

PAPER • OPEN ACCESS

Machine learning-based compression of quantum many body physics: PCA and autoencoder representation of the vertex function

To cite this article: Jiawei Zang *et al* 2024 *Mach. Learn.: Sci. Technol.* **5** 045076

View the [article online](#) for updates and enhancements.

You may also like

- [An efficient Wasserstein-distance approach for reconstructing jump-diffusion processes using parameterized neural networks](#)
Mingtao Xia, Xiangting Li, Qijing Shen et al.
- [Refinable modeling for unbinned SMEFT analyses](#)
Robert Schöfbeck
- [Parameter uncertainties for imperfect surrogate models in the low-noise regime](#)
Thomas D Swinburne and Danny Perez



OPEN ACCESS

RECEIVED
16 May 2024

REVISED
8 November 2024

ACCEPTED FOR PUBLICATION
12 December 2024

PUBLISHED
30 December 2024

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



PAPER

Machine learning-based compression of quantum many body physics: PCA and autoencoder representation of the vertex function

Jiawei Zang^{1,*} , Matija Medvidović^{1,2}, Dominik Kiese², Domenico Di Sante³, Anirvan M Sengupta^{2,4,5} and Andrew J Millis^{1,2}

¹ Department of Physics, Columbia University, 538 W 120th Street, New York, NY 10027, United States of America

² Center for Computational Quantum Physics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, United States of America

³ Department of Physics and Astronomy, Alma Mater Studiorum—University of Bologna, Bologna 40127, Italy

⁴ Department of Physics and Astronomy, Rutgers University, 136 Frelinghuysen Road, Piscataway, NJ 08854, United States of America

⁵ Center for Computational Mathematics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, United States of America

* Author to whom any correspondence should be addressed.

E-mail: jz3122@columbia.edu

Keywords: deep learning, machine learning, quantum system, many-body system

Supplementary material for this article is available [online](#)

Abstract

Theoretical approaches to quantum many-body physics require developing compact representations of the complexity of generic quantum states. This paper explores an interpretable data-driven approach utilizing principal component analysis (PCA) and *autoencoder* neural networks to compress the two-particle vertex, a key element in Feynman diagram approaches. We show that the linear PCA offers more physical insight and better out-of-distribution generalization than the nominally more expressive autoencoders. Even with ~ 10 – 20 principal components, we find excellent reconstruction across the phase diagram suggesting the existence of heretofore unrealized structures in the diagrammatic theory. We show that the principal components needed to describe the ferromagnetic state are not contained in the low rank description of the Fermi liquid (FL) state, unlike those for antiferromagnetic and superconducting states, suggesting that the latter two states emerge from pre-existing fluctuations in the FL while ferromagnetism is driven by a different process.

1. Introduction

One of the grand computational challenges in present-day quantum many-body physics is understanding large systems of interacting particles. Quantum physics is naturally formulated as a theory of linear operators acting on a Hilbert space whose dimension grows exponentially with the number of degrees of freedom. Exact diagonalization of any many-body Hamiltonian in this space quickly becomes unfeasible. Likewise, the infamous sign problem prevents statistically accurate solutions of generic fermionic models by means of quantum Monte Carlo methods [1] and it is generally believed that the solution of a general fermionic quantum many-body problem is NP hard [2]. An alternative approach to the problem is via *diagrammatic methods*, in which the two-particle vertex Γ (called the *vertex function*) plays a crucial role. The vertex function generally depends on three momenta $\mathbf{k}_i$ and three frequencies ν_i and can be written as $\Gamma(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)$ where $\mathbf{k}_i = (\mathbf{k}_i, \nu_i)$ denotes a frequency-momentum four-vector. The vertex function Γ describes the two-particle scattering from the initial states $\mathbf{k}_2$ and $\mathbf{k}_4 = \mathbf{k}_1 + \mathbf{k}_3 - \mathbf{k}_2$ into the final states $\mathbf{k}_1$ and $\mathbf{k}_3$ and may also be viewed as describing the scattering of a particle-hole pair of total momentum $\mathbf{q} = \mathbf{k}_2 - \mathbf{k}_1$ and relative momentum $2\mathbf{p} + \mathbf{q} = \mathbf{k}_1 + \mathbf{k}_2$ into another pair also of total momentum $\mathbf{q}$ but of relative momentum $2\mathbf{p}' + \mathbf{q} = \mathbf{k}_3 + \mathbf{k}_4$; this latter ‘particle-hole’ representation is shown in figure 1(a). Knowledge of the vertex function enables insights into the many-body properties of the system, including its response to external fields as well as its tendency to develop long-range order [1, 3–6]. However, the vertex function is in general

difficult to calculate and, although not exponentially large, demands large computational and memory resources with increasing system size and decreasing temperature. While there has been recent progress on compressing the frequency structure of single-particle [7, 8] and also two-particle Green's functions [9, 10], low-dimensional representation of the full vertex function remains an open question [11]. Deep learning approaches have recently had success in constructing such compact latent representations of the vertex function. Di Sante *et al* [12] have used an encoder-decoder structure on top of neural ordinary differential equations [13]. A latent space is learned that can represent the entire renormalization group at a fraction of the original computational cost. This naturally raises the question: is there a concise representation of the momentum structure of Γ ?

In this paper, we use principal component analysis (PCA) and a deep convolutional autoencoder [14] to address this question by compressing a large dataset of vertex functions assembled from functional renormalization group (fRG) calculations [15] for the 2D Hubbard model [16]. To determine the fidelity of the compressed representation we compute the pointwise mean square difference between the true and reconstructed vertices, and we compute physical quantities related to generalized susceptibilities and to the tendency of the system to order into different phases. We find in general that very low dimensional representations suffice to capture the physics accurately, suggesting the existence of a heretofore unsuspected simple underlying structure for the vertex function. We also find that PCA is markedly superior to the autoencoder in achieving the lowest possible reconstruction error for a fixed dimensionality of the representation.

Further, quantum many-body systems may be in ordered states (for example, ferromagnetic (FM), antiferromagnetic (AFM) or superconducting (SC)) or may be in quantum disordered states such as the conventional Fermi liquid (FL). One may ask whether the different forms of order emerge from the FL state by an amplification of pre-existing correlations as interaction parameters are varied, or whether an order emerges via the introduction of new physics not manifest in the FL. Our analysis enables us to address this question by quantifying the overlap of principal components of different states. Surprisingly, we find that the FM state is not naturally contained in the FL state, in the sense that the minimal number of principal components (4–20) that describe the FL with good accuracy do not provide an accurate description of the FM state vertex. On the other hand the vertex functions appropriate to the AFM and SC states can be described within the same minimal basis that describes the FL vertices. Our findings suggest that quantum many-body physics has hitherto unsuspected structure that may allow it to be reformulated in a compact and computationally efficient basis, and shows how comparing machine-learning-based data compressions across regimes of different physics and potentially across different systems may provide new insights. Our findings forge a path towards refining the computation of vertex functions, offering potential advantages in computational efficiency and physics discovery.

The rest of this paper is organized as follows. In section 2 we present the model. In section 3 we present the methods. In section 4.1 we introduce measures of the accuracy of the compression and use them to compare the reconstruction loss for PCA and autoencoder. In section 4.2 we study the reconstruction of ordered states using the information contained in FL states. In section 5 we give a summary, conclusion and outlook.


2. Model

We employ the two-dimensional square lattice Hubbard model, a paradigmatic model of interacting electrons on a lattice, as a test bed for our methodology. The model may be written in standard second quantized notation using a mixed momentum ($\mathbf{k}$) and lattice site (i) representation as

$$H = \sum_{\mathbf{k}, s} \xi_{\mathbf{k}} c_{\mathbf{k}s}^\dagger c_{\mathbf{k}s} - \mu \sum_{i, s} n_{is} + U \sum_i n_{i\uparrow} n_{i\downarrow}. \quad (1)$$

Here, $c_{\mathbf{k}s}^\dagger$ creates an electron with momentum $\mathbf{k}$ and spin s ; $\xi_{\mathbf{k}} = -2t[\cos(k_x) + \cos(k_y)] - 4t'\cos(k_x)\cos(k_y)$ describes how electrons move on the lattice (t and t' are respectively the quantum mechanical amplitudes for electrons to move from one site to its first or second neighbor), $n_{is} = c_{is}^\dagger c_{is}$ yields the number of electrons on site i , μ is the chemical potential that controls the electron density, and $U > 0$ is a repulsive interaction that correlates electron motion by disfavoring configurations with two electrons per site. In this paper, we set $U = 3t$ with $t = 1$ as our unit of energy.

For generic μ and t' the model is in a non-ordered FL state. However, for each t' there is a value of μ at which the Fermi surface passes through a 'van Hove' point at which the non-interacting electron density of states diverges. For chemical potentials near this point, the large density of states leads to an ordered state. Figure 1(b) shows the ground state phase diagram. The solid line shows the locus of van Hove points; i.e. the



t' dependence of the chemical potential value at which the non-interacting density of states diverges. For chemical potentials far from this line, one has a non-ordered FL phase, and in the vicinity of the line the model exhibits either AFM, SC, or FM order, depending on the value of t' [17] (the shaded regions denote approximate regimes for different phases). To probe the different phases illustrated in figure 1(b), we study the evolution of the vertex function along four specific trajectories: one follows the van Hove singularity along $\mu = 4t'$ for t' ranging from 0 to 0.5, depicted as a solid line; the other three track changes in μ , from $-2 + 4t'$ to $4t'$, at fixed t' values of 0.0, 0.25, and 0.5, represented by dashed lines. For each trajectory, we take 50 data points at intervals of $\Delta\mu = 0.04$, resulting in a total of 200 data points.

We employ the fRG method [16, 18] to calculate the vertex function $\Gamma(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)$. A detailed procedure is provided in the next section. Sample results, presented as two dimensional heat maps of a particular slice of the three argument vertex are shown in figure 1(c). It is noteworthy that the vertices have very substantial, nearly singular, $\mathbf{k}$ dependence, appearing as stripes in the plots. The calculated vertices enable one to infer the phase structure of the model via a calculation of the eigenvalues of the vertex in that particular channel where long-range order emerges. One defines an ordering eigenvector as the expectation value of a fermion particle-hole or particle particle bilinear $\langle c_{\mathbf{k}_1}^\dagger c_{\mathbf{k}_2} \rangle$ or $\langle c_{\mathbf{k}_1} c_{\mathbf{k}_2} \rangle$ (spin indices not explicitly notated) and forms an eigenvalue equation by contracting this object with appropriate indices of Γ as shown in figure 1(a). The onset of order is signaled by a divergent eigenvalue and the bilinear with the most divergent eigenvalue λ determines the preferred order parameter.

3. Methods

3.1. fRG

To calculate vertex functions Γ for the model, we employ the fRG method [16, 18], an established computational tool for studying Fermi surface instabilities of low-dimensional interacting electron systems. At the core of the fRG method is a set of coupled non-linear differential flow equations whose solution determines the *flow* of the vertex as function of the RG scale. If momentum space is discretized into N_k tiles along each dimension of d -dimensions (see [16] for details on the discretization), Γ is specified by N_k^{3d}

complex numbers determined by the solution of N_k^{3d} coupled nonlinear equations (as common practice in fRG we only consider the lowest frequency component of the vertex, and such a vertex, $\Gamma(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)$, is real-valued.). In this work, we chose $N_k = 24$ with $d = 2$, such that the vertex functions extracted at the end of the fRG flow are notably high-dimensional arrays with $24^6 \approx 2 \times 10^8$ entries in total. To enable a comparison between PCA and the autoencoder without excessive computational costs, we down-sampled the vertex by retaining every second point along each dimension. This down-sampling reduces the data to a manageable 144^3 entries, ensuring efficiency while preserving a substantial dimensionality of about 10^6 . Details about the generation of the input data are shown in the supplemental material.

The onset of order is signalled by a divergent eigenvalue and the bilinear with the most divergent eigenvalue λ determines the preferred order parameter. In our fRG calculations we stop the renormalization group flow if the magnitude of the vertex function exceeds some preset multiple of the electronic bandwidth D (we choose $4D$ as the threshold) and calculate the eigenvalues and eigenvectors for different long-range orders in a post-processing step. It is important to note that in the computational basis the divergence in the eigenvalue arises not from a divergence in a few entries of the matrix but rather from the coherent combination of a sum over many elements, and yet the combination of elements that produces a divergent eigenvalue is identified by the dimensional reduction procedure.

3.2. PCA and autoencoder

We executed PCA using a randomized SVD solver [19]. This linear method initially rotates the data into a basis where components are linearly uncorrelated, thereby identifying the principal directions, or components, of independent variation. The significance of each coordinate is ranked by the amount of variance in the data which can be captured. In effect, retaining the M most significant SVD principal components defines a M dimensional subspace of the original dimensional space of all possible vertices spanned by vertices $\Gamma_{1...M}$ (each with a complicated internal structure) and the compression is the approximation that any physical vertex Γ can be represented to sufficient accuracy as $\hat{\Gamma} = \sum_{i=1...M} a_i \Gamma_i$. The number of SVD principal axes is limited by the training sample size. However, our objective is not to capture all 10^6 dimensions but to extract a low-dimensional latent space that encapsulates the physically relevant features of the vertex functions for the Hubbard model. By verifying our model's accuracy on test data, we ensure that the retained components effectively capture this latent space within the original 10^6 -dimensional space.

Conversely, the autoencoder is a non-linear method designed to efficiently encode a given set of data. It compresses the data into a low-dimensional latent space by applying a flexible parametrized transformation (the *encoder*). A second model (the *decoder*) is trained to reconstruct the original dataset as accurately as possible. The vertex is then approximated as a nonlinear function (specified by the decoder) of a number of parameters equal to the latent space dimension. We use a five-layer convolutional neural-network [20] architecture for both the encoder and decoder with layer norms and GeLU activations [21]. In addition to transposed convolutions, we use non-parametric upsampling after each layer. Details are shown in supplementary material. The reconstruction error, defined in equation (2), serves as the cost function. Each autoencoder was trained with approximately 4 GPU-hours on a single Nvidia A100 card.


4. Result

4.1. Compression loss: comparison between PCA and the autoencoder

In this section we introduce measures of the accuracy of the compression. we use three error metrics: the direct (pointwise) reconstruction error of the vertex, given by

$$\epsilon = \frac{\|\hat{\Gamma} - \Gamma\|_2}{\|\Gamma\|_2}, \quad (2)$$

the reconstruction accuracy of the leading eigenvalue λ [22] for several instability channels, and the Landau parameters, averaged over the Fermi surface of the vertex functions weighted by angular and spin factors (see supplementary material for details). For simplicity, we restrict our study to the ordered states only, which is along the van Hove trajectory of $\mu = 4t'$. But the reconstruction errors for non-ordered states are comparable to the ordered states. Due to space limitations, we did not include these additional results. We use 50 vertex functions with t' ranging from 0 to 0.5 in increments of $\Delta t' = 0.01$ and took a random 80/20 train/test split of the input data, and subsequently evaluate the error metrics utilizing the entire dataset, as the error between the training and test sets was comparable and it provides a more comprehensive view of the trajectory of the ordered states. As shown in figure 2(a), PCA shows decreasing reconstruction errors with increasing latent dimension d , whereas the autoencoder's error remains roughly constant across



different d values. For details on the variance captured by each principal component from PCA, see the plot of the explained variance percentage in the supplementary material. Figure 2(b) shows the eigenvalues associated with AFM, SC, and FM order along the van Hove trajectory, with the smallest eigenvalues indicating the system's propensity towards the respective ground state. Transition points are marked by dashed vertical lines. For PCA, transitions in the reconstruction data align well with the original calculations when the latent dimension exceeds $d = 4$, and the accuracy keeps improving with higher d , becoming comparable to the autoencoder errors for $d \geq 4$ and clearly superior to the autoencoder errors for $d \geq 16$. In contrast, the autoencoder maintains roughly the same accuracy across various d .

Figure 2(d) shows the Landau parameters l for spin and charge channels. The Landau parameters [3] are canonical quantities defined in the theory of the FL that are related to susceptibilities and collective mode frequencies. For PCA, the amplitude and tendency of l are correctly captured even at $d = 1$, indicating that the first PCA axis contains substantial information about the development of the Landau parameter, which suggests the presence of FM states. However, the accuracy of l is significantly higher for $d = 8$ compared to $d = 1$. For the autoencoder, the accuracy remains high for different d .

Lastly, figure 2(c) presents a side-by-side comparison of a test slice of the original vertex function $\Gamma(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)$ and its reconstructions at $\mathbf{k}_3 = 0$ and $t' = 0$. We see that a relatively small number of dimensions is sufficient to represent both the vertex and physical quantities, including those relating to divergences of particular eigenvalues. As noted above, in the computational basis the divergent eigenvalues do not arise from divergences in the entries in the Γ matrix. Our findings strongly suggest that a limited number of dimensions is sufficient to gather the essential information contained in the vertex, including information about its singularities.

The exact reason behind the superior performance of the PCA when compared against a parameterized autoencoder neural network remains an open question. While we leave an exhaustive search over possible architecture choices for future research, our experiments with convolutional networks in three dimensions




Figure 3. (a) Reconstruction error ϵ (equation (2)) for all ordered states along van Hove trajectory ($t' = 0 - 0.5$) as a function of the latent dimension number d of PCA bases derived from all FL states. The data points are color-coded to indicate their association with FM, SC and AFM states. (b), (c) The eigenvalues associated with different physical orders and Landau parameters calculated along the van Hove trajectory. It compares the original vertex data with the FL PCA reconstruction that utilize different number (20, 50, 100) of bases.


suggest that local filters fail to capture essential features in interaction vertices with a strong global sparsity structure. In contrast, performing a global linear rotation, PCA is able to capture global patterns. For other tasks on larger vertex datasets, we conjecture that using PCA as a preprocessing or input layer would result in a more expressive model.

4.2. FL and ordered states

An interesting physics question concerns whether an ordered state may be understood as arising from an interaction-driven enhancement of fluctuations that are already present in the FL state, or whether the onset of the ordered state signals new physics not evident from an analysis of the FL state. The compression methods introduced here offer a new perspective on this question. We perform PCA again with only the FL vertices and obtain a reduced dimension subspace that accurately represents the FL vertices (117 vertices in total). We then assess the accuracy with which the vertices in the ordered state regions of the phase diagram may be represented within the subspace defined from the FL phase.

Panel (a) of figure 3 shows the pointwise reconstruction error ϵ as a function of the size of d of the PCA subspace used to represent the FL vertices computed along the van Hove line for a large number of t' values, shown as different lines color-coded according to the relevant type of order. We see that for relatively small $d \leq 10$ the FL-defined subspace provides a relatively poor compression of the ordered state vertices, but that as the subspace dimension is increased the vertices in the antiferromagnetic and SC phases become relatively well represented while the vertices in the FM region remain very poorly described. This analysis indicates that none of the ordered states are accurately represented by the leading terms of FL state, but that considering the vertex function's high dimension ($\sim 10^6$), the SC and AFM states are moderately well represented in terms of smaller but non-negligible contributions to the state. On the other hand, the FM state vertex function appears to require qualitatively new components: the structure learned or inferred from the FL phase does not transfer well to the FM phase. Additionally, the errors indicate a separation at about $t' = 0.34t$, coinciding with the FM and SC state transition.

Figure 3(b) presents the eigenvalues from the original and reconstructed vertices for various d . We see that the leading eigenvalue of the AFM phase is very well represented even at $d = 20$ consistent with the lower ϵ but that for the SC vertex the representation derived from the FL phase, while qualitatively reasonable at $d = 20$ only begins to become accurate by $d = 50$. The FM eigenvalue is qualitatively incorrect at $d = 20$ and becomes correct for $d = 50$. Figure 3(c) presents Landau parameters from the original and reconstructed vertices. The Landau parameter in the spin channel initially presents inaccuracies at $d = 20$, attaining precision only when $d \geq 50$, similar to the behavior observed for the FM eigenvalue. Conversely, in the charge channel, while the parameter is approximately correct starting from $d = 20$, it fails to achieve high



precision even at $d = 100$. This observation confirms that a low dimensional representation of the FL vertices does not transfer well to a representation of physically relevant quantities such as pairing eigenvalues and Landau parameters.

In figure 4 we present a more refined analysis. We consider a sequence of FL states defined by varying μ within the FL regime at fixed $t' = 0.0, 0.25, 0.5$ (vertical lines in figure 1(b)). The three sequences of μ terminate in different ordered states, so we expect that the sequence of FL states will be characterized by a growth of fluctuations in a given channel, leading eventually to the divergent fluctuations implying order, and that the evolution of the vertex will reflect this growth. We then define a PCA subspace for each sequence, and ask how well the ordered state vertices are described within the PCA subspace. We find that the reconstruction errors are larger than those obtained from PCA subspaces constructed from the entire FL data set. Specifically, for bases constructed from vertices at $t' = 0.0$, i.e. in the vicinity of AFM order, SC ordered states exhibit the lowest reconstruction error compared to AFM and FM orders. Similarly, for bases constructed from $t' = 0.25$, which is close to SC order, AFM ordered states yield the smallest ϵ . To develop a qualitative understanding of this circumstance, we calculate the eigenvalues of different orders along the FL to ordered state transition as shown in figure 4(b). For $t' = 0.0$ and $t' = 0.25$, we observe, that concomitantly with the incipient order at the van Hove line ($\mu = 4t'$), i.e. AFM order or singlet-pairing, the respective other ordering channel also shows an increase in the absolute value of λ . This indicates, that some fingerprint of fluctuations in a competing channel must be present in the vertex function, which apparently allows the PCA analysis to extract some information about the vertices in the ordered phase from the FL data. Note that this is not the case for $t' = 0.5$ (see the rightmost panel of figure 4(b)), where the AFM and SC eigenvalues decrease in absolute magnitude when FM order starts to set in.

5. Discussion

In this paper, we have shown that the two-particle vertex obtained from fRG calculations on the two dimensional Hubbard model, is well represented by a point in a low (i.e. 4–12) dimensional latent space derived from PCA or from an autoencoder. We observe that the sharp structures in the exactly calculated vertices even in the FL state show that the success of the compression does not arise simply because the

vertices are smooth (and hence approximated by a few low-order polynomials), and that the successful calculation of the pairing eigenvalues means that the compression can capture singularities arising from coherent sums over many non-divergent terms in the computational basis form of the vertex. The fact that the three-fold momentum dependence of the vertex can be so strongly compressed suggests that there is a hitherto unsuspected structure in the vertex functions and perhaps more generally in the correlation functions of the Hubbard model at moderate interaction strength, and also offers the hope that much more computationally efficient schemes for calculating the vertex can be devised. Our findings thus call for further investigation in this direction.

Note that the compression achieved with the two data-driven approaches discussed in this manuscript is even more efficient than some of the known techniques in the fRG literature. The truncated-unity (TU) approximation [11, 16], for example, assumes a weak-dependence of Γ on the two fermionic momenta in the particle-particle or particle-hole momentum/frequency convention and, consequently, represents it by a set of analytical form factors. Although the number of form factors required to converge TU-fRG calculations is usually small (order 10) [23, 24], fine resolution in the transfer momentum is required to resolve incipient long-range orders, resulting in many more degrees of freedom than identified in our data-driven analysis.

We also found that a low-rank representation of a FL state was not always transferable to the ordered states, suggesting that the ordered phases do not evolve in a simple way from structures of the non-ordered FL. The non-transferability was particularly acute in the FM phase, suggesting that this physics is simply qualitatively different from the physics of the FL phase. One important long-term goal would be the direct calculation of vertex functions in the compressed basis that we unraveled in this manuscript. This would require the projection of the fRG or other many-body equations into the basis suggested by the PCA analysis. Further, our finding that the PCA methods provide a substantially more efficient representation than the standard autoencoder methods suggests that research into the optimizations required in the autoencoder method may be beneficial.

Data availability statement

The data cannot be made publicly available upon publication because the cost of preparing, depositing and hosting the data would be prohibitive within the terms of this research project. The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgments

J Z and A J M acknowledge support from the NSF MRSEC program through the Center for Precision-Assembled Quantum Materials (PAQM)—DMR-2011738. The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 897276 (D D S). M M was supported by the CCQ graduate fellowship in computational quantum physics. The Flatiron Institute is a division of the Simons Foundation.

ORCID iD

Jiawei Zang  <https://orcid.org/0000-0002-3671-9282>

References

- [1] Sénéchal D, Tremblay A-M and Bourbonnais C 2006 *Theoretical Methods for Strongly Correlated Electrons* (Springer)
- [2] Troyer M and Wiese U-J 2005 Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations *Phys. Rev. Lett.* **94** 170201
- [3] Abrikosov A A, Dzyaloshinskii I, Gorkov L P and Silverman R A 1975 *Methods of Quantum Field Theory in Statistical Physics* (Dover)
- [4] Tahvildar-Zadeh A N, Freericks J K and Jarrell M 1997 Magnetic phase diagram of the Hubbard model in three dimensions: the second-order local approximation *Phys. Rev. B* **55** 942
- [5] Kuneš J 2011 Efficient treatment of two-particle vertices in dynamical mean-field theory *Phys. Rev. B* **83** 085102
- [6] Rohringer G, Valli A and Toschi A 2012 Local electronic correlation at the two-particle level *Phys. Rev. B* **86** 125114
- [7] Kaye J, Chen K and Parcollet O 2022 Discrete Lehmann representation of imaginary time Green's functions *Phys. Rev. B* **105** 235115
- [8] Shinaoka H, Otsuki J, Ohzeki M and Yoshimi K 2017 Compressing Green's function using intermediate representation between imaginary-time and real-frequency domains *Phys. Rev. B* **96** 035147
- [9] Shinaoka H, Otsuki J, Haule K, Wallerberger M, Gull E, Yoshimi K and Ohzeki M 2018 Overcomplete compact representation of two-particle Green's functions *Phys. Rev. B* **97** 205111
- [10] Moghadas E, Dräger N, Toschi A, Zang J, Medvidović M, Kiese D, Millis A J, Sengupta A M, Andergassen S and Di Sante D 2024 Compressing the two-particle Green's function using wavelets: theory and application to the Hubbard atom (arXiv:2402.13030)
- [11] Lichtenstein J, de la Peña D S, Rohe D, Di Napoli E, Honerkamp C and Maier S 2017 High-performance functional renormalization group calculations for interacting fermions *Comput. Phys. Commun.* **213** 100–10

- [12] Di Sante D, Medvidović M, Toschi A, Sangiovanni G, Franchini C, Sengupta A M and Millis A J 2022 Deep learning the functional renormalization group *Phys. Rev. Lett.* **129** 136402
- [13] Chen R T Q, Rubanova Y, Bettencourt J and Duvenaud D 2018 Neural ordinary differential equations *Conf. on Neural Information Processing Systems* vol 31 (arXiv:1806.07366v5)
- [14] Murphy K P 2022 *Probabilistic Machine Learning: An Introduction* (MIT Press)
- [15] Metzner W, Salmhofer M, Honerkamp C, Meden V and Schönhammer K 2012 Functional renormalization group approach to correlated fermion systems *Rev. Mod. Phys.* **84** 299–352
- [16] Beyer J, Hauck J B and Klebl L 2022 Reference results for the momentum space functional renormalization group *Eur. Phys. J. B* **95** 65
- [17] Honerkamp C and Salmhofer M 2001 Magnetic and superconducting instabilities of the Hubbard model at the van Hove filling *Phys. Rev. Lett.* **87** 187004
- [18] Honerkamp C and Salmhofer M 2001 Temperature-flow renormalization group and the competition between superconductivity and ferromagnetism *Phys. Rev. B* **64** 184516
- [19] Halko N, Martinsson P-G and Tropp J A 2011 Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions *SIAM Rev.* **53** 217
- [20] LeCun Y, Bengio Y and Hinton G 2015 Deep learning *Nature* **521** 436
- [21] Hendrycks D and Gimpel K 2023 Gaussian error linear units (GELUs) (arXiv:1606.08415 [cs.LG])
- [22] Zhai H, Wang F and Lee D-H 2009 Antiferromagnetically driven electronic correlations in iron pnictides and cuprates *Phys. Rev. B* **80** 064517
- [23] Hille C, Kugler F B, Eckhardt C J, He Y-Y, Kauch A, Honerkamp C, Toschi A and Andergassen S 2020 Quantitative functional renormalization group description of the two-dimensional Hubbard model *Phys. Rev. Res.* **2** 033372
- [24] Gneist N, Classen L and Scherer M M 2022 Competing instabilities of the extended Hubbard model on the triangular lattice: truncated-unity functional renormalization group and application to Moiré materials *Phys. Rev. B* **106** 125141