

Exploring Group Convolutional Networks for Sign Problem Mitigation via Contour Deformation

Christoph Gäntgen,^{a,b,c,*} Thomas Luu^{a,c} and Marcel Rodekamp^d

^a*Institute for Advanced Simulation, Forschungszentrum Jülich, 54245 Jülich, Germany*

^b*Center for Advanced Simulation and Analytics, Forschungszentrum Jülich, 52425 Jülich, Germany*

^c*Helmholtz-Institut für Strahlen- und Kernphysik, Rheinische Friedrich-Wilhelms-Universität Bonn, 53115 Bonn, Germany*

^d*Institut für Theoretische Physik, Universität Regensburg, 93040 Regensburg, Germany*

E-mail: c.gaentgen@fz-juelich.de

The sign problem that arises in Hybrid Monte Carlo calculations can be mitigated by deforming the integration manifold. While simple transformations are highly efficient for simulation, their efficacy systematically decreases with decreasing temperature and increasing interaction. Machine learning models have demonstrated the ability to push further, but require additional computational effort and upfront training. While neural networks possess the capacity to learn physical symmetries through proper training, there are anticipated advantages associated with encoding them into the network's structure. These include enhanced accuracy, accelerated training, and improved stability. The objective of the present study is twofold. First, we investigate the benefits of group convolutional models in comparison to fully connected networks, with a specific focus on the effects on the sign problem and on computational aspects. Second, we examine their capabilities for transfer learning, demonstrating the ability to further reduce training cost. We perform our investigations on the Hubbard model on select low-dimensional systems.

*The 41st International Symposium on Lattice Field Theory (LATTICE2024)
28 July - 3 August 2024
Liverpool, UK*

*Speaker

1. Introduction

The Hubbard model, introduced in [1], is a very well known condensed matter model used to describe electronic properties of metals. By combining tight binding with onsite interaction on a fixed spatial lattice, it creates a system of strongly correlated electrons.

$$H = - \sum_{x,y} \kappa_{x,y} \left(a_{x\uparrow}^\dagger a_{y\uparrow} + a_{x\downarrow}^\dagger a_{y\downarrow} \right) - \frac{U}{2} \sum_x (n_{x\uparrow} - n_{x\downarrow})^2 - \mu \sum_x (n_{x\uparrow} + n_{x\downarrow}) \quad (1)$$

Equation 1 consists of the tight-binding term with the hopping matrix $\kappa_{x,y}$, the interaction term with coupling parameter U and an optional chemical potential μ . Here a , a^\dagger and n are the electron-annihilation, -creation and -number operators acting on a specified spatial site and spin orientation. This hamiltonian can be transformed into the action

$$S = \sum_{x,t} \frac{\phi_{x,t}^2}{2\tilde{U}} - \log \det(M[\phi, \tilde{\kappa}, \tilde{\mu}] M[-\phi, -\tilde{\kappa}, -\tilde{\mu}]) \in \mathbb{C} \quad (2)$$

by Hubbard-Stratonovich transformation [2, 3]. Here the auxiliary scalar field ϕ is introduced alongside the fermion matrix M . The integral over imaginary time is discretized into N_t slices. As a result the physical parameters are rescaled ($\tilde{U} = U \times \delta$) by the temporal lattice spacing $\delta = \beta/N_t$, which depends on the inverse temperature β . The fermionic part of eq. 2 is generally complex valued with the exception of bipartite lattices at $\mu = 0$. Therefore expectation values obtained via the typical path integral formalism require additional care,

$$\langle \hat{O} \rangle = \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi \hat{O}[\phi] e^{-S[\phi_n]} \approx \frac{1}{N} \sum_{n=0}^N \hat{O}[\phi_n] \text{ with } \phi_n \sim e^{-S[\phi_n]}/\mathcal{Z} \quad (3)$$

As indicated in eq. 3 the standard procedure is to stochastically approximate the high dimensional integral with field configurations ϕ_n sampled from a probability distribution $e^{-S[\phi_n]}/\mathcal{Z}$. For complex valued actions reweighting is required to correct the otherwise complex probability distribution,

$$\langle \hat{O} \rangle = \frac{\langle \hat{O} e^{-iS_I} \rangle_R}{\langle e^{-iS_I} \rangle_R} \approx \frac{\sum_{n=0}^N \hat{O}[\phi_n] e^{-iS_I[\phi_n]}}{\sum_{n=0}^N e^{-iS_I[\phi_n]}}. \quad (4)$$

The complex phase is included in the observable and the expectation value is sampled with respect to the real part of the action. Though formally exact, the issue arises numerically from limited statistics. When the average phase is small the expectation value is more susceptible to statistical fluctuations. The sign problem refers to the slowed convergence of stochastic estimates due to the cancellation of complex phases/signs in the denominator [4]. The absolute value of this denominator is used as a measure for the severity of the sign problem, which we call the *statistical power* $\Sigma = |\langle e^{-iS_I} \rangle_R|$.

The separation of real and imaginary parts introduces non-holomorphicity in the RHS of eq. 4, meaning while the observable is a fixed value, the expectation values in the numerator and denominator are contour dependent. This contour deformation dependence provides a possible way to reduce the sign problem. In [5–9] contour deformation was applied successfully, making otherwise infeasible calculations due to the sign problem possible. Many of these works were facilitated by machine learning methods. This work explores the proper utilization of physical symmetries by

directly designing the symmetries within the network. We train the network on an approximation of the contributing Lefschetz thimbles [10], which is reached by integrating holomorphic flow equations from randomly sampled ϕ for finite flow time [11]. The network is trained to parametrize the imaginary part of the resulting manifold from the real part,

$$\phi^* = \phi_R + i\mathcal{N}(\phi_R). \quad (5)$$

Here \mathcal{N} stands for the neural network. The symmetries of our desired manifold follow from the properties of eq. 2:

- spatial geometry (translation, rotation, mirroring)
- temporal (translation, reversal)
- complex conjugate (real valued sign flip) .

Thus the ideal network is invariant under real valued sign flips and equivariant under the proper permutations of field components. The lattices we investigate in this work are an equilateral triangle and a square, combined with $N_t = 16$ time slices. In all cases we will consider the complex conjugation and temporal translation. On top of that we look at the 120° rotations of the triangle, resulting in a $3 \times 16 \times 2$ fold simplification. For the square we consider all spatial symmetries (except for the translation as it is redundant in this case), given by one mirror axis and 90° rotations. Resulting in up to $2 \times 4 \times 16 \times 2$ fold symmetry. The temporal reversal is not used here for practical reasons.

2. Group Equivariant Convolutional Neural Networks

For the translations of a square lattice a standard convolutional neural network (CNN) would suffice, but for more general symmetries their functionality has to be expanded. Cohen and Welling introduced exactly this expanded functionality, the group equivariant convolutional neural network (G-CNN), in [12]. They applied it to rotated images from CIFAR10 and MNIST achieving significantly better results than comparable CNNs of that time. With the additional symmetries accounted for, their G-CNNs reached a higher degree of weight sharing and as a consequence increased expressive capacity. For a detailed explanation we refer the reader to the original publication [12]. Here we only provide a brief overview. The most salient requirement is that the convolution operation commutes with the group symmetry transformations, thus ensuring *equivariance*. Combined sequentially with other equivariant operations or layers, a fully group convolutional neural network can be created. Equivariance is a property similar to invariance, when the input gets transformed, the output transforms accordingly.

Within the network, kernels (filters) are transformed under the group G , resulting in a stack of feature maps. Acting on the input with a transformed feature map is equivalent to acting on the inversely transformed input with the original feature map. Each element of the stack is associated with a group element, e.g. the $C4$ group creates a stack associated to e, r, r^2 and r^3 . This does not create additional parameters, each transformation shares the original parameters. The principle is the same as a conventional convolution layer creating an output value for each *translation* of the kernel.

3. Results

In this section we show the learning capabilities of different networks on flow generated training data. As suggested in equation 1 we use the real parts of flowed configurations as input and evaluate with a loss function, how well the network reproduced the imaginary parts. The loss values in this section refer to the mean squared error loss function evaluated on the validation set. For each network we created 24 randomly initialized samples. We find that the equivariant models are more susceptible to getting stuck in local minima during the training process, a phenomenon we had encountered before when experimenting with simple parametrizations, such as an elementwise Gaussian. To both showcase and counteract this, we show the minimum and maximum loss values out of all the samples.

3.1 Triangle

The triangle can be thought of as a linear grid of 3 points with periodic boundary conditions. Thus the rotations are equivalent to translation in this representation. The general N -D spatial square lattice combined with one temporal dimension form a $N + 1$ -D square grid. Those grids work with the default implementation of convolutional layers that are known for their application to images. Similar to images and videos, the spatial and temporal structure of field configurations has meaning. When using linear networks the information of which components are adjacent to each other is lost. In this subsection we compare selected architectures of convolutional networks with a fully connected linear network. Figure 1 shows a comparison of different trained models




Figure 1: **Left:** Final validation loss depending on amount of training configurations. **Right:** Final validation loss depending on tunable parameter count. The naming scheme for the labels uses C for a convolution layer and N for a layer in which only the channels interact with each other. Note that the error bars show the minimum and maximum results of the 24 samples, not the standard deviation.

depending on available training data and size of the model. The system is a triangle at $N_t = 16$, $\beta = 8$ and $U = 3$. Due to the odd number of sites it is non bipartite, hence it suffers from a sign problem. As expected the convolutional model has an advantage over the linear model in the case of limited training data. Also they need significantly fewer parameters to achieve the same performance. Evaluating the statistical power on those models gives the results in table 1. The integers behind the model name stand for the width of the model, in case of convolutional models this is the number of channels. Two integers separated with a colon stand for two intermediate layers

Model	Parameters	Statistical power
Fully-Connected 96	9360	0.158 ± 0.017
Convolutional 32	1697	0.117 ± 0.021
Convolutional $16_{3 \times 3}$	979	0.129 ± 0.018
Convolutional 32 : 32	2753	0.230 ± 0.020

Table 1: Statistical power evaluation of selected models.

and the subscript stands for a second smaller convolution kernel. For comparison we looked at two constant shifts, that we explored in depth in an earlier publication [13] and successfully applied in the calculation of relevant molecules [14]. Setting all imaginary parts to a constant, can greatly improve the sign problem. The offset that maximizes the statistical power came out to $\Sigma = 0.13$. Some models exceeded this value, thereby raising optimism for challenging sign problems where the optimized shift is insufficient. Further there is a trivial shift to the plane that is tangent to the critical point of the main Lefschetz Thimble. We can determine it analytically, it leads to $\Sigma = 0.03$. We consider $\Sigma > 0.1$ to be measurable with reasonable resources, showing once more that optimizing the offset expands measurable parameter space. Once this one parameter is optimized, it takes more advanced transformations to alleviate the sign problem.

3.2 Square

The square trivially fits a three dimensional space-time grid. We applied the $C4$ symmetry group, made up from rotations only. And the $D4$ group, containing reflections as well. As this lattice is bipartite it does not suffer from a sign problem at zero chemical potential. We induced a sign problem that is comparable to the one of the square lattice by introducing a chemical potential. Figure 2 shows the validation loss values for the four site square with $N_t = 16$ time slices, $\beta = 8$,

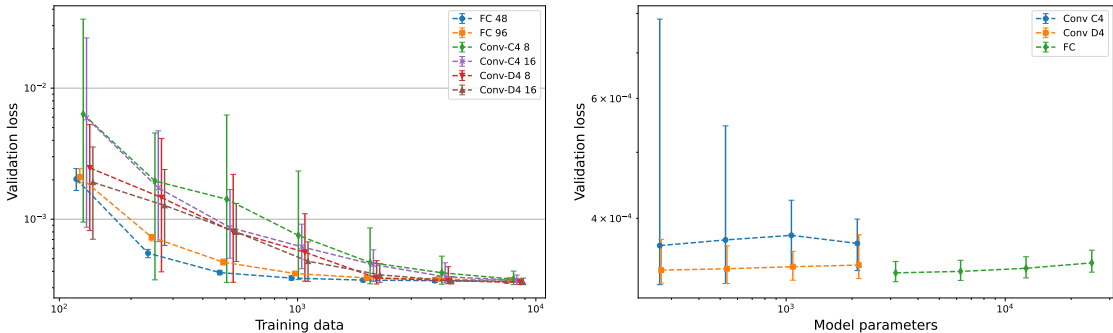


Figure 2: Left: Final validation loss depending on amount of training configurations. Right: Final validation loss depending on tunable parameter count. Note that the error bars show the minimum and maximum results of the 24 samples, not the standard deviation.

$U = 3$ and $\mu = 1.5$. Looking at the magnitude of the validation loss, it is clear that all the networks learned to reproduce the training data quite well. Unfortunately for the chosen system the generation of training data was extremely difficult, so the quality of which became a limiting factor. Still we see that the convolutional model performs as good as the fully connected network while requiring

a fraction of the parameters. Further the network utilizing the $D4$ group with both rotations and reflections, is slightly ahead of the one using rotations only.

4. Transfer-Learning

A major motivation for using convolution based networks was its ability to vary the input and output dimensions. The fully connected network is created for fixed input and output dimensions, so there is no obvious way to reuse a trained model for another system size. This would be of great help for larger systems, such as graphene sheets or carbon nanotubes. Similar to the physical behavior scaling with the system size, we expected find a transferable transformation for the field components. After some testing it became clear, that this transfer between lattice dimensions would not work, as the performance of a network declines significantly when it does not see the whole input at once. Hence the networks in this paper use a convolution kernel that covers the whole input, limiting the advantages over a fully connected model to the symmetry properties only.

Nonetheless as the symmetric networks have proven to require less training data, they pose as promising candidates for classical transfer learning. Assuming that the Lefschetz thimbles (or more generally the action values $S(\phi)$) change smoothly with the model parameters β and U , we can assume that the trained contour deformation for one system, is a beneficial starting point for training data generation for nearby system parameters. This introduces the cost of transforming each configuration in the process of generating training data, so we are facing a trade off between quantity and quality (presumably). Figure 3 shows both linear and convolutional models that were




Figure 3: **Left:** Statistical power for β transfer learning at $U = 3$. **Right:** Statistical power for U transfer learning at $\beta = 8$. The system is a triangle with $N_t = 16$.

initially trained on standard training data (2^{15} samples) for one β or U . Then from the learned transformation new training data was generated (2^{12} samples). The model was then retrained on the new data, without resetting the already learned parameters. With few exceptions the transfer learned models could not keep up with the original training method. We observe though that generally the smaller transfer steps are more successful than further steps (for example $\beta = 7 \rightarrow 8$ is better than $\beta = 6 \rightarrow 8$), so step size clearly matters. We conclude that the applicability of this method is situational. If the task at hand is a fine analysis of a range of parameters this method can save resources and promote stability.

Further we looked at the same principle to iteratively refine one model, by repeatedly generating training data from the transformation of a trained model with holomorphic flow and retraining

the model. We find that this approach has the potential to generate very good training data and models, but is unreliable. Considering the additional cost from calculating the transformations, it is preferable to start from an offset and flow further with higher precision.

5. Summary & Outlook

In this proceeding we revisited a method for reducing the Hubbard model Sign problem via contour deformation that has been successfully applied in previous works [5–7, 9, 11, 13]. A neural network learns to parametrize an optimized manifold obtained by computationally expensive holomorphic flow. The original work featured fully connected neural networks which we were now able to surpass with convolutional networks. They are connected to the structure of the inputs, ensuring and leveraging the symmetries of the physical system. We observed that these *informed* models excel for limited training data availability and require less tunable parameters. Further it drastically limits the risk of introducing unphysical biases, by guaranteeing that the output parametrization follows the proper symmetries.

We investigated these models also with the expectation, they would excel at transfer learning tasks too. However we could not confirm this assumption. Generating the training data from a related system does make sense intuitively, but introduces a layer of variability to the process of generating training data. Thus we found it had the potential to outperform our established practices, but did so unreliably.

In the future we plan to implement the symmetries of larger carbon nano structures, such as fullerenes. The major weakness of a fully connected network was its scaling. With symmetric model architectures we hope to limit this issue, especially for systems with lots of symmetries to leverage.

Acknowledgements

We gratefully acknowledge the computing time granted by the JARA Vergabegremium and provided on the JARA Partition part of the supercomputer JURECA at Forschungszentrum Jülich. Both C.G. and T.L. were support in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the CRC 1639 NuMerIQS–project no. 511713970. M.R. were supported through MKW NRW under the funding code NW21-024-A.

References

- [1] J. Hubbard, *Calculation of Partition Functions*, *Phys. Rev. Lett.* **3** (1959) 77.
- [2] R.L. Stratonovich, *On a Method of Calculating Quantum Distribution Functions*, *Sov. Phys. Dokl.* **2** (1957) 416.
- [3] T. Luu and T.A. Lähde, *Quantum Monte Carlo Calculations for Carbon Nanotubes*, *Phys. Rev. B* **93** (2015) 155106 [1511.04918].
- [4] E.Y. Loh, J.E. Gubernatis, R.T. Scalettar, S.R. White, D.J. Scalapino and R.L. Sugar, *Sign problem in the numerical simulation of many-electron systems*, *Phys. Rev. B* **41** (1990) 9301.
- [5] J.-L. Wynen, E. Berkowitz, S. Krieg, T. Luu and J. Ostmeyer, *Leveraging Machine Learning to Alleviate Hubbard Model Sign Problems*, *Phys. Rev. B* **103** (2021) 125153 [2006.11221v1].
- [6] A. Alexandru, P.F. Bedaque, H. Lamm, S. Lawrence and N.C. Warrington, *Fermions at Finite Density in 2+1 Dimensions with Sign-Optimized Manifolds*, *Phys. Rev. Lett.* **121** (2018) 191602 [1808.09799].
- [7] M. Cristoforetti, F. Di Renzo, A. Mukherjee and L. Scorzato, *Monte Carlo simulations on the Lefschetz thimble: Taming the sign problem*, *Phys. Rev.* **D88** (2013) 051501 [1303.7204].
- [8] W. Detmold, G. Kanwar, M.L. Wagman and N.C. Warrington, *Path integral contour deformations for noisy observables*, *Phys. Rev. D* **102** (2020) 014514 [2003.05914].
- [9] M. Rodekamp, E. Berkowitz, C. Gäntgen, S. Krieg, T. Luu and J. Ostmeyer, *Mitigating the Hubbard Sign Problem with Complex-Valued Neural Networks*, *Phys. Rev. B* **106** (2022) 125139 [2203.00390v2].
- [10] S. Lefschetz, *On Certain Numerical Invariants of Algebraic Varieties with Application to Abelian Varieties*, *Trans. Amer. Math. Soc.* **22** (1921) 327.
- [11] A. Alexandru, G. Basar, P.F. Bedaque, G.W. Ridgway and N.C. Warrington, *Monte Carlo calculations of the finite density Thirring model*, *Phys. Rev. D* **95** (2017) [1609.01730v2].
- [12] T.S. Cohen and M. Welling, *Group Equivariant Convolutional Networks*, *Proceedings of the International Conference on Machine Learning (ICML), 2016* (2016) [1602.07576].
- [13] C. Gäntgen, E. Berkowitz, T. Luu, J. Ostmeyer and M. Rodekamp, *Fermionic Sign Problem Minimization by Constant Path Integral Contour Shifts*, *Phys. Rev. B* **109** (2023) [2307.06785v1].
- [14] M. Rodekamp, E. Berkowitz, C. Gäntgen, S. Krieg, T. Luu, J. Ostmeyer et al., *Single Particle Spectrum of Doped C₂₀H₁₂-Perylene*, 2406.06711.