Learning latent variable evolution for the functional renormalization group

Matija Medvidović
Flatiron Institute, Columbia University
matija.medvidovic@columbia.edu

Alessandro Toschi
TU Vienna

Giorgio Sangiovanni
University of Würzburg

Cesare Franchini
University of Bologna, University of Vienna

Andrew J. Millis
Flatiron Institute, Columbia University

Anirvan M. Sengupta
Flatiron Institute, Rutgers University

Domenico Di Sante
University of Bologna, Flatiron Institute
domenico.disante@unibo.it

Abstract

We perform a data-driven dimensionality reduction of the 4-point vertex function characterizing the functional Renormalization Group (fRG) flow for the widely studied two-dimensional $t-t'$ Hubbard model on the square lattice. We show that a deep learning architecture based on a Neural Ordinary Differential Equations efficiently learns the evolution of low-dimensional latent variables in all relevant magnetic and d-wave superconducting regimes of the Hubbard model.Ultimately, our work uses an encoder-decoder architecture to extract compact representations of the 4-point vertex functions for correlated electrons, a goal of utmost importance for the success of cutting-edge methods for tackling the many-electron problem.

1 Introduction

Interacting electron systems exhibit a rich variety of distinct phenomena at different energy and temperature scales. Upon lowering these scales, new effective degrees of freedom and collective behaviors emerge, typically including competing spin, charge and pairing fluctuations. The difficulties inherent in treating these competing, scale-dependent phenomena on an equal footing represent one of the major obstacles to the numerical solution of theoretical models.

The renormalization group (RG) provides a powerful approach to study these problems [1–5]. The RG property of keeping only relevant degrees of freedom, as a scale parameter is reduced, makes it a valuable tool to study interacting fermions. In its *exact* or *functional* ("fRG") form, the RG is formulated as an exact functional flow equation which provides an effective-action description of the underlying microscopic model [6–9]. In contrast to standard RG, the common formulation of fermionic fRG keeps track of the entire frequency/momentum-dependence of the interaction vertices during the flow [10, 9], introducing the need for alternate data representations [11–17].

The advent of machine learning (ML) techniques and data-driven approaches applied to many body quantum physics has triggered enormous interest [18, 19]. In this paper, we present a data-driven approach for dimensionality reduction of the fRG vertex function $\mathrm{V}(\mathrm{k}_1, \mathrm{k}_2, \mathrm{k}_3)$, whose description traditionally requires computation and storage of a function of three continuous momentum variables. We use a neural network architecture known as Neural Ordinary Differential Equations

Machine Learning and the Physical Sciences workshop, NeurIPS 2022.

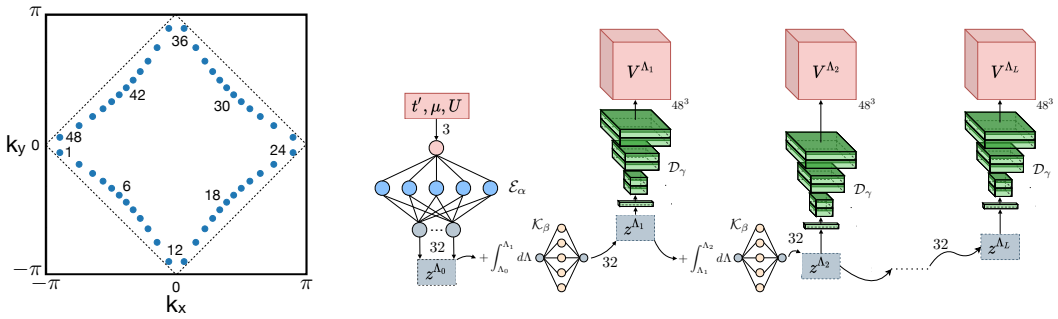

Figure 1: **Left:** The Fermi surface (FS) of the tight-binding model. The blue points indicate the 48 momenta used to patch the FS. The black dashed lines are the Umklapp surface of perfect nesting at $t' = 0$. **Right:** The deep learning architecture described in Eqs. 3.

(NODE) [20], augmented with an encoder-decoder architecture. This approach provides us with additional insight into the low-dimensional structure of the fRG flow [21].

Thus our ML approach to fRG fits within the framework of reduced order models [22] without making unnecessary assumptions. We expect such flexible representation learning techniques to be useful in other vertex-based numerical methods, besides the fRG, which suffer the bottleneck of dealing with high-dimensional data sets.

2 The fRG ground states of the Hubbard model

The microscopic Hamiltonian we consider is

$$H = -t \sum_{\text{nn,s}} c^\dagger_{i,s} c_{j,s} - t' \sum_{\text{nnn,s}} c^\dagger_{i,s} c_{j,s} + U \sum_i n_{i,\uparrow} n_{i,\downarrow} \tag{1}$$

with hopping amplitudes t and t' between nearest neighbours (nn) and next-nearest neighbours (nnn) on the 2D square lattice, and onsite Coulomb repulsion U. The 2-particle properties of this model are investigated through the temperature-flow one-loop fRG scheme [23, 24], where the RG flow of $\mathrm{V}^\Lambda(\mathrm{k}_1, \mathrm{k}_2, \mathrm{k}_3)$ is

$$\frac{\mathrm{dV}^\Lambda}{\mathrm{d}\Lambda} = \mathrm{V}^\Lambda \circ \mathrm{L}^\Lambda \circ \mathrm{V}^\Lambda \tag{2}$$

with the RG scale Λ given by the temperature T, and $\circ$ defining integration over the internal degrees of freedom. Comprehensive reviews of the fRG scheme applied to fermionic problems are found in Refs. [25–27]. In essence, for spin-rotation invariant systems, the right-hand-side of Eq. 2 splits into the sum of three contributions, which describe the particle-particle, direct particle-hole and crossed particle-hole channels [10] necessary to account on similar footing for the SC and density-wave instabilities. L^Λ is a scale-dependent loop kernel that contains information on the single-particle properties of the microscopic model.

Neglecting the frequency dependences of the vertex couplings, which have a negative scaling dimension (irrelevant couplings) under the RG flow, but keeping a full momentum description in terms of a discrete set of N_k wave vectors on the Fermi surface (FS), Eq. 2 is recasted into a set of N_k^3 coupled ordinary differential equations (ODE). The solution to this problem, with initial conditions $\mathrm{V}^{\Lambda_0}(\mathrm{k}_1, \mathrm{k}_2, \mathrm{k}_3) = U$ when $\Lambda_0 = 8t$ is the bandwidth, yields the gradual evolution of the 2-particle vertex function $\mathrm{V}^\Lambda(\mathrm{k}_1, \mathrm{k}_2, \mathrm{k}_3)$ as $\Lambda \to 0$ approaches the FS. For a typical $N_k = 48$ FS discretization, as depicted in Fig. 1, Eq. 2 gives coupled ODEs with more than 10^5 variables.

When Eq. 2 is numerically solved at varying t' for chemical potential fixed at the van Hove filling ($\mu = 4t'$) and at weak-coupling $U = 3t$, the Hubbard model in Eq. 1 experiences three different regimes [23]: i) The first regime is close to half-filling, with $t' > -0.2t$ and dominant AF scattering processes between FS regions connected by wave vectors $\sim (\pi, \pi)$. These can be seen in Fig. 2

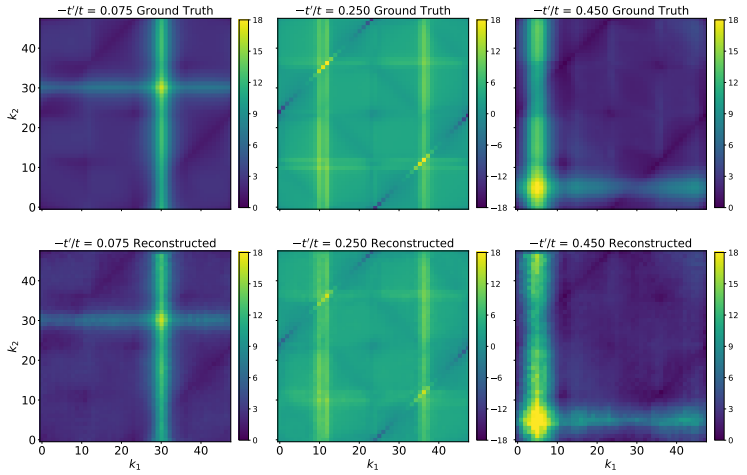

Figure 2: **Top:** False-color representation of amplitude of vertex function $\mathrm{V}^{\Lambda}(\mathrm{k}_1, \mathrm{k}_2, \mathrm{k}_3)$ at a late stage of the renormalization process, for different representative initial conditions. To represent a function of three momenta we fix the first outgoing wave vector k_3 and present the vertex as a function of the two independent incoming vectors. These choices correspond to AFM, d-wave SC and FM instabilities respectively [24, 23]. **Bottom:** The same as the top but for the predicted data $\hat{\mathrm{V}}^{\Lambda}(\mathrm{k}_1, \mathrm{k}_2, \mathrm{k}_3)$. We highlight that these data points belong to the test set.

as bright features corresponding to repulsive couplings on the line $\mathrm{k}_2 - \mathrm{k}_3 \sim (\pi, \pi)$. ii) Further decreasing t', d-wave SC takes over, with the dominant $d_{x^2-y^2}$ symmetry of the pairing scattering, as seen from the sign profile of the diagonal features of Fig. 2. iii) After a quantum critical point for $t' \sim -0.33t$, scattering processes with small momentum transfer $\mathrm{k}_2 - \mathrm{k}_3 \sim (0, 0)$ dominate, see bright features in Fig. 2, leading to a change of ground state from d-wave singlet superconductivity (SC) to ferromagnetism (FM).

3 The Deep Learning fRG: results and interpretation

By inspecting the $\mathcal{O}(10^5)$ couplings of the 2-particle vertex functions of Fig. 2, just before the fRG flow runs to strong coupling and the one-loop approximation breaks down, we recognize that many of them either have remained nearly constant or have become vanishingly small under the RG flow. Only *few* of them have grown positively or negatively (bright features) under the RG evolution. However, as mentioned before, contrary to the standard RG procedure for traditional critical phenomena [3], fermionic fRG does not discard any coupling in the vertex $\mathrm{V}^{\Lambda}(\mathrm{k}_1, \mathrm{k}_2, \mathrm{k}_3)$ during the flow. Our approach is to find a simpler representation in a data-driven manner, using the power of neural nets to find useful features.

In recent years, there have been many developments in utilizing neural networks for predicting sequence data [28]. Since we are interested in finding latent variables whose dynamics itself is governed by an ODE, the natural candidate is a flexible dimensionality reduction scheme based on the Parameterized NODE architecture [20, 29]. The method, sketched in Fig. 1, focuses on three deep neural networks – the encoder $\mathcal{E}$, the NODE $\mathcal{K}$ and the decoder $\mathcal{D}$. The complete action of our model is defined by:

$$\mathbf{z}^{\Lambda_0} = \mathcal{E}_{\alpha}(t', U, \mu) \; ; \; \frac{\mathrm{d}\mathbf{z}^{\Lambda}}{\mathrm{d}\Lambda} = \mathcal{K}_{\beta}(\mathbf{z}^{\Lambda}) \; ; \; \hat{\mathrm{V}}^{\Lambda} = \mathcal{D}_{\gamma}(\mathbf{z}^{\Lambda}) \tag{3}$$

where α, β, γ are parameter sets corresponding to each neural network. The ground truth data are generated by solving the fRG problem in Eq. 2 for 35 values of t' in the range $0 \leq -t'/t < 0.5$ (and $U = 3t$), and storing for each t' the whole vertex dynamics, for a total of ~ 7 thousands collected vertices, each with $\mathcal{O}(N_k^3 \sim 10^5)$ elements.

The encoder maps the Hubbard model parameters and fRG initial condition to a low-dimensional latent representation $\mathbf{z}^{\Lambda_0}$ of drastically smaller dimension than N_k^3. All the results here are obtained

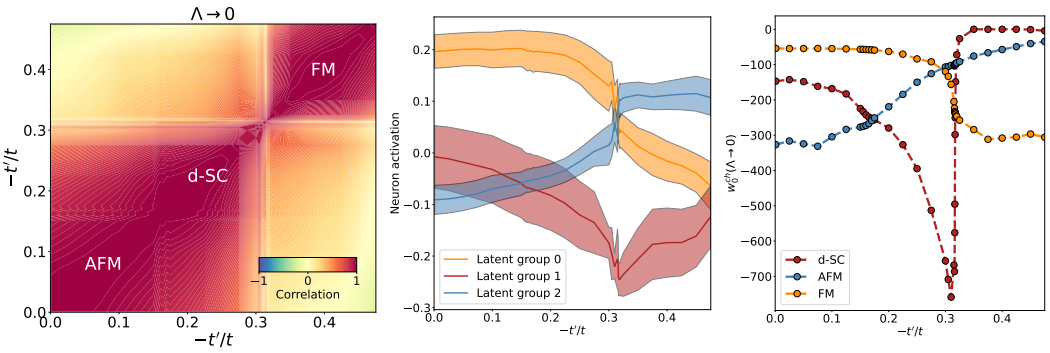

Figure 3: **Left:** Correlation matrix of the latent vectors $\mathbf{z}^\Lambda$ at different t' values for $\Lambda \to 0$. The correlation matrix is defined as the scalar product kernel $[\bar{\mathbf{z}}(t'_1) \cdot \bar{\mathbf{z}}(t'_2)]$ between normalized latent variables $\bar{\mathbf{z}}$, where t'_1 and t'_2 are any of the $0 \leq -t'/t < 0.5$. Red (blue) features correspond to a high (low) degree of statistical correlation. **Middle:** Three-groups K-Means clustering [30] of neuron activation in $\mathcal{K}_\beta$ as a function of t' for $\Lambda \to 0$. **Right:** The leading channel eigenvalues $w_0^{ch}(\Lambda)$ as functions of t' at limit of validity of one-loop approximation.

with a latent space dimension of 32, but are robust against the use of either smaller or larger values, as we show in Ref. [21]. The NODE then defines a differential equation propagation rule for latent variables in Λ. Finally, at each step of the flow, a decoder network is employed to map the latent representation $\mathbf{z}^\Lambda$ to a reconstructed 4-point vertex function $\hat{\mathrm{V}}^\Lambda(\mathrm{k}_1, \mathrm{k}_2, \mathrm{k}_3)$. We use mean squared error (MSE) between $\hat{\mathrm{V}}^\Lambda$ and V^Λ, in conjunction with a gradient-based optimizer [31]. All three networks, $\mathcal{E}_\alpha, \mathcal{K}_\beta$, and, $\mathcal{D}_\gamma$, are optimized simultaneously.

The learned dynamics successfully captures the final instability for the entire range of next-nearest neighbour couplings t'. Satisfactory prediction of the qualitative features of the vertex data is achieved in the limit $\Lambda \to 0$, as presented in Fig. 2. More interestingly, Fig. 3 shows that, during the fRG dynamics in the latent space, three highly statistically correlated latent space representations $\mathbf{z}$ emerge as a learned feature of the NODE neural network. At $\Lambda = \Lambda_0$, a first classification task is performed by the encoder $\mathcal{E}_\alpha$, which produces highly-correlated latent variables according to the value of t'. The NODE $\mathcal{K}_\beta$ takes it over to finite RG-time $\ln \Lambda_0/\Lambda$, and during the final stages of the fRG evolution in Λ, three markedly correlated areas appear, as shown in Fig. 3 (left).

The boundaries of these three regions roughly coincide with the values of t' at which the fRG predicts a change in the leading two-particle instability [23]. It is also interesting to notice that while the anti-ferromagnetic (AFM) and $d-$wave SC areas show similar normalized $\bar{\mathbf{z}}$ and are thus well-aligned in the latent space (the scalar product kernel $[\bar{\mathbf{z}}(t'_1) \cdot \bar{\mathbf{z}}(t'_2)] \sim 1$), reflecting their common origin in the dominant spin-fluctuations, the FM region stands on its own, separated from the other two phases by either a quantum critical point [23] or a first order transition [32]. The neural network distinguishes between these three many-body regimes by learning specific low-dimensional hidden representations. This is accomplished by activating three different groups of neurons in $\mathcal{K}_\beta$ as a function of t', as shown in Fig. 3. Each instability ground state corresponds indeed to a specific pattern of active neurons. This is manifestly evident when the neuron activation patterns of Fig. 3 (middle) are contrasted to Fig. 3 (right), where we show the dependence on t' of the most negative eigenvalues $w_0^{ch}(\Lambda)$ of the fRG channel-couplings $W^{\Lambda,ch}(\mathrm{k}_1, \mathrm{k}_2) = \sum_i w_i^{ch}(\Lambda) f_i^{ch}(\mathrm{k}_1)^* f_i^{ch}(\mathrm{k}_2)$ [26], with channels $ch =$ AFM, SC, FM. These leading eigenvalues are the ones associated with the highest ordering temperature T_c for their specific channel [26].

4 Outlook

Our work presents an application of artificial intelligence to fRG, which successfully unveils a dimensionality reduced dynamics for the Hubbard model on a square lattice at specific sets of electron filling. Nonetheless, the relevance of the procedure outlined in our paper goes beyond the testbed cases considered. In particular, the identification of how to extract and manipulate relevant informa-

tion encoded in the 4-point (or two-particle) vertex functions of many-electron problems, separating it from the non-relevant ones, represents a goal of utmost importance for the success of several cutting-edge methods for quantum materials [25, 33–37]. Numerical manipulation and storage of these vertex functions requires large-scale memory and computational resources, forming bottleneck for applications and necessitating learning of compressed representations of such vertex functions.

For this purpose, it will be important to explore whether transfer learning [38] could mitigate the burden of training deep nets for similar Hamiltonians. More generally, our approach, with the essential computations depending only on a "small" number of internal variables of the network, is expected to scale better than conventional implementation of most many-body algorithms. Thus, it might help us access currently hard-to-reach parts of phase diagrams of quantum many-body systems.

5 Acknowledgments

The authors are grateful to Ronny Thomale for providing the Fortran N-patch fRG code used to generate the ground truth data, and to Daniel Springer and Tobias Müller for stimulating discussions. 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 (BITMAP). This work is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through Project-ID 258499086 - SFB 1170 and through the Würzburg-Dresden Cluster of Excellence on Complexity and Topology in Quantum Matter-ct.qmat Project- ID 390858490 - EXC 2147 as well as the Austrian Science Fund (FWF) through the project I 2794-N35. MM acknowledges support from the CCQ graduate fellowship in computational quantum physics. The Flatiron Institute is a division of the Simons Foundation.

A Broader impact statement

Since our work tackles an abstract physics problem, it is difficult to estimate any societal impacts that are not already associated with underlying models (Neural ODEs and Convolutional Neural Networks) or AI in general (e. g. environmental impacts).

Regardless, as noted in the main text, we conjecture that large compression of high-dimensional vertex functions can reduce overall compute and storage needed for related calculations. Such reduction in required computational resources can contribute to energy savings for large-scale calculations and, consequently, to reduction of overall emissions associated with high-performance computation.

B Computational resources

The code and the corresponding license needed to reproduce results in this work is available by clicking through the following URL: github.com/BITMAPdds/NeuralFRG. Numerical simulations were carried out using PyTorch [39], (BSD), NumPy [40] (BSD), SciPy [41] (BSD). Plots were generated using Matplotlib [42] (PSF). Data storage and loading was handled in the HDF5 format [43, 44] (BSD).

The Fortran N-patch fRG code used to produce ground-truth data was obtained through private communication with Ronny Thomale.

The amount of compute used is difficult to report accurately because the data wasn't collected during the process itself. In total, we estimate that the results presented in this paper used approximately 72 GPU-hours for neural-network training and 12 CPU-hours for fRG data generation. These resources were used at a local cluster belonging to the authors' institutions.

References

[1] Wilson, K. G. & Kogut, J. G. The renormalization group and the ϵ expansion. *Phys. Rep.* **12**, 75 (1974).

[2] Wilson, K. G. The renormalization group: Critical phenomena and the kondo problem. *Rev. Mod. Phys.* **47**, 773 (1975).

[3] Fisher, M. E. *Scaling, Universality and Renormalization Group Theory* (Springer, Berlin, 1983).

[4] Shankar, R. Renormalization-group approach to interacting fermions. *Rev. Mod. Phys.* **66**, 129 (1994).

[5] Polchinski, J. Effective field theory and the fermi surface (1992). URL https://arxiv.org/abs/hep-th/9210046.

[6] Wetterich, C. Exact evolution equation for the effective potential. *Phys. Lett. B* **301**, 90 (1993).

[7] Morris, T. R. The exact renormalisation group and approximate solutions. *Int. J. Mod. Phys. A* **9**, 2411 (1994).

[8] Salmhofer, M. & Honerkamp, C. Fermionic Renormalization Group Flows: Technique and Theory. *Progr. Theoret. Phys.* **105**, 1 (2001).

[9] Kopietz, P. & Busche, T. Exact renormalization group flow equations for nonrelativistic fermions: Scaling toward the fermi surface. *Phys. Rev. B* **64**, 155101 (2001).

[10] Honerkamp, C., Salmhofer, M., Furukawa, N. & Rice, T. M. Breakdown of the landau-fermi liquid in two dimensions due to umklapp scattering. *Phys. Rev. B* **63**, 035109 (2001).

[11] Wentzell, N. *et al.* High-frequency asymptotics of the vertex function: Diagrammatic parametrization and algorithmic implementation. *Phys. Rev. B* **102**, 085106 (2020).

[12] Kaye, J., Chen, K. & Parcollet, O. Discrete lehmann representation of imaginary time green's functions. *Phys. Rev. B* **105**, 235115 (2022).

[13] Hille, C. *et al.* Quantitative functional renormalization group description of the two-dimensional hubbard model. *Phys. Rev. Research* **2**, 033372 (2020).

[14] Hille, C., Rohe, D., Honerkamp, C. & Andergassen, S. Pseudogap opening in the two-dimensional hubbard model: A functional renormalization group analysis. *Phys. Rev. Research* **2**, 033068 (2020).

[15] Husemann, C. & Salmhofer, M. Efficient parametrization of the vertex function, Ω scheme, and the $t, t^{'}$ hubbard model at van hove filling. *Phys. Rev. B* **79**, 195125 (2009).

[16] Wang, W.-S. *et al.* Functional renormalization group and variational monte carlo studies of the electronic instabilities in graphene near $\frac{1}{4}$ doping. *Phys. Rev. B* **85**, 035414 (2012).

[17] Lichtenstein, J. *et al.* High-performance functional renormalization group calculations for interacting fermions. *Computer Physics Communications* **213**, 100 (2017).

[18] Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602 (2017).

[19] Carleo, G. *et al.* Machine learning and the physical sciences. *Rev. Mod. Phys.* **91**, 045002 (2019).

[20] Chen, R. T. Q., Rubanova, Y., Bettencourt, J. & Duvenaud, D. Neural Ordinary Differential Equations. *NIPs* **109**, 31 (2018).

[21] Di Sante, D. *et al.* Deep learning the functional renormalization group. *Phys. Rev. Lett.* **129**, 136402 (2022).

[22] Quarteroni, A. & Rozza, G. (eds.) *Reduced Order Methods for Modeling and Computational Reduction* (Springer, Cham, 2014).

[23] Honerkamp, C. & Salmhofer, M. Magnetic and superconducting instabilities of the hubbard model at the van hove filling. *Phys. Rev. Lett.* **87**, 187004 (2001).

[24] Honerkamp, C. & Salmhofer, M. Temperature-flow renormalization group and the competition between superconductivity and ferromagnetism. *Phys. Rev. B* **64**, 184516 (2001).

[25] Metzner, W., Salmhofer, M., Honerkamp, C., Meden, V. & Schönhammer, K. Functional renormalization group approach to correlated fermion systems. *Rev. Mod. Phys.* **84**, 299 (2012).

[26] Platt, C., Hanke, W. & Thomale, R. Functional renormalization group for multi-orbital fermi surface instabilities. *Advances in Physics* **62**, 453 (2013).

[27] Kopietz, P., Bartosch, L. & Schütz, F. *Introduction to the Functional Renormalization Group* (Springer-Verlag, Berlin Heidelberg, 2010).

[28] Goodfellow, I., Bengio, Y. & Courville, A. *Deep learning* (MIT press, 2016).

[29] Lee, K. & Parish, E. J. Parameterized neural ordinary differential equations: applications to computational physics problems. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **477**, 20210162 (2021).

[30] Lloyd, S. Least squares quantization in pcm. *IEEE Transactions on Information Theory* **28**, 129 (1982).

[31] Kingma, D. P. & Ba, J. L. Adam: A method for stochastic optimization. In *3rd Int. Conf. Learn. Represent. ICLR 2015 - Conf. Track Proc.* (International Conference on Learning Representations, ICLR, 2015).

[32] Lin, H. Q. & Hirsch, J. E. Two-dimensional hubbard model with nearest- and next-nearest-neighbor hopping. *Phys. Rev. B* **35**, 3359 (1987).

[33] Bickers, N. E. In Sénéchal, D., Tremblay, A. & Bourbonnais, C. (eds.) *Theoretical Methods for Strongly Correlated Electrons* (Springer, New York, 2004). Chap. 6.

[34] Kugler, F. B. & von Delft, J. Multiloop functional renormalization group that sums up all parquet diagrams. *Phys. Rev. Lett.* **120**, 057403 (2018).

[35] Kiese, D., Müller, T., Iqbal, Y., Thomale, R. & Trebst, S. Multiloop functional renormalization group approach to quantum spin systems. *Phys. Rev. Research* **4**, 023185 (2022).

[36] Taranto, C. *et al.* From infinite to two dimensions through the functional renormalization group. *Phys. Rev. Lett.* **112**, 196402 (2014).

[37] Rohringer, G. *et al.* Diagrammatic routes to nonlocal correlations beyond dynamical mean field theory. *Rev. Mod. Phys.* **90**, 025003 (2018).

[38] Tan, C. *et al.* A survey on deep transfer learning. In *International conference on artificial neural networks*, 270–279 (Springer, 2018).

[39] Paszke, A. *et al.* Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems* **32** (2019).

[40] Harris, C. R. *et al.* Array programming with NumPy. *Nature* **585**, 357 (2020).

[41] Virtanen, P. *et al.* SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261 (2020).

[42] Hunter, J. D. Matplotlib: A 2D graphics environment. *Comput. Sci. Eng.* **9**, 99–104 (2007).

[43] The HDF Group. Hierarchical data format, version 5 (1997-NNNN). URL https://www.hdfgroup.org/HDF5/.

[44] Collette, A. *Python and HDF5* (O'Reilly, 2013).

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] We note and discuss that these results hold for the specific dataset and physical model under consideration here.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See section A.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code, data and instructions are included in Section B.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [No] There was no space due to the strict 4-page limit and strong discouragement of appendices. However, we have prepared supplementary material that can be attached to the final version, if accepted.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] Due to the numerical cost of training the model, we present only a single trained model in this paper. We have performed experiments investigating the effect of different latent space dimensions and random but have had to omit them due to limited space. However, we do offer an open-source code repository that can be used to reproduce the results and investigate effects of any tweaks.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] The the total amount of compute and the type of resources used are included in Section B.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See Sec. B.
   (b) Did you mention the license of the assets? [Yes] See Sec. B.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We have provided a URL to the open-source code repository we used to produce results presented in this work.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] See Sec. B.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]