Sign Structure of Many-Body Wavefunctions and Machine Learning

Tom Westerhout, Andrey A. Bagrov,[*] Mikhail I. Katsnelson
Institute for Molecules and Materials, Radboud University,
6525 AJ Nijmegen, The Netherlands
t.westerhout@student.ru.nl
a.bagrov@science.ru.nl
m.katsnelson@science.ru.nl

Nikita Astrakhantsev[†]
Physik-Institut, Universität Zürich,
Winterthurerstrasse 190, CH-8057 Zürich, Switzerland
nikita.astrakhantsev@physik.uzh.ch

Konstantin Tikhonov[‡]
Institut für Nanotechnologie, Karlsruhe Institute of Technology,
76021 Karlsruhe, Germany
tikhonov@itp.ac.ru

Following fascinating success in image and speech recognition tasks, machine learning (ML) methods have recently proven themselves to be very useful in physical sciences. For example, ML has been used to classify phases of matter [1], to enhance quantum state tomography [2], to bypass expensive dynamic ab initio calculations [3], and much more [4]. The ML ideas also helped to take a novel and fresh look at the concept of variational Monte Carlo [5]. A simple yet very generic variational ansatz that inherits the structure of a certain neural network (Restricted Boltzmann Machine) was suggested. For the test cases of one- and two-dimensional Heisenberg and transverse field Ising models, it was demonstrated that, optimizing this ansatz with the Stochastic Reconfiguration (SR) scheme [6], one can achieve high accuracy in approximating ground states of systems of up to hundreds of spins, sometimes outperforming the state-of-the-art methods.

Although a variational wave function with a network structure may be able to approximate the ground state really well, in some cases the desired point in the space of variational parameters can be hard to reach, and learning algorithm hits a saddle point before approaching the solution. This results in a large relative energy error and a low overlap between the NQS and the actual exact ground state, making the obtained solution almost useless for computing physical observables such as electric conductivity or spin-spin correlation functions. This problem is particularly pronounced for systems where the energy gap between the ground state and the first excited state is very small, like for frustrated spin systems such as $J_1 - J_2$ antiferromagnetic Heisenberg model on square lattice [7], or the Fermi-Hubbard model away from the neutrality point [8].

So far, significant effort has been put into the search for neural quantum states architectures that have good expressibility, – a potential capacity to represent a many-body wave function with high accuracy using a moderate number of parameters [9, 10, 11]. At the same time, there is another issue that is

[*] Also: Department of Physics and Astronomy, Uppsala University, Box 516, SE-75120 Uppsala, Sweden

[†] Also: Moscow Institute of Physics and Technology, 141700 Dolgoprudny, Russia and Institute for Theoretical and Experimental Physics NRC "Kurchatov Institute", Moscow, 117218 Russia

[‡] Also: Skolkovo Institute of Science and Technology, 143026 Skolkovo, Russia and Landau Institute for Theoretical Physics RAS, 119334 Moscow, Russia

Second Workshop on Machine Learning and the Physical Sciences (NeurIPS 2019), Vancouver, Canada.

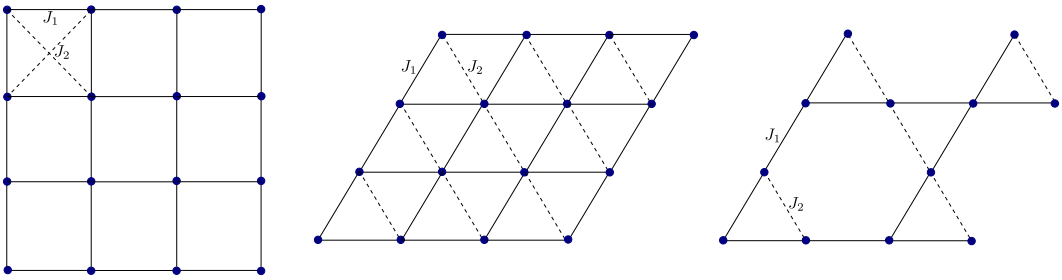

Figure 1: The three studied cases of the frustrated antiferromagnetic Heisenberg model: next-nearest neighbor $J_1 - J_2$ model on a square lattice (left), nearest-neighbor model on an anisotropic triangular lattice (middle), spatially anisotropic Kagome lattice (right). In all cases $J_2 = 0$ corresponds to the absence of frustrations.

not widely discussed in this context. In the variational optimization schemes, an ansatz is adjusted iteratively in a certain way, so that we expect the system to end up in the lowest energy state allowed by the form of the ansatz [12, 6, 13]. At each step of this iterative procedure, one has to evaluate the change of the trial wavefunction parameters, induced by the evolution operator: $\psi_n(\sigma) \to \psi_{n+1}(\sigma)$. Evaluation of the NN weights describing the state $\psi_{n+1}(\sigma)$ relies on MC sampling from basis of the Hilbert space of the model, and for large systems the total number of samples is negligibly small in comparison with the dimension of the Hilbert space. Hence, it is of crucial importance for the NN to accurately *generalize* onto a larger subspace that was not sampled in the course of learning and correctly predict phases and amplitudes of the wave function on the full set of basis vectors.

While the generalization issue concerns both phases and amplitudes, it turns out that these two components of the wavefunction behave differently in this respect. Already from the first works in the field, it seemed plausible that effectiveness of NN as variational ansatz is somehow connected to the sign structure of the models. For instance, in [5], even for the unfrustrated Heisenberg antiferromagnet on a square lattice, the Hamiltonian must first be brought into stoquastic (sign-definite) form by a unitary transformation in order to reduce noise and attain proper level of convergence (see also [14]). As another example, let us note that in recent study [13] it was stressed that biasing the NQS anzats with certain predefined (heuristic) sign structures is very important for performance of the method. Therefore, although we will study both aspects, special attention will be paid to the sign structure.

1 Models of interest

We consider several antiferromagnetic spin models described by the Heisenberg Hamiltonian:

$$\hat{H} = J_1 \sum_{\langle a,b \rangle} \hat{\boldsymbol{\sigma}}_a \otimes \hat{\boldsymbol{\sigma}}_b + J_2 \sum_{\langle\langle a,b \rangle\rangle} \hat{\boldsymbol{\sigma}}_a \otimes \hat{\boldsymbol{\sigma}}_b \ , \tag{1}$$

where, for each lattice geometry, the first sum is taken over the unfrustrated sublattice (solid lines in Fig. 1), and the second sum is taken over the sublattice that brings in frustrations (dashed lines in Fig. 1). Namely, we consider $J_1 - J_2$ model on a square lattice [15, 16] and the nearest-neighbor antiferromagnetic on spatially anisotropic triangular [17] and Kagome [18, 19] lattices.

For every model, its ground state belongs to the sector of minimal magnetization, thus the dimension of the corresponding Hilbert space is $K = C^N_{[N/2]}$ (where N is the number of spins). It is convenient to work in the basis of eigenstates of $\hat{\sigma}^z$ operator: $|\mathcal{S}\rangle \sim |\uparrow\downarrow \ldots \downarrow\uparrow\rangle$. In this basis the Hamiltonian $\hat{H}$ is real-valued. The ground state is thus also real-valued, and every coefficient in its basis expansion is characterized by a sign $s_i = \text{sign}(\psi_i)$ (instead of a continuous phase):

$$|\Psi_{GS}\rangle = \sum_{i=1}^{K} \psi_i |\mathcal{S}_i\rangle = \sum_{i=1}^{K} s_i |\psi_i| |\mathcal{S}_i\rangle. \tag{2}$$

By running numerical experiments, we shall demonstrate that it is indeed the lack of generalization that prevents a neural quantum state from accurately learning the signs of the wavefunction, even

though expressibility of the corresponding ansätze could be good enough. Our strategy is to consider exact ground states of the models and test how well NN can predict sign structure of the whole state when they are given only a small fraction of it during training. We shall demonstrate that, when the models are interpolated between unfrustrated and fully frustrated regimes, networks' generalization abilities change in a non-trivial way, becoming very poor in certain cases. This motivates a search for neural quantum states architectures with better generalization capacity.

For the three models defined above, we systematically study the generalization properties of NNs (separating the signs from the amplitudes) of different architectures varying the degree of frustration (controlled by J_2/J_1) and the size of the training dataset.

2 Methods

In this study, we use feed-forward networks of three different architectures (dense 1-layer, dense 2-layer, and convolutional 2-layer) to encode wavefunction coefficients via splitting them into amplitudes and signs. All of our networks have the same input format: spin configuration $|\mathcal{S}_i\rangle = |\sigma_1\sigma_2\ldots\sigma_N\rangle$ represented as a binary sequence, $\sigma_k = \pm 1$. Network encoding amplitudes outputs a real number, natural logarithm of the amplitude. Network encoding sign structure outputs a real number $p \in [0, 1]$ interpreted as a probability for the corresponding sign to be plus. Thus, unlike the approach of [5], we represent wavefunction sign using a binary classifier. Both networks are trained on data obtained from exact diagonalization. We sample $\varepsilon_{\text{train}} \cdot K$ spin configuration from the Hilbert space basis according to probability distribution $P(i) = \frac{|\psi_i|^2}{\sum_j |\psi_j|^2}$. They constitute the training dataset. In practical applications of NQS [5, 20, 7], SR [21, 6], Stochastic Gradient Descent [22, 23], or Generalized Lanczos [12], the training dataset is generated by Monte Carlo sampling from basis of the Hilbert space of the model, and, since dimension of the latter grows exponentially with the number of spins, only a tiny fraction of it can be covered with a Monte Carlo chain in reasonable time. Therefore it is natural to mimic this incomplete coverage with $\varepsilon_{\text{train}} \ll 1$.

To assess the performance of the NNs we evaluate overlap (scalar product) between exact eigenstate and the trial state. A trial state for sign NN is defined as a state with exact amplitudes but with sign structure encoded in a NN. Analogously, a trial state for amplitude NN is obtained by superimposing the exact sign structure onto the positive amplitudes, predicted by an amplitude NN. Code to reproduce our results and more details about the training scheme can be found on GitHub [24].

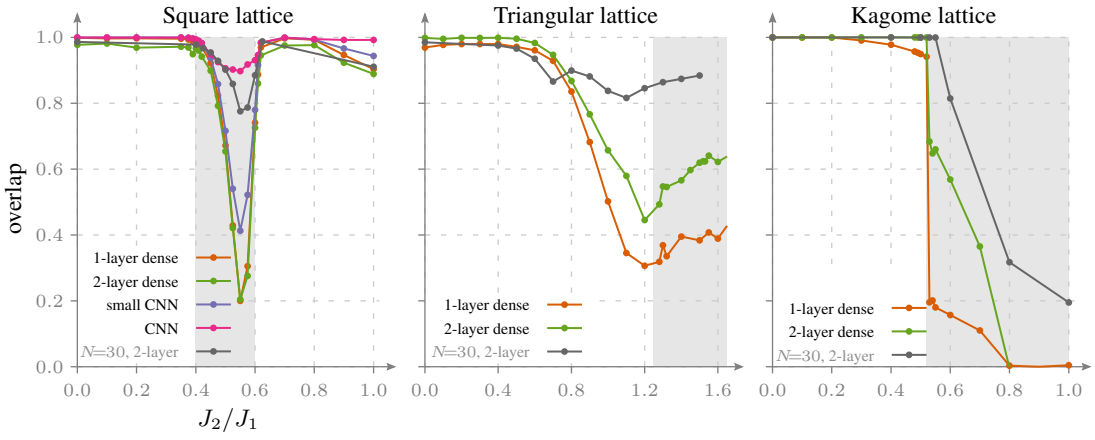

Figure 2: Overlap of the variational wave function with the exact ground state as a function of J_2/J_1 computed for the square (left), triangular (center), and Kagome (right) lattices. Overlap was computed on the rest dataset (not included into training and validation datasets). Note that generalization is poor in the frustrated regions (which are shaded on the plots). We also show some preliminary results for larger clusters.

3 Results

We have analyzed how NNs learn ground state structures of three lattice models, in each case considering periodic clusters of 24 spins[4]. Effective dimension of Hilbert space in the zero-magnetization sector is $d = C_{12}^{24} \simeq 2.7 \cdot 10^6$.

Fig. 2 shows how well a model trained on a small subset (1%) of the Hilbert space basis, predicts the sign structure on the remaining basis vectors unavailable to it during training. For all three models we see that behavior of generalization quality, albeit different, reflects very well the known phase transitions with generalization being easy in ordered phases and becoming notoriously hard in disordered phases. Note also that different neural networks may generalize very differently: in particular, as shown on the left panel of Fig. 2, dips in performance of convolutional NNs are much smaller than those for dense networks for the square lattice model. We believe that experiments of this kind would help to choose proper architectures to be used in iterative diagonalization schemes.

It is very important to distinguish the ability to represent the data from generalization. In the context of NQS, the former means that a NN is able to express complex quantum states well if training was conducted in a perfect way. We observe that perfect expressibility (overlap $= 1$) is indeed possible if the training set is large enough – meaning that the network can represent the target state very well. However, this ability does not automatically make a neural network useful if it cannot generalize well. It is important to study how generalization quality changes when size of the training dataset is increased. Results for Kagome lattice are shown on the left panel of Fig. 3. Interestingly, even in the frustrated phase ($J_2 = 0.6$) it is possible to generalize reasonably well from a relatively small subset of the basis states, but the required $\varepsilon_{\text{train}}$ becomes substantially larger than in the magnetically ordered phase. Most importantly, the ability of the NN to generalize establishes in an abrupt manner contrary to more typical smooth behavior observed in statistical models of learning [25, 26, 27].

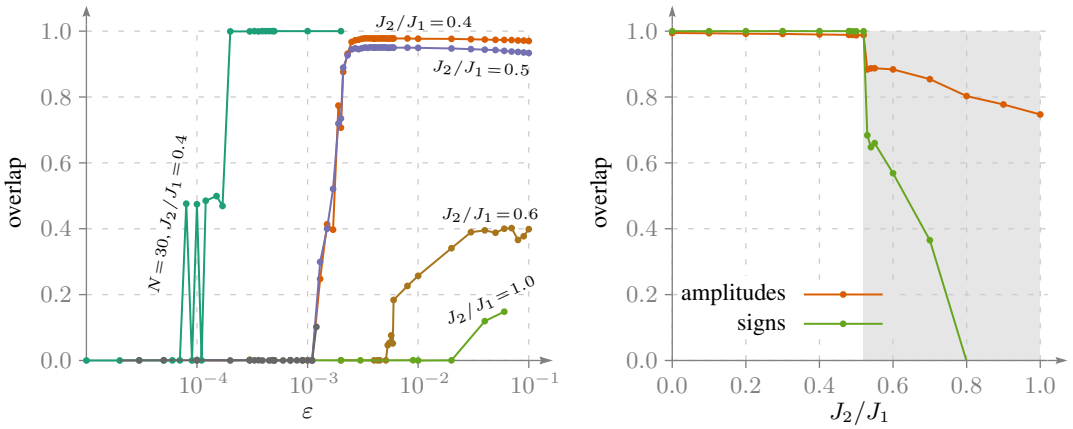

Figure 3: Left: Dependence of generalization quality measured by overlap between the variational and the exact states on the size of the training data set $\varepsilon_{\text{train}}$ for Kagome model for $J_2 = 0.4, 0.5, 0.6, 1.0$. Right: Comparison of generalization quality as measured by overlap for learning the sign structure (red) and amplitude structure (blue) for Kagome lattice for 2-layer dense architecture. Note that both curves decrease in the frustrated region, but sign structure is much harder to learn.

One may wonder whether it is indeed the signs rather than amplitudes which are responsible for the difficulty of learning the wavefunction as a whole. To prove this statement, we conduct the following analysis. In the context of learning, overlap between a trial wavefunction and the target state can be used to characterize the effectiveness of NNs in two different ways. First, one can fix the amplitudes of the wavefunction and use a NN to predict the signs. This produces a trial wavefunction ψ_{sign}. Alternatively, one can fix the sign structure, and encode the amplitudes in a NN to get a trial wavefunction ψ_{amp}. Using the Kagome lattice as an example, on the right panel of Fig. 3 we show that although prediction of both signs and amplitudes becomes harder at the point of phase transition $J_2/J_1 = 0.51$, drop in the sign curve is much larger, and at even higher J_2 the quality of the learned

[4]For some architectures, we also provide results for 30-spin clusters

states becomes too poor to approximate the target wavefunction. At the same time, even deeply in the frustrated regime generalization of amplitudes, given the exact sign structure, leads to a decent result.

4 Discussion

We have demonstrated that generalization may indeed be an essential factor that is likely responsible for spoiling the convergence of NQS in a number of physically interesting cases such as frustrated quantum spin systems. Our main observation which is qualitatively valid for all the studied models and NN architectures is that a NN fails to *generalize* the distribution of signs in the ground state of a many-body system with competing interactions in the regime of strong frustrations if the training is done on a small fraction of basis states. At the same time, even simple neural networks seem to have no problem in generalizing amplitudes from the training dataset onto the full Hilbert space, and have very good capacity to *express* both sign and amplitude distributions of the studied states. This understanding gives us a possibility to formulate a very concrete and simple test for future NQS architectures that will be used for studying ground state physics of many-body quantum systems:

A neural quantum state can be trained to approximate ground state of a large-scale many-body system only if it is capable of generalizing sign structure of a moderately-sized (exactly solvable) system ground state.

Finally, it is worth mentioning that, while the dip in generalization is not desirable in the context of variational energy optimization, it could be used as a tool to identify – in a completely unsupervised manner – the position of the phase transitions, similarly in spirit to approaches of [28, 29, 30, 31].

References

[1] Juan Carrasquilla and Roger G. Melko. Machine learning phases of matter. *Nature Physics*, 13(5):431, 2017.

[2] Giacomo Torlai, Guglielmo Mazzola, Juan Carrasquilla, Matthias Troyer, Roger Melko, and Giuseppe Carleo. Neural-network quantum state tomography. *Nature Physics*, 14(5):447, 2018.

[3] Felix Brockherde, Leslie Vogt, Li Li, Mark E. Tuckerman, Kieron Burke, and Klaus-Robert Müller. Bypassing the kohn-sham equations with machine learning. *Nature Communications*, 8(1):872, 2017.

[4] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *arXiv:1903.10563*, 2019.

[5] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017.

[6] Sandro Sorella, Michele Casula, and Dario Rocca. Weak binding between two aromatic rings: Feeling the van der waals attraction by quantum monte carlo methods. *The Journal of Chemical Physics*, 127(1):014105, 2007.

[7] Xiao Liang, Wen-Yuan Liu, Pei-Ze Lin, Guang-Can Guo, Yong-Sheng Zhang, and Lixin He. Solving frustrated quantum many-particle models with convolutional neural networks. *Physical Review B*, 98(10):104426, 2018.

[8] Di Luo and Bryan K. Clark. Backflow transformations via neural networks for quantum many-body wave functions. *Physical Review Letters*, 122(22):226401, 2019.

[9] Xun Gao and Lu-Ming Duan. Efficient representation of quantum many-body states with deep neural networks. *Nature Communications*, 8(1):662, 2017.

[10] Zi Cai and Jinguo Liu. Approximating quantum many-body wave functions using artificial neural networks. *Physical Review B*, 97(3):035116, 2018.

[11] Giuseppe Carleo, Yusuke Nomura, and Masatoshi Imada. Constructing exact representations of quantum many-body systems with deep neural networks. *Nature Communications*, 9(1):5322, 2018.

[12] Sandro Sorella. Generalized lanczos algorithm for variational quantum monte carlo. *Physical Review B*, 64(2):024512, 2001.

[13] Kenny Choo, Titus Neupert, and Giuseppe Carleo. Study of the two-dimensional frustrated j1-j2 model with neural network quantum states. *arXiv:1903.06713*, 2019.

[14] Giacomo Torlai, Juan Carrasquilla, Matthew T. Fishman, Roger G. Melko, and Matthew Fisher. Wavefunction positivization via automatic differentiation. *arXiv:1906.04654*, 2019.

[15] H.J. Schulz, T.A.L. Ziman, and Didier Poilblanc. Magnetic order and disorder in the frustrated quantum heisenberg antiferromagnet in two dimensions. *Journal de Physique I*, 6(5):675–703, 1996.

[16] J. Richter and J. Schulenburg. The spin-1/2 j 1–j 2 heisenberg antiferromagnet on the square lattice: Exact diagonalization for n= 40 spins. *The European Physical Journal B*, 73(1):117–124, 2010.

[17] Seiji Yunoki and Sandro Sorella. Two spin liquid phases in the spatially anisotropic triangular heisenberg model. *Physical Review B*, 74(1):014408, 2006.

[18] Andreas P. Schnyder, Oleg A. Starykh, and Leon Balents. Spatially anisotropic heisenberg kagome antiferromagnet. *Physical Review B*, 78(17):174420, 2008.

[19] P.H.Y. Li, R.F. Bishop, C.E. Campbell, Damian J.J. Farnell, O. Götze, and J. Richter. Spin-1/2 heisenberg antiferromagnet on an anisotropic kagome lattice. *Physical Review B*, 86(21):214403, 2012.

[20] Kenny Choo, Giuseppe Carleo, Nicolas Regnault, and Titus Neupert. Symmetries and many-body excited states with neural-network quantum states. *arXiv:1807.03325*, 2018.

[21] Michele Casula, Claudio Attaccalite, and Sandro Sorella. Correlated geminal wave function for molecules: An efficient resonating valence bond approach. *The Journal of Chemical Physics*, 121(15):7110–7126, 2004.

[22] A. Harju, B. Barbiellini, S. Siljamäki, Risto M. Nieminen, and G. Ortiz. Stochastic gradient approximation: An efficient method to optimize many-body wave functions. *Physical Review Letters*, 79(7):1173, 1997.

[23] Hiroki Saito and Masaya Kato. Machine learning technique to find quantum many-body ground states of bosons on a lattice. *Journal of the Physical Society of Japan*, 87(1):014001, 2017.

[24] `https://github.com/nikita-astronaut/nqs_frustrated_phase`.

[25] Haim Sompolinsky, Naftali Tishby, and H. Sebastian Seung. Learning from examples in large neural networks. *Physical Review Letters*, 65(13):1683, 1990.

[26] Hyunjune Sebastian Seung, Haim Sompolinsky, and Naftali Tishby. Statistical mechanics of learning from examples. *Physical Review A*, 45(8):6056, 1992.

[27] Carlo Baldassi, Alessandro Ingrosso, Carlo Lucibello, Luca Saglietti, and Riccardo Zecchina. Subdominant dense clusters allow for simple learning and high computational performance in neural networks with discrete synapses. *Physical Review Letters*, 115(12):128101, 2015.

[28] Evert P.L. Van Nieuwenburg, Ye-Hua Liu, and Sebastian D. Huber. Learning phase transitions by confusion. *Nature Physics*, 13(5):435, 2017.

[29] Peter Broecker, Fakher F. Assaad, and Simon Trebst. Quantum phase recognition via unsupervised machine learning. *arXiv:1707.00663*, 2017.

[30] Wenjian Hu, Rajiv R.P. Singh, and Richard T. Scalettar. Discovering phases, phase transitions, and crossovers through unsupervised machine learning: A critical examination. *Physical Review E*, 95(6):062122, 2017.

[31] Elmer V.H. Doggen, Frank Schindler, Konstantin S. Tikhonov, Alexander D. Mirlin, Titus Neupert, Dmitry G. Polyakov, and Igor V. Gornyi. Many-body localization and delocalization in large quantum chains. *Physical Review B*, 98(17):174202, 2018.