E_L(R_{N_s}) \end{bmatrix}. \quad (13)$$

Next, let $\mathbf{1}$ be the column vector of length N_s whose entries are all equal to 1 and let $P = \frac{1}{N_s} \mathbf{1}\mathbf{1}^T$ represent the orthogonal projector onto the span of $\mathbf{1}$. When acting on the left, the operator P replaces each row of the matrix or vector on its right with the average of all of its rows. Our empirical estimates of $\bar{O}(R)$ and $\bar{\epsilon}(R)$ can thus be collected as

$$\bar{O} = (I - P)O, \quad \bar{\epsilon} = (I - P)\epsilon. \quad (14)$$

The estimate of the Fubini-Study distance of the update $d\theta$ can then be written as

$$D^2(|\psi'\rangle, |\psi_{\theta+d\theta}\rangle) \approx \|\bar{O}d\theta - \bar{\epsilon}\|^2, \quad (15)$$

and the update takes the form

$$d\theta = \operatorname{argmin}_{d\theta'} \|\bar{O}d\theta' - \bar{\epsilon}\|^2. \quad (16)$$

In traditional SR, many more samples are taken than parameters in the ansatz, meaning that the least-squares problem is overdetermined and should have a unique solution. To protect against the case when \bar{O} is ill-conditioned or singular, a Tikhonov regularization is generally added to yield the regularized problem

$$d\theta = \operatorname{argmin}_{d\theta'} \frac{1}{\lambda} \|\bar{O}d\theta' - \bar{\epsilon}\|^2 + \|d\theta'\|^2. \quad (17)$$

In the case of real parameters and a real wavefunction, the solution to (17) is given by

$$d\theta = (\bar{O}^T \bar{O} + \lambda I)^{-1} \bar{O}^T \bar{\epsilon}. \quad (18)$$

This is more traditionally written as

$$d\theta = -\frac{\delta\tau}{2} (S + \lambda I)^{-1} g, \quad (19)$$

where $S = \bar{O}^T \bar{O}$, $g = \nabla_\theta L(\theta) = -2\bar{O}^T \bar{\epsilon} / \delta\tau$, and the imaginary time step has been halved to cancel the factor of two in the gradient.

2.3. Minimum-step stochastic reconfiguration

For a problem with N_p variational parameters, the traditional SR update scales as $O(N_p^3)$ due to the need to invert the $N_p \times N_p$ matrix S . The insight of MinSR is that the parameter update can be calculated much more efficiently when $N_s \ll N_p$. In fact the linear system is highly underdetermined in this setting. The MinSR method addresses this issue by choosing the solution with minimal norm. For regularization, MinSR utilizes a pseudoinverse with a cutoff for small eigenvalues. Subsequently, Rende et al. [27] suggested using a Tikhonov regularization instead. This approach amounts to solving the same regularized equation as in traditional SR, namely (17). In the underdetermined setting, the application of the Sherman-Morrison-Woodbury formula to (18) yields the solution

$$d\theta = \bar{O}^T (\bar{O} \bar{O}^T + \lambda I)^{-1} \bar{\epsilon}. \quad (20)$$

This is the version of MinSR that we use as the basis for SPRING. The key point is that the update can now be calculated efficiently since $\bar{O} \bar{O}^T$ is only $N_s \times N_s$ rather than $N_p \times N_p$.

It is worth noting that methods very similar to MinSR were introduced independently in the machine learning community as early as 2019. For example, Zhang, Martens, and Grosse described a pseudoinverse-based approach to natural gradient descent for the overparameterized case when the number of parameters is greater than the total number of data-points available [36]. Subsequently, Ren and Goldfarb proposed an efficient subsampled natural gradient method [25] for the case when the minibatch size is smaller than the number of parameters. In fact, the Tikhonov-regularized form of MinSR can be viewed as a streamlined implementation of the method of Ren and Goldfarb which applies whenever the gradient of the loss function is a linear combination of the model gradients at the sampled points. We elaborate on this connection in Appendix A.

2.4. Kaczmarz method for solving linear systems

MinSR enables the efficient solution of an underdetermined subsample from the SR equation, but the solution of such an equation is not necessarily a good approximation to the solution of the original SR equation. Our goal is to develop an algorithm that can converge to the solution of the original SR equation using only a series of underdetermined subsamples. This is the motivation behind the Kaczmarz algorithm, and in particular its blocked variants.

To understand the Kaczmarz algorithm, consider an overdetermined system of linear equations $Ax = b$, where A is $m \times n$ and

$$A = \begin{bmatrix} a_1^T \\ \vdots \\ a_m^T \end{bmatrix}, b = \begin{bmatrix} b_1 \\ \vdots \\ b_m \end{bmatrix}. \quad (21)$$

Assume that A has full column rank and that the system is consistent, meaning that there exists a unique x^* such that $Ax^* = b$. The earliest form of the Kaczmarz algorithm, which dates back to the first half of the twentieth century [37], solves this problem iteratively. A starting guess x_0 is required for initialization and is then honed by iterations of the form

$$x_k = x_{k-1} + \frac{a_i}{\|a_i\|^2} (b_i - a_i^T x_{k-1}), \quad (22)$$

where i is chosen to cycle through all the rows of A one at a time. The interpretation is that x_k is attained by projecting x_{k-1} onto the solution space of the sampled equation $a_i^T x = b_i$.

The seminal work of Strohmer and Vershynin [28] proposed a randomized variant of the algorithm which samples row i with a probability proportional to $\|a_i\|^2$. With this randomized sampling strategy, they showed that the algorithm converges as

$$\mathbb{E} \left[\|x_k - x^*\|^2 \right] \leq (1 - \kappa_D(A)^{-2})^k \|x_0 - x^*\|^2, \quad (23)$$

where $\kappa_D(A) = \|A\|_F \|A^+\|_2$ is the Demmel condition number of A and A^+ is the Moore-Penrose pseudoinverse of A .

Since single row updates are very inefficient on modern hardware, blocked versions of the randomized Kaczmarz algorithm were subsequently proposed and analyzed, i.e. [29]. Such randomized block Kaczmarz algorithms take the form

$$x_k = x_{k-1} + A_\sigma^+ (b_\sigma - A_\sigma x_{k-1}), \quad (24)$$

where σ represents a randomly selected subset of the rows of A and b and A_σ^+ represents the pseudoinverse of A_σ . The interpretation is similar to before: x_k represents the projection of x_{k-1} onto the solution space of the sampled equation $A_\sigma x = b_\sigma$. For consistent equations the randomized block Kaczmarz algorithm converges linearly to the true solution x^* , though the convergence rate is more complicated than in the single row case and has a more nuanced dependence on the sampling procedure.

3. Methods

We now present our main contribution, the Subsampled Projected-Increment Natural Gradient Descent (SPRING) algorithm for optimizing neural network wavefunctions. To start, recall the formula for the SR parameter update:

$$d\theta = \underset{d\theta'}{\operatorname{argmin}} \mathbb{E}_{R \sim p} \left[\left| \bar{O}(R) d\theta' - \bar{\epsilon}(R) \right|^2 \right]. \quad (25)$$

Since R is drawn from a continuous and high-dimensional space this equation can be seen as a highly overdetermined linear least-squares problem, with each configuration R corresponding to a single row. Now, let \bar{O}_k and $\bar{\epsilon}_k$ denote the subsamples from $\bar{O}(R)$ and $\bar{\epsilon}(R)$ that are available at VMC training iteration k . Our approach hinges on the assumption that the imaginary time-step is small, meaning the parameter vector θ changes only slightly at each iteration. Thus, although the sampled subproblems $(\bar{O}_{k-1}, \bar{\epsilon}_{k-1})$, $(\bar{O}_k, \bar{\epsilon}_k)$ can be completely different, the underlying SR equations are nearly identical. This inspires us to proceed as follows:

1. Make the approximation that the recent iterates $(\bar{O}_k, \bar{\epsilon}_k), (\bar{O}_{k-1}, \bar{\epsilon}_{k-1}), \dots$ are all random block samples from the current SR equation

$$\operatorname{argmin}_{d\theta'} \mathbb{E}_{R \sim p} \left[|\bar{O}(R)d\theta' - \bar{\epsilon}(R)|^2 \right]. \quad (26)$$

2. Apply the randomized block Kaczmarz method to these iterates to yield an approximate solution ϕ_k to the current SR equation.
3. Apply some scaling and clipping procedures, to be described later, to ϕ_k to calculate the parameter update $d\theta_k$.

In this way SPRING takes advantage of the entire optimization history to inform each parameter update, distinguishing it from MinSR which leverages only a single minibatch of data at a time.

Following the structure of a single step of the randomized block Kaczmarz method [29], we calculate ϕ_k by projecting ϕ_{k-1} onto the solution space of the underdetermined equation $\bar{O}_k \phi = \bar{\epsilon}_k$. Recall that the basic form for the Kaczmarz projection is

$$\phi_k = \phi_{k-1} + \bar{O}_k^+ (\bar{\epsilon}_k - \bar{O}_k \phi_{k-1}). \quad (27)$$

Equivalently we can say that ϕ_k satisfies

$$\phi_k = \operatorname{argmin}_{\phi} \|\phi - \phi_{k-1}\|^2 \text{ s.t. } \bar{O}_k \phi = \bar{\epsilon}_k. \quad (28)$$

Thus, at a conceptual level, our algorithm differs from MinSR only in how we break the indeterminacy of the sampled subproblem: MinSR chooses the solution of minimal norm, whereas SPRING chooses the solution that is nearest to the previous approximate solution.

In practice \bar{O}_k can be singular or very ill-conditioned, so it is beneficial add some form of regularization to this projection. We choose to regularize the projection by incorporating the linear equations using a penalty term rather than a hard constraint. To make this modification, we introduce a regularization parameter λ and use the formula

$$\phi_k = \operatorname{argmin}_{\phi} \frac{1}{\lambda} \|\bar{O}_k \phi - \bar{\epsilon}_k\|^2 + \|\phi - \phi_{k-1}\|^2. \quad (29)$$

Even with this regularization, we find that directly using (29) can result in unstable optimization trajectories (see Fig. 10b). We do not yet have an explanation for this phenomenon. To stabilize the method, we decay the previous gradient by a small amount before projecting it. Formally, this stabilization comes in the form of a new regularization parameter μ and a modified update formula

$$\phi_k = \operatorname{argmin}_{\phi} \frac{1}{\lambda} \|\bar{O}_k \phi - \bar{\epsilon}_k\|^2 + \|\phi - \mu \phi_{k-1}\|^2. \quad (30)$$

An explicit formula for ϕ_k can now be derived. We first define $\pi_k = \phi_k - \mu \phi_{k-1}$, $\pi = \phi - \mu \phi_{k-1}$, and $\bar{\zeta}_k = \bar{\epsilon}_k - \mu \bar{O}_k \phi_{k-1}$ and note that (30) can be recast as

$$\pi_k = \operatorname{argmin}_{\pi} \frac{1}{\lambda} \|\bar{O}_k \pi - \bar{\zeta}_k\|^2 + \|\pi\|^2. \quad (31)$$

This formula now has the same form as MinSR with Tikhonov regularization, so we know that the solution is given by

$$\pi_k = \bar{O}_k^T (\bar{O}_k \bar{O}_k^T + \lambda I)^{-1} \bar{\zeta}_k. \quad (32)$$

Finally, we can translate back to ϕ_k yielding

$$\phi_k = \bar{O}_k^T (\bar{O}_k \bar{O}_k^T + \lambda I)^{-1} \bar{\zeta}_k + \mu \phi_{k-1}. \quad (33)$$

To understand this formula, it is helpful to explicitly write out the three terms:

$$\phi_k = \bar{O}_k^T (\bar{O}_k \bar{O}_k^T + \lambda I)^{-1} \bar{\epsilon}_k + \mu \phi_{k-1} - \mu \bar{O}_k^T (\bar{O}_k \bar{O}_k^T + \lambda I)^{-1} \bar{O}_k \phi_{k-1}. \quad (34)$$

The first term alone corresponds to MinSR, while the first two terms together correspond to a scaled version of MinSR with momentum (to be presented in Section 3.4). It is the addition of the third term that concretely distinguishes SPRING from these other methods.

Equation (33) is almost the formula that we use in practice. The only missing piece is an extra stabilization procedure that we introduce in Section 3.2, which is not essential but should be included for best performance. Finally, to calculate $d\theta_k$, we also include a learning rate schedule η_k and a norm constraint C which ensures that the parameters are not changed too much at each iteration (see Section 3.1). To ensure that the Kaczmarz algorithm is implemented consistently, regardless of the learning rate schedule, we

always calculate $\bar{\epsilon}_k$ using $\delta\tau = 1$ and then apply the learning rate η_k only when we calculate $d\theta_k$. For convenience we present the full procedure for SPRING including these extra details in Algorithm 1.

Algorithm 1 SPRING.

Require: Hamiltonian H , ansatz ψ_θ ,

Require: Initialization θ_0 , iteration count K , batch size N_s , learning rate schedule η_k

Require: Tikhonov damping λ , decay factor μ , norm constraint C

```

1:  $\theta \leftarrow \theta_0$ 
2:  $\phi \leftarrow 0$ 
3: for k in 1:K do
4:   Sample  $R_1, \dots, R_{N_s}$  from  $p(R) = \psi_\theta(R)^2 / \langle \psi_\theta | \psi_\theta \rangle$  ▷ i.e. using MCMC
5:   Calculate  $\bar{O}$ ,  $\bar{\epsilon}$  using  $\delta\tau = 1$  ▷ via (13), (14)
6:    $\bar{\zeta} \leftarrow \bar{\epsilon} - \mu \bar{O} \phi$ 
7:    $\phi \leftarrow \bar{O}^T (\bar{O} \bar{O}^T + \lambda I + \frac{1}{N_s} \mathbf{1}\mathbf{1}^T)^{-1} \bar{\zeta} + \mu \phi$  ▷ following (39) and using Cholesky
8:    $d\theta \leftarrow \phi \cdot \min(\eta_k, \sqrt{C} / \|\phi\|)$  ▷ See Section 3.1
9:    $\theta \leftarrow \theta + d\theta$ 
10: end for
11: return  $\theta$ 

```

It is worth noting that due to several technical details, we cannot directly transfer the convergence results for the classical Kaczmarz method to the VMC setting. First of all, the inclusion of the parameter μ in SPRING distinguishes it substantially from the Kaczmarz method. We can understand this distinction as arising from the fact that in SPRING, the underlying least-squares system changes a small amount after each iteration. Thus, it makes sense to progressively forget the information from previous iterations, which is what we accomplish by setting $\mu < 1$. Additionally, the least-squares problem arising in VMC need not be consistent, and in the inconsistent setting the Kaczmarz method only converges to within a ball of the true solution, with the radius of the ball depending on how inconsistent the system is [38]. For these reasons, we cannot derive a rigorous convergence guarantee for SPRING with regard to either $d\theta$ or θ . Still, the connection to the Kaczmarz method serves to motivate the algorithm and explain its superior performance relative to MinSR.

3.1. Norm constraint

As an additional stabilization procedure, we include a constraint on the norm of the MinSR, MinSR+M, and SPRING updates inspired by the norm constraint used by KFAC. We have found that including this norm constraint helps to stabilize the optimization procedure and enable larger learning rates to be applied, and it also improves the ability to transfer learning rates from one system to another. For reference we briefly describe how the norm constraint used by KFAC works before introducing the form of the norm constraint that we use for MinSR and SPRING.

The norm constraint in KFAC is “natural” in the sense that it applies to the norm induced by the Fisher information matrix as opposed to the Euclidean norm. To define this scheme let g be the estimated energy gradient, η be the learning rate, S be KFAC’s approximation to the Fisher information matrix, and $\phi = S^{-1}g$ be the unscaled and unconstrained KFAC parameter update. The final update will be $d\theta = c\phi$ for some scalar c . The norm constraint relies upon the fact that $\|d\theta\|_S^2 = c^2\phi^T S \phi = c^2\phi^T g$ and as a result it enforces the constraint

$$c^2\phi^T g \leq C. \quad (35)$$

Concretely this is achieved by choosing

$$d\theta = \phi \cdot \min(\eta, \sqrt{C} / \sqrt{\phi^T g}). \quad (36)$$

In SPRING the update $d\theta$ and the gradient g are not related by a positive definite matrix S ; hence the above scheme is not justified. Instead we implement a simple Euclidean norm constraint given by

$$\|d\theta\|^2 \leq C. \quad (37)$$

Just like in KFAC, if this constraint is violated before the norm constraint is applied, we scale the update down to ensure that it is satisfied. Concretely, we calculate the update $d\theta_k$ from the Kaczmarz iterate ϕ_k and the learning rate η_k using the formula

$$d\theta_k = \phi_k \cdot \min(\eta_k, \sqrt{C} / \|\phi_k\|).$$

3.2. Numerically stable inversion

In exact arithmetic, $T = \bar{O}\bar{O}^T$ is positive semidefinite, $T + \lambda I$ is positive definite, and $(T + \lambda I)^{-1}\bar{\zeta}$ can be computed efficiently using the Cholesky decomposition. However, in single precision we have found that the T matrix can become indefinite. As a result, the Cholesky decomposition can fail if the regularization is chosen to be too small, and we have even observed this to occur when using our default value of $\lambda = 0.001$.

We now note that T is *always* singular due to the construction of \bar{O} . Recall that $\mathbf{1}$ represents the column vector whose entries are all equal to 1 and $P = \frac{1}{N_s} \mathbf{1}\mathbf{1}^T$ represents the orthogonal projector onto the span of $\mathbf{1}$. Since $\bar{O} = (I - P)O$, we have that

$$T\mathbf{1} = \bar{O}\bar{O}^T\mathbf{1} = \bar{O}O^T(I - P)\mathbf{1} = \bar{O}O^T(\mathbf{1} - \mathbf{1}) = 0. \quad (38)$$

This means that any numerical perturbation to T can make it indefinite. If the numerical perturbation is large enough, the resulting T matrix can have an eigenvalue smaller than $-\lambda$, in which case the Cholesky factorization fails. This can happen even when all of the other eigenvalues of T are well-separated from 0.

We alleviate this issue by replacing the matrix $T + \lambda I$ with the further regularized matrix $T + \lambda I + \omega P$ for some positive real ω . We thus calculate the Kaczmarz iterate for iteration k as

$$\phi_k = \bar{O}_k^T(T_k + \lambda I + \omega P)^{-1}\bar{\zeta}_k + \mu\phi_{k-1}. \quad (39)$$

This does not affect the value of ϕ_k in exact arithmetic. To see this, note that $\mathbf{1}$ is an eigenvector of $(T_k + \lambda I)$ with eigenvalue λ , and is thus also an eigenvector of $(T_k + \lambda I + \omega P)$ with eigenvalue $\omega + \lambda$. As a result, we have

$$(T_k + \lambda I + \omega P)^{-1} = (T_k + \lambda I)^{-1} + \left(\frac{1}{\omega + \lambda} - \frac{1}{\lambda}\right)P. \quad (40)$$

Acting on the left with \bar{O}_k^T then annihilates the extra term:

$$\begin{aligned} \bar{O}_k^T(T_k + \lambda I + \omega P)^{-1} &= \bar{O}_k^T(T_k + \lambda I)^{-1} + \left(\frac{1}{\omega + \lambda} - \frac{1}{\lambda}\right)\bar{O}_k^T P \\ &= \bar{O}_k^T(T_k + \lambda I)^{-1}. \end{aligned}$$

Thus, using (39) reduces the likelihood that numerical errors cause the Cholesky solver to fail without otherwise affecting the computation. In our experiments we have used the value $\omega = 1$ in all cases. However, it is theoretically possible for the eigenvalues of T to all be smaller than 1, in which case setting $\omega = 1$ would increase the condition number of the regularized T matrix. In such a case ω could instead be chosen for example to be the mean of the eigenvalues of the original T matrix.

3.3. MinSR implementation details

We now briefly discuss the version of MinSR that we use for our numerical experiments, which is very similar but not identical to the basic version presented in Section 2.3. Functionally, the only differences are the inclusion of a norm constraint and the extra regularization of the matrix $\bar{O}_k\bar{O}_k^T$ as described in Section 3.2. Additionally, for consistency with SPRING we set $\delta\tau = 1$ when calculating $\bar{\epsilon}_k$ and we instead control the step size with a learning rate parameter η_k . With these modifications the MinSR algorithm is described by the following formulas:

$$\phi_k = \bar{O}_k^T(\bar{O}_k\bar{O}_k^T + \lambda I + P)^{-1}\bar{\epsilon}_k, \quad (41)$$

$$d\theta_k = \phi_k \cdot \min(\eta_k, \sqrt{C}/\|\phi_k\|). \quad (42)$$

3.4. MinSR with momentum

We motivate SPRING primarily via its relation to the Kaczmarz method. An alternative viewpoint is to see SPRING as a modification of a momentum method with an added projection step to remove the component of the momentum that conflicts with the current subproblem $\bar{O}_k d\theta = \bar{\epsilon}$. While we do not endorse this perspective, it is natural to wonder how SPRING would compare to MinSR with naive momentum, which we refer to as MinSR+M.

Concretely, MinSR+M iteratively updates a momentum direction ϕ_k by taking a convex combination of the previous value ϕ_{k-1} and the current MinSR solution. Just like in SPRING, we assume an imaginary time step of $\delta\tau = 1$ when we calculate $\bar{\epsilon}_k$ and instead control the step size using a learning rate parameter η_k and a norm constraint. The method is thus described by the formulas:

$$\phi_k = (1 - \mu)\bar{O}_k^T(\bar{O}_k\bar{O}_k^T + \lambda I + P)^{-1}\bar{\epsilon}_k + \mu\phi_{k-1}, \quad (43)$$

$$d\theta_k = \phi_k \cdot \min(\eta_k, \sqrt{C}/\|\phi_k\|). \quad (44)$$

As in SPRING, the μ parameter acts as a decay factor on the previous parameter update. However, MinSR+M differs significantly from SPRING in that no regularized projector is applied to the previous update. It also differs in that the MinSR solution is scaled by a factor of $(1 - \mu)$ to achieve a convex combination, which means μ must be chosen to be strictly less than one. See Section 4.3 and Fig. 10a for further discussion of the choice of μ for MinSR+M. We demonstrate in our numerical experiments that SPRING outperforms MinSR+M substantially. This supports the viewpoint that SPRING is not simply a momentum method, but is rather a way of improving the quality of the parameter updates as approximate solutions to the SR equation.

4. Results

We now turn to our numerical results. Our main finding is that SPRING outperforms KFAC, MinSR, and MinSR+M for several small atoms and molecules. We present results for the carbon, nitrogen, and oxygen atoms in Section 4.1 and for the N_2 and CO

Table 1

Summary of final energies attained with tuned learning rates for all systems studied. Energies are reported in Hartrees, to four decimal places.

System	KFAC	MinSR	MinSR+M	SPRING	Benchmark	Benchmark source
C	-37.8445	-37.8445	-37.8448	-37.8449	-37.8450	[40], Table XI
N	-54.5877	-54.5885	-54.5889	-54.5890	-54.5892	[40], Table XI
O	-75.0643	-75.0652	-75.0659	-75.0668	-75.0673	[40], Table XI
N ₂ , equilibrium	-109.5301	-109.5268	-108.5294	-109.5322	-109.5423	[41], Table II
N ₂ , 4.0 Bohr	-109.1862	-109.1794	-109.1829	-109.1906	-109.2021	[42], MLR ₄ (6, 8)
CO	-113.3140	-113.3092	-113.3117	-113.3169	-113.3255	[3], Table II

molecules in Section 4.2. We summarize the final energies attained in our main experiments in Table 1. We use VMCNet [30] to run all our experiments, which is a neural network VMC framework based on JAX [39]. For all of our experiments, we use a standard FermiNet architecture with 16 dense determinants; see Table B.2 for details.

For our training phase we use one hundred thousand optimization iterations with one thousand MCMC samples per iteration. For our MCMC proposals we use Gaussian all-electrons moves with a step-size that is tuned dynamically to maintain an acceptance ratio of approximately 50%. We take ten MCMC steps between optimization iterations to reduce correlations between the samples, and we use mean-centered local energy clipping to improve the stability of the optimization. We use a decaying learning rate with a decay rate of $r = 1e - 4$ and learning rate schedule $\eta_k = \frac{\eta}{1+r^k}$. For MinSR, MinSR+M, and SPRING, we assume $\delta\tau = 1$ and apply the learning rate η_k as described in Section 3. For more details on our VMC settings, see Table B.3. To obtain final energies we run a separate inference phase with the parameters from the last iteration of each optimization run; see Table B.4 for details.

To compare the methods fairly, we do not start them from fully random initializations. Rather, for each system in question, we first run a short preliminary optimization and save the result to a checkpoint. The rest of our experiments then load the starting parameters from the end of the preliminary optimization phase, switch out the optimizer and hyperparameters to the desired settings, and run one hundred thousand further optimization iterations. This procedure removes the chaotic early optimization stage as a factor from the comparison between the different methods, ensuring that no method is unfairly advantaged by randomly having a favorable start. Our preliminary optimization is different from the Hartree-Fock pretraining procedure (see [3,6] for example) because we optimize the variational energy from the beginning. However, our results should translate well to a setting with pretraining as a result of this preliminary optimization. See Table B.5 for more details on the preliminary optimization phase.

For all of our optimizers we use a Tikhonov regularization parameter of $\lambda = 0.001$ and a norm constraint of $C = 0.001$ unless indicated otherwise. For MinSR+M we use a momentum parameter of $\mu = 0.9$, and for SPRING we use the regularization parameter $\mu = 0.99$. In Section 4.3 we present several hyperparameter studies which demonstrate that these values provide a fair basis for comparing the methods. We also investigate the effect of the preliminary optimization phase and demonstrate that SPRING outperforms KFAC significantly without preliminary optimization.

In all of our numerical results we report energy errors relative to the benchmark energies listed in Table 1. For the energy and variance trajectories, we smooth out the curves by reporting averages over a sliding window of ten thousand iterations. We note that we are using a somewhat smaller network, fewer MCMC samples, and fewer optimization iterations than were used in most state-of-the-art calculations with KFAC; hence it should not be surprising that some of our results do not reach the same level of accuracy as those presented in the literature for comparable systems. Our results can be viewed as proof-of-principle for the value of SPRING on small calculations, which will need to be extended to larger calculations by future works.

4.1. Small atoms

We first present results for KFAC, MinSR, MinSR+M, and SPRING on the carbon, nitrogen, and oxygen atoms. We find that it is critical to tune the learning rate of each method in order to achieve the best performance and obtain a fair comparison. Thus, for each method, we test a number of learning rates on the carbon atom, with results presented in Fig. 1. Our experiments lead us to choose a learning rate of $\eta = 0.02$ for KFAC, $\eta = 0.1$ for MinSR, $\eta = 0.2$ for MinSR+M, and $\eta = 0.02$ for SPRING.

After tuning the learning rates on the carbon atom, we use the selected learning rates to compare the methods on carbon, nitrogen, and oxygen. The results with optimized learning rates are shown in Fig. 2. On the carbon atom, we find that SPRING reaches chemical accuracy significantly faster than the other methods. Relative to KFAC and MinSR, its final energy error is about a factor of four lower, and its final local energy variance that is about an order of magnitude lower. On nitrogen and oxygen, SPRING continues to outperform the other methods. In the case of the oxygen atom, SPRING reaches chemical accuracy after forty thousand iterations, whereas MinSR+M takes eighty thousand iterations and the other methods never reach chemical accuracy.

4.2. Molecules

We next test KFAC, MinSR, MinSR+M, and SPRING on two somewhat larger systems: the N₂ and CO molecules. We begin by tuning the learning rates for all four methods on the N₂ molecule at an equilibrium bond distance of 2.016 Bohr. We present the results of the learning rate sweeps in Fig. 3. Based on these results, we pick learning rates of $\eta = 0.05$ for KFAC, $\eta = 0.02$ for MinSR, $\eta = 0.02$ for MinSR+M, and $\eta = 0.002$ for SPRING. These learning rates are not always chosen strictly for the best final energy since the

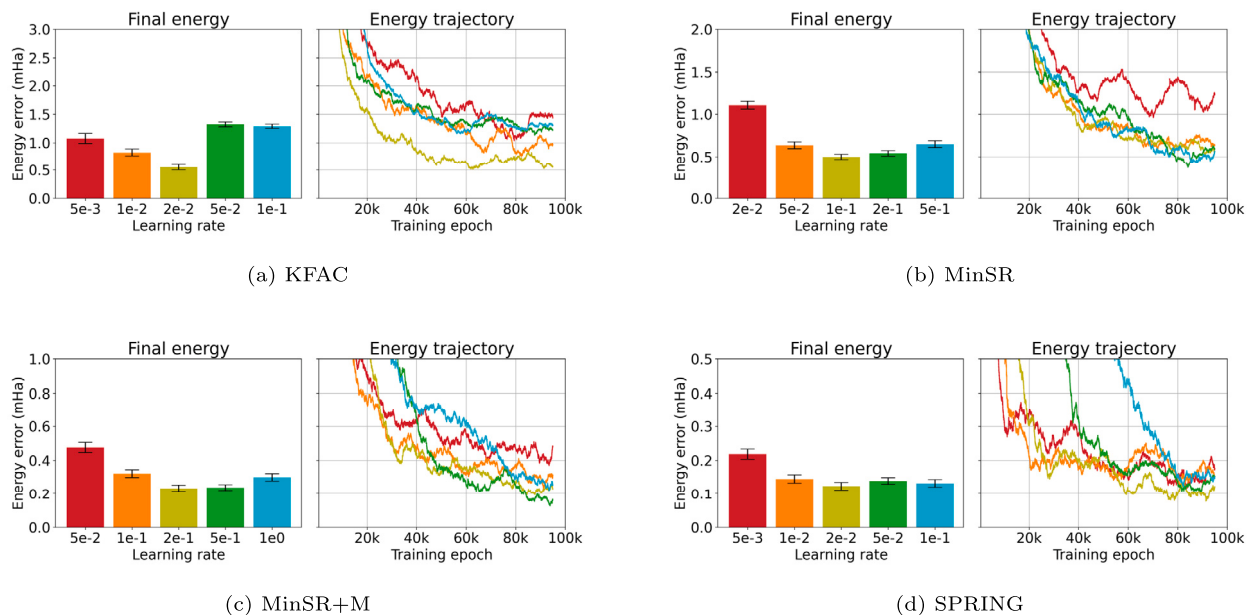


Fig. 1. Learning rate sweeps on the carbon atom with four different optimizers.

final energy can vary due to statistical fluctuations in the parameters and statistical errors in the energy estimation. We instead look at the optimization trajectories and the final energies and pick the learning rate that appears to be optimal based on the combination of the two. It is worth noting that the optimal learning rates for MinSR+M and SPRING change by a factor of 10 between the carbon atom and the N_2 molecule, whereas the optimal learning rate for MinSR changes by only a factor of 5, and the optimal learning rate for K FAC changes by only a factor of 2. The optimal learning rate may thus be more system-dependent for MinSR+M and SPRING compared to the other methods.

Using these optimized learning rates, we then compare the four methods on the N_2 molecule both at equilibrium and at a stretched bond distance of 4.0 Bohr, and on the CO molecule at an equilibrium bond distance of 2.173 Bohr. The results of these comparisons are presented in Fig. 4 and Fig. 5. For these systems none of the optimizers are able to reach chemical accuracy, which is expected due to the relatively small network, few MCMC samples, and few optimization iterations that we are using. Nonetheless, SPRING still displays a marked advantage over the other methods, converging much faster and to a lower energy and energy variance in all three cases. The MinSR+M optimizer only improves upon MinSR marginally for these systems, and it is outperformed by K FAC. This provides evidence that the SPRING algorithm is distinctly better than simply adding momentum on top of MinSR.

4.3. Hyperparameter studies

In this section, we provide several experiments that demonstrate that our main results are robust to the choices of C and μ . We also provide evidence that the results are not significantly skewed by the preliminary optimization phase. For simplicity we focus on the carbon atom for the purposes of this section.

We first show in Fig. 6 how the results of the learning rate experiments differ if we do not apply a norm constraint to any of the methods. In all four cases, the performance is much more sensitive to the learning rate when the constraint is turned off, and we are prevented from displaying results with larger learning rates because they produce completely unstable optimization trajectories. The combination of these factors means that without the norm constraints, it is much more difficult to find an effective learning rate, and our ability to transfer the learning rate from system to system is hindered. Nonetheless, even with the norm constraint off, the tuned version of SPRING significantly outperforms the tuned versions of the other methods, as shown in Fig. 7. In Fig. 8, we show that neither increasing nor decreasing the norm constraint used for K FAC enables it to compete with the performance of SPRING on the carbon atom. This provides further evidence that the advantage of SPRING is not related to the choice or the form of the norm constraint.

To better understand the impact of the norm constraint on SPRING, we return to the learning rate sweep for SPRING on the carbon atom and provide some additional data from these five optimization runs. In particular, we plot the scale factor $q_k = \max(1, \|\eta_k \phi_k\| / \sqrt{C})$. For example, if $q_k = 5$ then this means that at iteration k the norm constraint resulted in a scaling down of the parameter update by a factor of 5, and if $q_k = 1$ then the norm constraint at iteration k had no effect. We see in Fig. 9 that q_k tends to grow over the course of the first one hundred optimization iterations and then decay to 1 as the optimization progresses. We can understand this in the following way: q_k first rises as more optimization history is incorporated into SPRING; it then decays due to both the learning rate decay and the decay of the energy gradient. As should be expected, with a larger learning rate, q_k reaches a larger peak and takes longer to decay to 1.

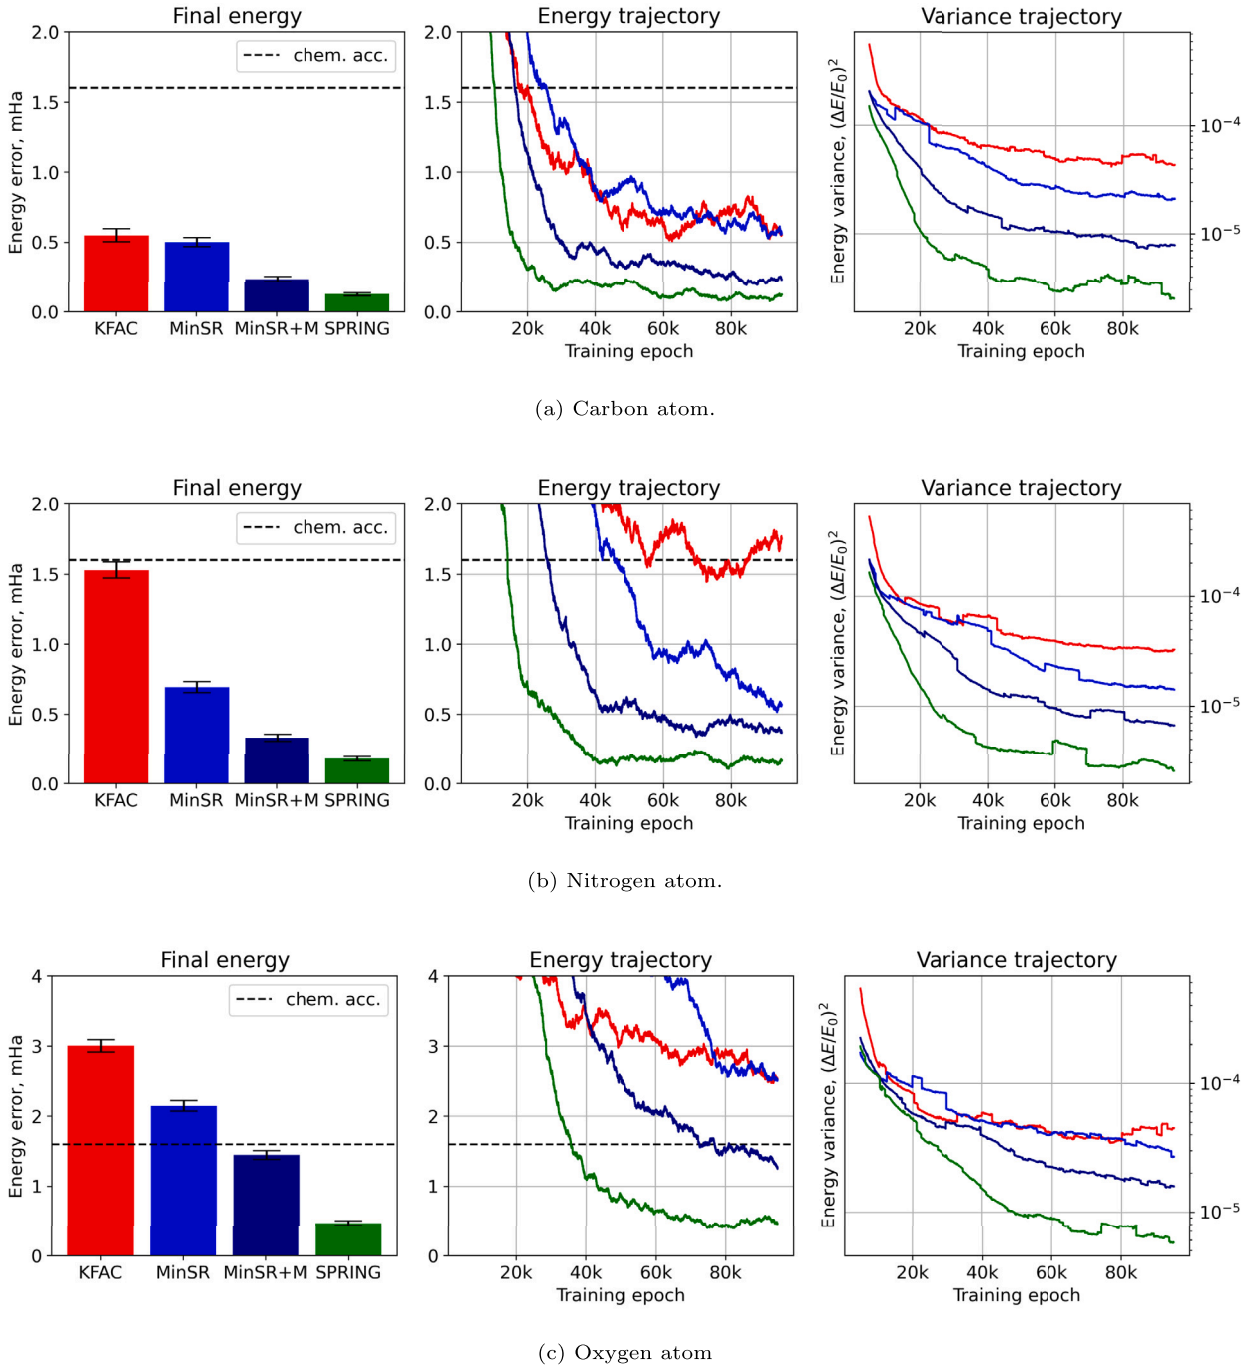


Fig. 2. Comparison of methods on three small atoms, with learning rates tuned on the carbon atom.

Next, we test the MinSR+M scheme with several different values of the momentum parameter μ , using the previously optimized learning rate of $\eta = 0.2$. For values larger than $\mu = 0.95$ we find that the optimization is unstable and we do not report numerical results. We find that choosing μ anywhere between 0.8 and 0.95 produces approximately optimal results, which justifies our use of $\mu = 0.9$ for our main experiments. Results are shown in Fig. 10a. Similarly, we test SPRING with several values of its regularization parameter μ , again using the previously optimized learning rate of $\eta = 0.02$. Results are shown in Fig. 10b. We find that choosing μ anywhere between 0.98 and 0.999 produces approximately optimal results, which justifies our choice of $\mu = 0.99$ for our main experiments. Interestingly, if $\mu = 1.0$ then the method becomes unstable. We do not yet have an explanation for the source of this instability.

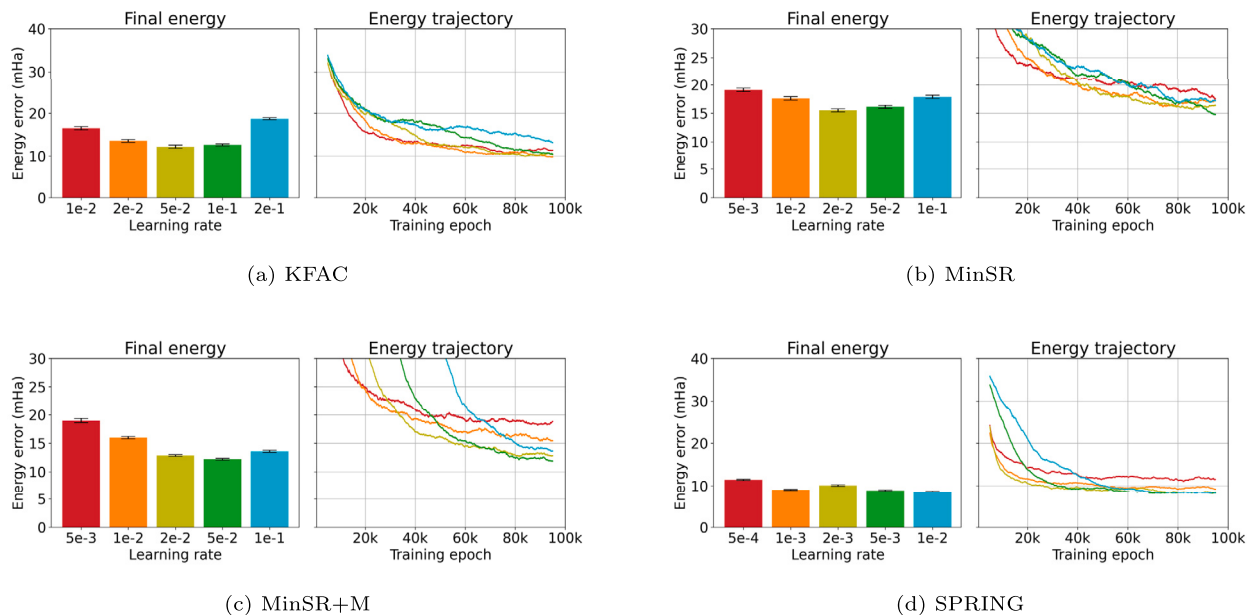




Fig. 3. Learning rate sweeps on the nitrogen molecule at equilibrium bond distance with four different optimizers.



(a) Equilibrium configuration, bond distance 2.016 Bohr.



(b) Stretched configuration, bond distance 4.0 Bohr.

Fig. 4. Comparison of methods on N_2 molecule at two bond distances, with learning rates tuned at equilibrium.


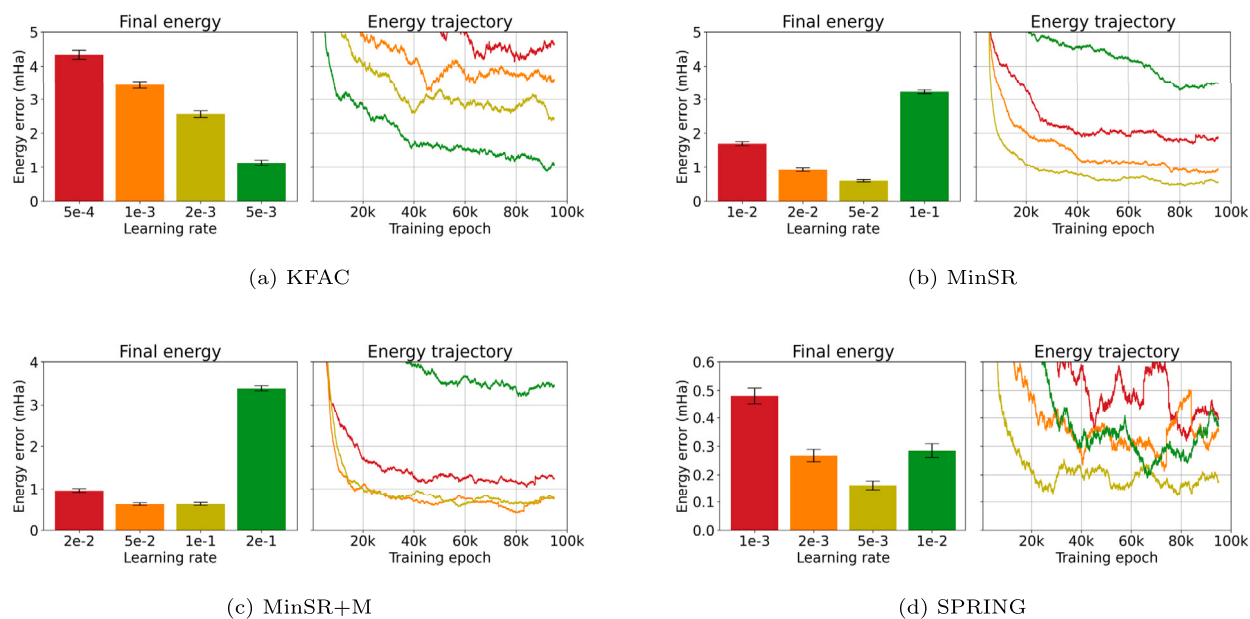
Fig. 5. Comparison of methods on CO molecule, with learning rates tuned on N₂.

Fig. 6. Learning rate sweeps on the carbon atom without norm constraints with four different optimizers.




Fig. 7. Comparison of methods on carbon atom, with tuned learning rates and no norm constraint.

Finally, we consider the effects of the preliminary optimization phase. Since we have used KFAC for the preliminary optimization in all cases, it is unclear thus far whether SPRING can perform well when run from a random initialization. To ameliorate this concern, we compare the performance of SPRING against KFAC without any preliminary optimization, with results in Fig. 11. We find that

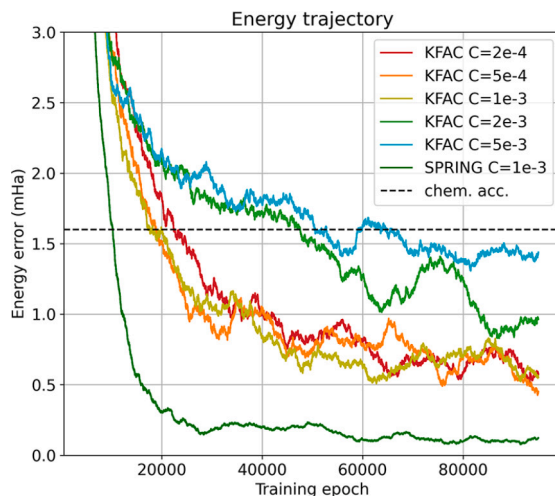


Fig. 8. Testing KFac on the carbon atom with different values of the norm constraint C . No value performs better than the default value of $C = 0.001$ and all settings are outperformed by SPRING with $C = 0.001$. All runs use the optimized learning rate of $\eta = 0.02$.

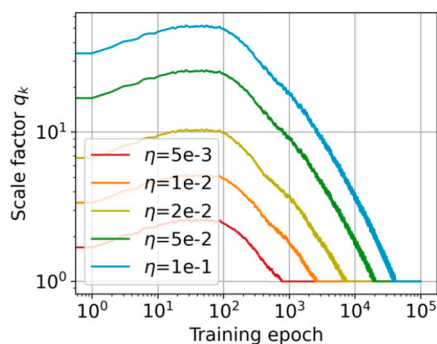


Fig. 9. Effect of the norm constraint on the SPRING parameter update during the learning rate sweep on the carbon atom. The quantity $q_k = \max(1, \|\eta_k \phi_k\| / \sqrt{C})$ represents the extent to which the parameter update is scaled down as a result of the norm constraint.

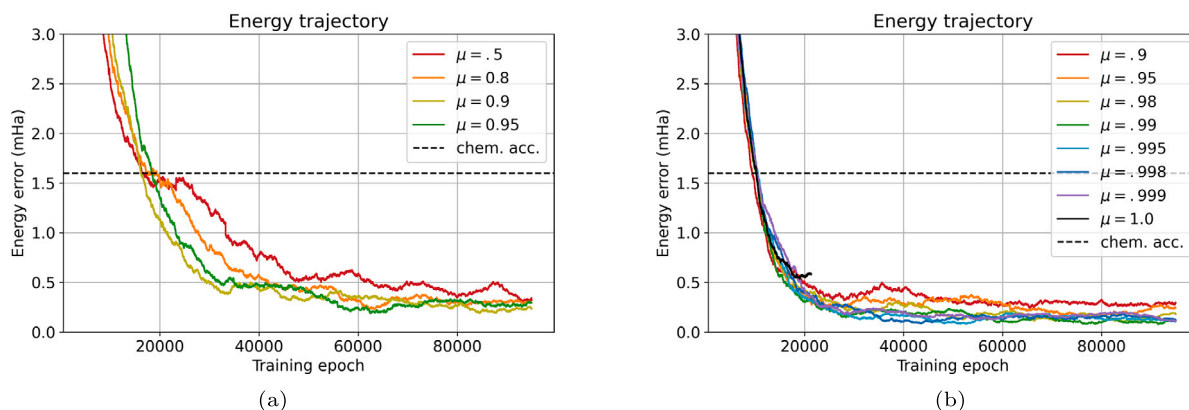


Fig. 10. Hyperparameter studies for the parameter μ for both MinSR+M and SPRING. (a) MinSR+M with several values of the momentum parameter μ , on the carbon atom. The learning rate is held fixed at the previously tuned value of $\eta = 0.2$. (b) SPRING on the carbon atom with several values of the regularization parameter μ . All values converge well except for the unregularized case $\mu = 1.0$, shown in black, for which the optimization is unstable and encounters NaNs after approximately twenty thousand epochs. The learning rate is held fixed at the previously tuned value of $\eta = 0.02$.

KFac outperforms SPRING in the very early stages of the optimization. However, SPRING overtakes KFac after about seven thousand training iterations and reaches chemical accuracy in fewer than half as many iterations as KFac. We conclude that SPRING is also effective starting from a random initialization.

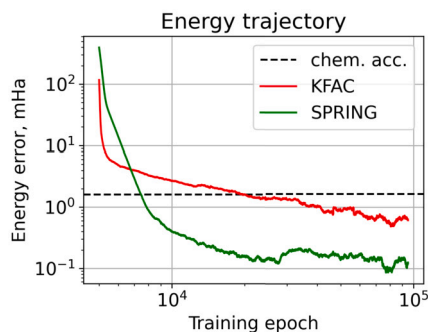


Fig. 11. Comparison of SPRING versus KFAC on the carbon atom without preliminary optimization. We use the previously tuned learning rate of $\eta = 0.02$ for both KFAC and SPRING.

4.4. Computational cost

The computational cost of SPRING is essentially identical to that of MinSR, as the extra computations for SPRING are negligible in cost. The same applies to MinSR+M. For all three of these methods, the computational bottleneck lies in computing the matrix $T = \tilde{O}\tilde{O}^T$, which has asymptotic complexity $O(N_p \cdot N_s^2)$. This cost dominates the $O(N_s^3)$ cost of the Cholesky factorization.

Relative to KFAC, we expect the cost of SPRING to be highly dependent on the system size, the number of MCMC samples per iteration, and the number of GPUs used to parallelize the calculation, and we do not attempt to perform a systematic study along these lines. In our experiments, we use a single GeForce RTX 2080Ti GPU with 1000 MCMC walkers. With these particular settings, we find that SPRING results in VMC iterations that are about 50% slower than KFAC for our smaller systems such as the carbon atom. As system size grows, the difference is reduced, and for the N_2 molecule the VMC iterations with SPRING are only about 5% slower than with KFAC.

5. Discussion

Neural network wavefunctions represent a promising avenue towards highly accurate simulations of small but challenging molecular systems. The major bottleneck of applying such wavefunctions is the high cost associated with their optimization. In this work we introduce a new optimizer called SPRING to alleviate this bottleneck. SPRING combines ideas from the recently proposed MinSR optimizer [26] with ideas from the randomized Kaczmarz method for solving overdetermined least-squares problems [28]. By doing so, SPRING is able to utilize optimization history to improve upon MinSR in a principled way at essentially no extra cost. We test SPRING on several small atoms and molecules, comparing it against MinSR, MinSR with momentum, and KFAC. We find that SPRING consistently outperforms the alternatives across all tested systems. We hope that these findings will be extended to larger systems by future works.

Due to several discrepancies between the VMC setting and the traditional Kaczmarz setting, we are not yet able to furnish a rigorous proof of convergence for the SPRING optimizer. One interesting direction for future research is to modify SPRING in a way that makes it possible to prove its convergence. For example, the exponentially convergent randomized Kaczmarz method of Strohmer and Vershynin requires that the rows are sampled with probability proportional to the square of their Euclidean norm [28]. Furthermore, in the randomized block Kaczmarz method, Needell and Tropp find that for optimal performance it is critical to partition the constraints of the least-squares problem into well-conditioned blocks [29]. Such sampling schemes would represent a significant departure from traditional VMC methods that sample from the probability density of the wavefunction, but they could lead to rigorous convergence guarantees, better performance, or both. There are also variants of the Kaczmarz method such as the randomized extended Kaczmarz algorithm of Zouzias and Freris [43] that can converge to the solution of inconsistent least-squares problems. This technique cannot be directly applied to the VMC setting since it requires sampling the columns as well as the rows of the system. However, it could serve as inspiration for further development of the algorithm.

We also draw a connection between the MinSR and SPRING methods for optimizing neural network wavefunctions and the efficient subsampled natural gradient descent method of Ren and Goldfarb [25]. In particular, we show in Appendix A that MinSR can be viewed as a simplified implementation of the method of Ren and Goldfarb which applies as long as the gradient of the loss function is a linear combination of the model gradients at the sampled points. This is quite a common scenario in machine learning, occurring for example in supervised learning with a mean-squared error loss function. In such a setting, SPRING can then be viewed as a potential improvement to existing subsampled natural gradient descent methods. We leave it to future works to determine whether SPRING can yield performance improvements for applications outside of VMC.

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed,

or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

CRedit authorship contribution statement

Gil Goldshlager: Writing – review & editing, Writing – original draft, Software, Methodology, Investigation, Conceptualization. **Nilin Abrahamsen:** Writing – review & editing, Software, Methodology, Investigation. **Lin Lin:** Writing – review & editing, Supervision, Resources, Methodology, Conceptualization.

Declaration of competing interest

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

Data availability

Data will be made available on request.

Acknowledgements

This research used the Savio computational cluster resource provided by the Berkeley Research Computing program at the University of California, Berkeley (supported by the UC Berkeley Chancellor, Vice Chancellor for Research, and Chief Information Officer). This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number(s) DE-SC0023112 (G.G.). This effort was supported by the SciAI Center, and funded by the Office of Naval Research (ONR), under Grant Number N00014-23-1-2729 (N.A.). LL is a Simons Investigator in Mathematics. We thank Yixiao Chen, Zhiyan Ding, Yuehaw Khoo, Michael Lindsey, Eric Neuscamman, and Zaiwen Wen for their helpful discussions, and the anonymous reviewers for their valuable comments.

Appendix A. Connection to efficient subsampled natural gradient descent

Recall that MinSR with Tikhonov regularization uses the update formula

$$d\theta = \bar{O}^T (\lambda I + \bar{O}\bar{O}^T)^{-1} \bar{\epsilon}, \quad (\text{A.1})$$

where λ is the damping parameter. This is the same formula that arises from the “simple linear algebra trick” of Rende et al. [25]. We now introduce the efficient subsampled natural gradient method of Ren and Goldfarb [25] and show how it is equivalent to MinSR in the VMC setting. The key idea of Ren and Goldfarb is that a damped Fisher matrix arising from a small minibatch can be viewed as a low-rank perturbation of the identity and is thus amenable to inversion using the Sherman-Morrison-Woodbury formula. Using the notation of the later work by Yang et al. [44], the resulting update is

$$d\theta = \frac{1}{\lambda} \left(-I + U (\lambda I + U^T U)^{-1} U^T \right) g,$$

where g is the gradient estimator and the columns of U contain the model gradients divided by the square-root of the number of samples used.

Now suppose, as is often the case, that the gradient of the loss function is simply a linear combination of the gradients at each input, with some input-dependent scalar weights. Then we can write the gradient estimator as $g = Uv$ for some column vector v that depends on the sampled inputs. This allows us to simplify the NGD-SMW formula as follows:

$$\begin{aligned} d\theta &= \frac{1}{\lambda} \left(-I + U (\lambda I + U^T U)^{-1} U^T \right) Uv \\ &= \frac{1}{\lambda} \left(-U + U (\lambda I + U^T U)^{-1} U^T U \right) v \\ &= \frac{1}{\lambda} \left(-U + U (\lambda I + U^T U)^{-1} (\lambda I + U^T U) - U (\lambda I + U^T U)^{-1} (\lambda I) \right) v \\ &= -U (\lambda I + U^T U)^{-1} v. \end{aligned}$$

Comparing with (A.1), we see that this is equivalent to the MinSR formula when we identify \bar{O}^T with U and $\bar{\epsilon}$ with v . Furthermore, this identification is the natural one, since the columns of \bar{O}^T represent the gradients of the logarithm of the normalized wavefunction and the formula $g = \bar{O}^T \bar{\epsilon}$ holds in the VMC setting up to constant factors. Thus, the MinSR method can be viewed as a simplified way of implementing the efficient subsampled natural gradient method in the context of VMC. Furthermore, SPRING can be viewed as a potential improvement to subsampled natural gradient descent and may have applications outside the VMC setting.

Table B.2

List of FermiNet architecture hyperparameters used for all experiments.

Hyperparameter	Value
One-electron stream width	256
Two-electron stream width	16
Number of equivariant layers	4
Backflow activation function	tanh
Number of determinants	16
Exponential envelope structure	Isotropic
Determinant type	Dense

Table B.3

List of settings for the training phase.

Setting	Value
Standard deviations for local energy clipping	5
Number of walkers	1000
MCMC burn-in steps	5000
Training iterations	1e5
MCMC steps between updates	10

Table B.4

List of settings for the inference phase.

Setting	Value
Number of walkers	2000
MCMC burn-in steps for C,N,O	5000
MCMC burn-in steps for N ₂	1e4
Inference iterations	2e4
MCMC steps between local energy measurements	10

Table B.5

List of settings for the preliminary optimization phase.

Setting	Value
Optimizer	KFAC
Learning rate	0.05
Kernel initializers	Orthogonal
Bias initializers	Random normal
Standard deviations for local energy clipping	5
Number of walkers	1000
MCMC burn-in steps	5000
Training iterations for C,N,O	1000
Training iterations for N ₂	5000
MCMC steps between updates	10

Appendix B. Network architecture and hyperparameters

We list in Table B.2 the hyperparameters of the FermiNet architecture that we use in all our experiments. We list in Table B.3 the other settings used for our VMC training phase, in Table B.4 the settings used for our inference phase, and in Table B.5 the settings used for our preliminary optimization.

References

- [1] W.M.C. Foulkes, Lubos Mitas, R.J. Needs, Guna Rajagopal, Quantum Monte Carlo simulations of solids, *Rev. Mod. Phys.* 73 (1) (2001) 33.
- [2] Becca Federico, Sandro Sorella, *Quantum Monte Carlo Approaches for Correlated Systems*, Cambridge University Press, 2017.
- [3] David Pfau, James S. Spencer, Alexander GDG Matthews, W. Matthew, C. Foulkes, Ab initio solution of the many-electron Schrödinger equation with deep neural networks, *Phys. Rev. Res.* 2 (3) (2020) 033429.
- [4] Jan Hermann, Zeno Schätzle, Frank Noé, Deep-neural-network solution of the electronic Schrödinger equation, *Nat. Chem.* 12 (10) (2020) 891–897.
- [5] Zeno Schätzle, P.B. Szabó, Matěj Mezera, Jan Hermann, Frank Noé, Deepqmc: an open-source software suite for variational optimization of deep-learning molecular wave functions, *J. Chem. Phys.* 159 (9) (2023).
- [6] Ingrid von Glehn, James S. Spencer, David Pfau, A self-attention ansatz for ab-initio quantum chemistry, arXiv preprint, arXiv:2211.13672, 2022.
- [7] Leon Gerard, Michael Scherbela, Philipp Marquetand, Philipp Grohs, Gold-standard solutions to the Schrödinger equation using deep learning: how much physics do we need?, *Adv. Neural Inf. Process. Syst.* 35 (2022) 10282–10294.

- [8] Gino Cassella, Halvard Sutterud, Sam Azadi, N.D. Drummond, David Pfau, James S. Spencer, W.M.C. Foulkes, Discovering quantum phase transitions with fermionic neural networks, *Phys. Rev. Lett.* 130 (Jan 2023) 036401.
- [9] Jane Kim, Gabriel Pescia, Bryce Fore, Jannes Nys, Giuseppe Carleo, Stefano Gandolfi, Morten Hjorth-Jensen, Alessandro Lovato, Neural-network quantum states for ultra-cold Fermi gases, 2023.
- [10] Wan Tong Lou, Halvard Sutterud, Gino Cassella, W.M.C. Foulkes, Johannes Knolle, David Pfau, James S. Spencer, Neural wave functions for superfluids, 2023.
- [11] Gabriel Pescia, Jannes Nys, Jane Kim, Alessandro Lovato, Giuseppe Carleo, Message-passing neural quantum states for the homogeneous electron gas, 2023.
- [12] Xiang Li, Zhe Li, Ji Chen, Ab initio calculation of real solids via neural network ansatz, *Nat. Commun.* 13 (1) (2022) 7895.
- [13] Sandro Sorella, Generalized Lanczos algorithm for variational quantum Monte Carlo, *Phys. Rev. B* 64 (Jun 2001) 024512.
- [14] M.P. Nightingale, Vilen Melik-Alaverdian, Optimization of ground- and excited-state wave functions and van der Waals clusters, *Phys. Rev. Lett.* 87 (4) (2001) 043401.
- [15] Sandro Sorella, Michele Casula, Dario Rocca, Weak binding between two aromatic rings: feeling the van der Waals attraction by quantum Monte Carlo methods, *J. Chem. Phys.* 127 (1) (2007).
- [16] C.J. Umrigar, Julien Toulouse, Claudia Filippi, S. Sorella, R.G. Hennig, Alleviation of the fermion-sign problem by optimization of many-body wave functions, *Phys. Rev. Lett.* 98 (Mar 2007) 110201.
- [17] Julien Toulouse, Cyrus J. Umrigar, Optimization of quantum Monte Carlo wave functions by energy minimization, *J. Chem. Phys.* 126 (8) (2007).
- [18] Eric Neuscamman, C.J. Umrigar, Garnet Kin-Lic Chan, Optimizing large parameter sets in variational quantum Monte Carlo, *Phys. Rev. B* 85 (4) (2012) 045103.
- [19] Luning Zhao, Eric Neuscamman, A blocked linear method for optimizing large parameter sets in variational Monte Carlo, *J. Chem. Theory Comput.* 13 (6) (2017) 2604–2611.
- [20] Iliya Sabzevari, Ankit Mahajan, Sandeep Sharma, An accelerated linear method for optimizing non-linear wavefunctions in variational Monte Carlo, *J. Chem. Phys.* 152 (2) (2020).
- [21] Robert J. Webber, Michael Lindsey, Rayleigh-Gauss-Newton optimization with enhanced sampling for variational Monte Carlo, *Phys. Rev. Res.* 4 (3) (2022) 033099.
- [22] James Martens, Roger Grosse, Optimizing neural networks with Kronecker-factored approximate curvature, in: *International Conference on Machine Learning*, PMLR, 2015, pp. 2408–2417.
- [23] Shun-ichi Amari, Natural gradient works efficiently in learning, *Neural Comput.* 10 (2) (1998) 251–276.
- [24] Frederik Benzing, Gradient descent on neurons and its link to approximate second-order optimization, in: *International Conference on Machine Learning*, PMLR, 2022, pp. 1817–1853.
- [25] Yi Ren, Donald Goldfarb, Efficient subsampled Gauss-Newton and natural gradient methods for training neural networks, arXiv preprint, arXiv:1906.02353, 2019.
- [26] Ao Chen, Markus Heyl, Efficient optimization of deep neural quantum states toward machine precision, 2023.
- [27] Riccardo Rende, Luciano Loris Viteritti, Lorenzo Bardone, Federico Becca, Sebastian Goldt, A simple linear algebra identity to optimize large-scale neural network quantum states, arXiv preprint, arXiv:2310.05715, 2023.
- [28] Thomas Strohmer, Roman Vershynin, A randomized Kaczmarz algorithm with exponential convergence, *J. Fourier Anal. Appl.* 15 (2) (2009) 262–278.
- [29] Deanna Needell, Joel A. Tropp, Paved with good intentions: analysis of a randomized block Kaczmarz method, *Linear Algebra Appl.* 441 (2014) 199–221.
- [30] Jeffmin Lin, Gil Goldshlager, Lin Lin, VMCNet: flexible, general-purpose VMC framework, built on JAX, <http://github.com/jeffminlin/vmcnet>, 2021.
- [31] M. Born, R. Oppenheimer, Zur Quantentheorie der Molekeln, *Ann. Phys.* 4 (84) (1927) 457–484.
- [32] Jeffmin Lin, Gil Goldshlager, Lin Lin, Explicitly antisymmetrized neural network layers for variational Monte Carlo simulation, *J. Comput. Phys.* 474 (2023) 111765.
- [33] Giuseppe Carleo, Matthias Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* 355 (6325) (2017) 602–606.
- [34] Giuseppe Carleo, Kenny Choo, Damian Hofmann, James Smith, Tom Westerhout, Fabien Alet, Emily J. Davis, Stavros Efthymiou, Ivan Glasser, Sheng-Hsuan Lin, et al., Netket: a machine learning toolkit for many-body quantum systems, *SoftwareX* 10 (2019) 100311.
- [35] Chae-Yeun Park, Michael J. Kastoryano, Geometry of learning neural quantum states, *Phys. Rev. Res.* 2 (2) (2020) 023232.
- [36] Guodong Zhang, James Martens, Roger B. Grosse, Fast convergence of natural gradient descent for over-parameterized neural networks, *Adv. Neural Inf. Process. Syst.* 32 (2019).
- [37] Stefan Kaczmarz, Angenaherte auflösung von systemen linearer glei-chungen, *Bull. Int. Acad. Pol. Sic. Let., Cl. Sci. Math. Nat.* (1937) 355–357.
- [38] Deanna Needell, Randomized Kaczmarz solver for noisy linear systems, *BIT Numer. Math.* 50 (2010) 395–403.
- [39] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, Qiao Zhang, JAX: composable transformations of Python+NumPy programs, 2018.
- [40] Subhas J. Chakravorty, Steven R. Gwaltney, Ernest R. Davidson, Farid A. Parpia, Charlotte Froese p Fischer, Ground-state correlation energies for atomic ions with 3 to 18 electrons, *Phys. Rev. A* 47 (5) (1993) 3649.
- [41] Claudia Filippi, Cyrus J. Umrigar, Multiconfiguration wave functions for quantum Monte Carlo calculations of first-row diatomic molecules, *J. Chem. Phys.* 105 (1) (1996) 213–226.
- [42] Robert J. Le Roy, Yiye Huang, Calvin Jary, An accurate analytic potential function for ground-state n_2 from a direct-potential-fit analysis of spectroscopic data, *J. Chem. Phys.* 125 (16) (2006).
- [43] Anastasios Zouzias, Nikolaos M. Freris, Randomized extended Kaczmarz for solving least squares, *SIAM J. Matrix Anal. Appl.* 34 (2) (2013) 773–793.
- [44] Minghan Yang, Dong Xu, Zaiwen Wen, Mengyun Chen, Pengxiang Xu, Sketch-based empirical natural gradient methods for deep learning, *J. Sci. Comput.* 92 (3) (2022) 94.