Regressive and generative neural networks for scalar field theory

Kai Zhou,^{1,2,*} Gergely Endrődi,² Long-Gang Pang,^{1,3,4} and Horst Stöcker^{1,2,5}

¹Frankfurt Institute for Advanced Studies, 60438 Frankfurt am Main, Germany

²Institut für Theoretische Physik, Goethe Universität, 60438 Frankfurt am Main, Germany

³Department of Physics, University of California, Berkeley, California 94720, USA

⁴Nuclear Science Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA

⁵GSI Helmholtzzentrum für Schwerionenforschung, 64291 Darmstadt, Germany

(Received 12 March 2019; published 9 July 2019)

We explore the perspectives of machine learning techniques in the context of quantum field theories. In particular, we discuss two-dimensional complex scalar field theory at nonzero temperature and chemical potential—a theory with a nontrivial phase diagram. A neural network is successfully trained to recognize the different phases of this system and to predict the values of various observables, based on the field configurations. We analyze a broad range of chemical potentials and find that the network is robust and able to recognize patterns far away from the point where it was trained. Aside from the regressive analysis, which belongs to supervised learning, an unsupervised generative network is proposed to produce new quantum field configurations that follow a specific distribution. An implicit local constraint fulfilled by the physical configurations was found to be automatically captured by our generative model. We elaborate on potential uses of such a generative approach for sampling outside the training region.

DOI: 10.1103/PhysRevD.100.011501

I. INTRODUCTION

Deep learning with a hierarchical structure of artificial neural networks is a branch of machine learning aiming at understanding and extracting high-level representations of big data [1]. It is particularly effective in tackling complex nonlinear systems with a high level of correlations that cannot be captured easily by conventional techniques. Traditionally employed for tasks like pattern recognition in images or speech, automated translation, or board game playing, applications of deep learning have been found recently in many areas of physics including nuclear [2–5], particle [6–10], and condensed matter [11–20] physics.

Significant progress has been made in utilizing machine learning methods for condensed matter systems like classical or quantum spin models. Specific tasks in these settings include the discrimination between certain phases and the identification of phase transitions [11–15], the compressed representation of quantum wave functions [16], or the acceleration of Monte Carlo algorithms [17–19]. Recently, deep neural networks have also been considered in particle physics, for the processing of experimental heavy-ion collision datasets [2] and in the context of algorithmic development for numerical lattice field theory simulations [21–23].

Pattern recognition, especially in classification and regression tasks, has been discussed previously in interacting many-body systems for condensed matter physics. In the present paper, we generalize the application of deep learning for the classification of phases in a lattice quantum field theoretical setting. We further demonstrate the capability of deep neural networks in learning physical observables, even with highly nonlinear dependence on the field configurations and with only limited training dataproviding an effective high-dimensional nonlinear regression method. In addition, we proceed by implementing, for the first time, a generative adversarial network (GAN) [24] for lattice field theory to generate field configurations following and generalizing the training set distribution. This is an unsupervised learning framework that uses unlabeled data to perform representation learning. Such a GAN-powered approach is not a full-fledged alternative to the Monte Carlo algorithm, which possesses desired properties like ergodicity, reversibility, and detailed balance. However, it can result in a one-pass direct sampling network where no Markov chain is needed. Our aim here is to provide a proof of principle that generative networks, if trained adequately, are capable of capturing and representing the distribution of configurations in a strongly correlated quantum field theory. On the practical side, generative networks would prove useful when combined with traditional approaches to accelerate simulation algorithms, e.g., by improving decorrelation for proposals in a Markov chain

zhou@fias.uni-frankfurt.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

process. Further potential use of such setups would be for reducing large ensembles of field configurations into a single (highly trained) network as an efficient representation for the quantum statistical field ensembles, thereby significantly reducing storage requirements.

II. OBSERVABLES IN SCALAR FIELD THEORY

We consider a complex scalar field ϕ with mass *m* and quartic coupling λ in (1 + 1)-dimensional Euclidean spacetime at nonzero temperature *T*. This system is studied in the grand canonical approach, introducing a chemical potential μ that controls how the charge density *n* fluctuates. For low temperatures, two different regimes of the parameter space can be distinguished: At low μ the density is suppressed, called the "silver blaze" behavior [25], whereas above a threshold $\mu > \mu_{\text{th}}$ the density increases considerably.¹ It is widely believed that the QCD phase diagram also exhibits such a behavior in the region at low temperatures and medium to high densities [25].

This interesting behavior is a nonperturbative phenomenon and cannot be observed directly, as the action becomes complex for $\mu \neq 0$, hindering standard simulations in terms of the field ϕ . However, using the worldline formalism, the partition function can be reexpressed using dual variables and the action rendered real and positive [26–28]. The dual variables are the integers $k_{\nu}(x)$ and $\ell_{\nu}(x)$ that are associated with the links starting at the point $x = (x_1, x_2)$ and lying in the direction $\nu = 1$ (space) or $\nu = 2$ (time). The total number of variables is therefore $\mathcal{N} \equiv 2 \times 2 \times N_1 \times N_2$, where N_{ν} denotes the number of lattice sites in the direction ν . While the ℓ -integers can take arbitrary values, the *k*integers satisfy a zero-divergence-type constraint and thus always form closed loops,

$$\sum_{\nu} [k_{\nu}(x) - k_{\nu}(x - a\hat{\nu})] = 0, \qquad (1)$$

where $\hat{\nu}$ is the unit vector in the ν direction, and *a* is the lattice spacing.

We consider a low-temperature ensemble $N_1 \times N_2 = 10 \times 200$ (the dimensionality of the configuration space is therefore $\mathcal{N} = 8000$) generated with mass m = 0.1, coupling $\lambda = 1.0$ and a range of chemical potentials $0.91 \le \mu \le 1.05$ around the threshold value $\mu_{\text{th}} \approx 0.94$ (all dimensionful quantities are understood in lattice units). For $\mu < \mu_{\text{th}}, \langle n \rangle$ is almost zero and $\langle \phi^2 \rangle$ is constant. In contrast, both observables rise approximately linearly beyond the threshold. This is demonstrated in Fig. 1.

III. SCALAR FIELD THEORY IN A NEURAL NETWORK

In the following section, we will apply deep neural networks for the complex scalar field configurations generated with the dualization approach using standard Monte Carlo methods.

A. Classification of phases

We first use the network to detect the transition between the low- and high-density phases of the system by performing a classification task. In particular, we train the neural network to identify the threshold chemical potential $\mu_{\rm th}$ without specific physical guidance. We train a convolutional neural network (CNN) to target at a binary classification: the configurations are either in the low-density "silver blaze" region ($\langle n \rangle \approx 0$) or in the condensed region ($\langle n \rangle \neq 0$). To perform a semisupervised training, we feed the lattice configurations at $\mu = 0.91 \ll \mu_{\rm th}$ and at $\mu = 1.05 \gg \mu_{\rm th}$ as input to the CNN. The training points are also highlighted in Fig. 1.

Once trained, we test the CNN by scanning through the configurations at different values of the chemical potential $0.91 < \mu < 1.05$. The output of the network for each configuration is identified as the probability *P* that the configuration in question corresponds to the condensed phase. In Fig. 2, we show the expectation value of *P* predicted by the network as a function of the chemical potential. We observe that along with increasing chemical potential, the first nonzero point of $\langle P \rangle$ reveals the transition at $\mu = \mu_{\text{th}}$. More interestingly, we find that the number density's nonzero value is perfectly indicated by the nonzero probability *P* from the trained network.

The particle number is given by the sum of all the k_2 variables. This simple pattern might be easily learned by



FIG. 1. The expectation values of the density and of the squared field, as functions of the chemical potential, connected by lines to guide the eye. The dashed vertical line marks the threshold chemical potential $\mu_{\rm th}$. The filled symbols indicate ensembles that are used in the training of our neural network (see details in the text).

¹In the following, we will refer to the pronounced change in the density at the threshold as a transition, keeping in mind that—in accordance with the Mermin-Wagner theorem—it is not connected to spontaneous symmetry breaking.



FIG. 2. The expectation value $\langle P \rangle$ of the condensation probability as a function of the chemical potential using full and restricted inputs. The error bar shows the statistical error in one ensemble. The dashed vertical line marks the threshold chemical potential, and the lines connecting the points serve to guide the eye.

the network. We thus perform the same binary classification task with a restricted training input, including only one of the remaining three variable sets: either k_1 , ℓ_1 , or ℓ_2 . The results for the three restricted inputs are visualized in Fig. 2. Clearly, the network succeeded in learning the same features also using the restricted inputs, as $\langle P \rangle$ starts to rise at around the same threshold chemical potential $\mu_{\text{th}} \approx$ 0.94 for all four cases. This inspires us to analyze the correlation between the number density and the similarly defined observables involving either k_1 , ℓ_1 , or ℓ_2 . To this end we consider the normalized correlation coefficient

$$R[A,B] \equiv \frac{\langle AB \rangle - \langle A \rangle \langle B \rangle}{\sqrt{\langle A^2 \rangle - \langle A \rangle^2} \sqrt{\langle B^2 \rangle - \langle B \rangle^2}}, \qquad (2)$$

which vanishes for decorrelated data and equals unity for complete correlation. As shown in Fig. 3, $\sum l_1$, $\sum l_2$, and also $|\phi|^2$ are all strongly correlated with *n*, while $\sum k_1$ is fully decorrelated. Still, the machine succeeds in classifying the configurations based only on k_1 , as indicated in Fig. 2. Here the alternating sum $\Sigma_x(-1)^x k_1(x)$ is also scrutinized, which equals $-\Sigma_x(-1)^x k_2(x)$ due to the closedloop constraint in Eq. (1), but it is not correlated with neither. With conventional techniques, neither of the physical observables (*n* and $|\phi|^2$) sensitive to the transition can be constructed using only the k_1 variables. The excellent performance of our CNN, shown in Fig. 2, indicates the existence of strong hidden features in the k_1 variables that correlate with the phase of the system. According to these results, the network can decode these hidden correlations in a highly effective manner.

B. Nonlinear regression of observables

Supervised learning is applied with a CNN to regress the thermodynamic observables of the system, including the



PHYS. REV. D 100, 011501 (2019)



FIG. 3. The normalized correlation coefficient [Eq. (2)] of the number density and various observables including $|\phi|^2$ and the sum of the k_1 , l_1 , l_2 variables or $\Sigma_x(-1)^x k_1(x)$ over all lattice sites.

particle number density and the squared field, based on the lattice configurations. As in Sec. III A, the training dataset consists of configurations at $\mu = 0.91$ and $\mu = 1.05$. The generalization ability of the machine is investigated by testing the network predictions on previously unseen configurations at different values of the chemical potential. The results for the density and for the squared field are plotted in Fig. 4, showing the true values of the observables (from physical formula) against the network predictions. Notice that the density, being an average of integers, only takes discrete values on a finite lattice.

As is visible in Fig. 4, the network performs well over a broad range of chemical potentials, predicting n and $|\phi|^2$ accurately, the maximal deviation being around 5% for the density and around 7% for the squared field. The training was performed using only two far-away segments of the range of the target observables, corresponding to configurations at $\mu = 0.91$ and $\mu = 1.05$ (the expectation values of the observables for these ensembles are also indicated in the plots). On the one hand, the high quality of the regression for the density may seem natural owing to the linear dependence of n on the individual variables. On the other hand, the squared field is a highly nonlinear function of the high-dimensional input $(\mathbb{R}^{200 \times 10 \times 4} \rightarrow \mathbb{R}^1)$ —even if $|\phi|^2$ and n are correlated (see Fig. 3), the squared field is not discrete, and thus its fluctuations are not completely captured by those of the density. This makes the excellent predictive ability of the network nontrivial and surprising. Put differently, using limited training data (covered small range of the target domain), our CNN network can correctly reproduce the whole target space mapping, which is curved and even dramatically changing (close to the transition point). This means that the network has effectively encoded the configuration into a plainer and a more abstract latent space (intermediate layers in the network). A linear interpolation in these layers can result in nonlinear regression in the final output layer.



FIG. 4. Comparison of the true values of the observables (vertical axis) to the network predictions (horizontal axis) for the particle density *n* (upper panel) and for the squared field $|\phi|^2$ (lower panel). Each filled circle represents one configuration, and the dotted crosses indicate the expectation values of the observables on the training ensembles.

C. Configuration production using the generative adversarial network

The generative adversarial network (GAN) [24] is a deep generative model that aims to learn the distribution of input variables from the training data. It belongs to the unsupervised learning category within deep learning approaches. Such generative modeling-assisted approaches have been tested in various scientific contexts, including medicine [29,30], particle physics [31–33], cosmology [34–36], and condensed matter physics [37,38]. Here we employ, for the first time, the generative modeling GAN application in strongly correlated quantum field theory. To ensure training stability, we consider the Wasserstein-GAN architecture [39] with gradient penalty (WGAN-gp) [40] in this study.

To verify the effectiveness of our GAN and, in particular, whether the generated configurations are indeed physical, we will first check the divergence-type constraint [Eq. (1)] for the complex scalar field. As shown in Fig. 5, the absolute divergence per site for the generated outputs is not



FIG. 5. The absolute divergence [using the lhs of Eq. (1)] per site for configurations from the GAN generator as a function of training epochs, with (blue) and without (red) rounding configuration entries to its nearest discrete value.

exactly zero but is decreasing and converges to zero as the number of training epochs grows. Equation (1) represents a highly implicit physical constraint inside the training dataset, which is not provided as supervision to the training of the GAN. Instead, the network automatically recognized this constraint for the configurations in a converging way. The generation time for a single configuration using the GAN (on an Nvidia TitanXp GPU) is 0.2 ms.

Next, we turn to the distribution of observables in the samples generated by the GAN and check to what extent it agrees with the training distribution. In the top panel of Fig. 6, we visualize the probability density distribution of the number density *n* from the GAN after training with one ensemble of configurations at $\mu = 1.05$. We observe that the GAN has captured the discrete distribution of n well. The ensemble average of the particle number density from GAN is estimated (using 1000 random samples) to be $\langle n \rangle_{\rm GAN} = 0.578$, also close to the Monte Carlo value $\langle n \rangle_{\rm MC} = 0.580$. As mentioned earlier, the particle number density is simply the average of the time component of the k variables in the configuration. In contrast, the squared field $|\phi|^2$ is highly nonlinear as a function of the input variables. Nevertheless, the multimodal distribution of $|\phi|^2$ is also well reproduced by the generative network; see the bottom panel of Fig. 6. The ensemble average of $|\phi|^2$ from GAN (for the same 1000 samples above), $\langle \phi^2
angle_{
m GAN} = 0.449$, is also close to the Monte Carlo result $\langle \phi^2 \rangle_{\rm MC} = 0.447$. Figures 5 and 6 demonstrate that the generative adversarial network can be trained to capture the statistical distribution of the field configurations even on the level of physical observables.

The above GAN structure is designed to reproduce certain distributions in the training dataset. Next, we attempt to use the network to generalize the distribution



FIG. 6. The probability density distribution of the number density *n* (top panel) and of the squared field $|\phi|^2$ (bottom panel) from the GAN (green) along with the training data distribution obtained from the Monte Carlo simulation (blue) for fixed chemical potential $\mu = 1.05$ with 1000 samples.



FIG. 7. The mean particle number density on the configurations generated by the cGAN with (blue) and without (red) rounding configuration entries to the nearest discrete value, against the specified condition values for n.

that it was trained on. The discriminator is provided with relevant labels (in this case, the value of the number density *n*) for the training dataset to condition the network (cGAN) [41]. Specifically, the training sample at $\mu = 1.05$ contains the cases n = 0.4, 0.5, 0.6, and 0.7. After the training, we test the generalization ability of the network by specifying desired number densities outside the above set of values. Figure 7 shows the performance of the cGAN for this generalization task. We stress that for the training only $0.4 \le n \le 0.7$ values were provided, but the agreement between the desired n (condition) and the measured n on the generated configurations is spectacular over a broader range of density values. This generalization task might be viewed as converting the grand-canonical ensemble of configurations (at fixed μ) to a series of canonical ensembles (at various values of n).

IV. CONCLUSIONS

In this paper, we proposed a set of novel techniques for the investigation of a lattice-regularized quantum field theory by deep neural networks, including discovering hidden correlations, learning observables, and producing field configurations. Specifically, our analysis was carried out for the dualized representation of complex scalar field theory in 1 + 1 dimensions.

We first showed that a convolutional neural network can be used in a semisupervised manner to detect the phase transition in this strongly correlated quantum field theory based on the microscopic configurations. We found that the network is capable of recognizing correlations in the system between various observables and phase classifications without the specific knowledge guidance. Interestingly, the network discovered a correlation beyond the conventional analysis, which enabled it to use a restricted subset of the input variables (in particular, the k_1 variables) alone to decode information about the phase transition.

We continued by designing a regressive neural network to learn physical observables (*n* and $|\phi|^2$) with limited training samples. The network achieved remarkable agreement with the physical observables and also revealed a great generalization ability when tested at chemical potentials beyond the training set. This approach provides an effective high-dimensional nonlinear regression method even with limited data points (compared to the huge Hilbert space, i.e., the number of possible configurations), where traditional interpolation or regression would require higher statistics that grow exponentially with input dimensionality.

Finally, we proposed to generate new configurations following a specific distribution by adapting the modern deep generative modeling technique GAN. We found that the generator in the GAN has the ability of automatically recognizing the implicit but crucial physical constraint on the configurations in an unsupervised manner, and it can represent the distribution of prominent observables with direct sampling. The generalization of configuration production to different parameter domains, e.g., towards a critical region, where conventional techniques slow down considerably, is a fascinating feature that deserves further investigation.

ACKNOWLEDGMENTS

This work is supported by the funding of AI research at Frankfurt Institute for Advanced Studies (FIAS) from

- Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, Nature (London) **521**, 436 (2015).
- [2] L.-G. Pang, K. Zhou, N. Su, H. Petersen, H. Stöcker, and X.-N. Wang, An equation-of-state-meter of quantum chromodynamics transition from deep learning, Nat. Commun. 9, 210 (2018).
- [3] R. Utama and J. Piekarewicz, Validating neural-network refinements of nuclear mass models, Phys. Rev. C 97, 014306 (2018).
- [4] T. Bayram, S. Akkoyun, and Ş. Şentürk, Adjustment of nonlinear interaction parameters for relativistic mean field approach by using artificial neural networks, Phys. At. Nucl. 81, 288 (2018).
- [5] L. Neufcourt, Y. Cao, W. Nazarewicz, and F. Viens, Bayesian approach to model-based extrapolation of nuclear observables, Phys. Rev. C 98, 034318 (2018).
- [6] P. Baldi, P. Sadowski, and D. Whiteson, Searching for exotic particles in high-energy physics with deep learning, Nat. Commun. 5, 4308 (2014).
- [7] P. Baldi, P. Sadowski, and D. Whiteson, Enhanced Higgs Boson to $\tau^+\tau^-$ Search with Deep Learning, Phys. Rev. Lett. **114**, 111801 (2015).
- [8] J. Barnard, E. N. Dawe, M. J. Dolan, and N. Rajcic, Parton shower uncertainties in jet substructure analyses with deep neural networks, Phys. Rev. D 95, 014018 (2017).
- [9] I. Moult, L. Necib, and J. Thaler, New angles on energy correlation functions, J. High Energy Phys. 12 (2016) 153.
- [10] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel, A. Aurisano, K. Terao, and T. Wongjirad, Machine learning at the energy and intensity frontiers of particle physics, Nature (London) 560, 41 (2018).
- [11] L. Wang, Discovering phase transitions with unsupervised learning, Phys. Rev. B 94, 195105 (2016).
- [12] E. P. L. van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Learning phase transitions by confusion, Nat. Phys. 13, 435 (2017).
- [13] K. Ch'ng, J. Carrasquilla, R. G. Melko, and E. Khatami, Machine Learning Phases of Strongly Correlated Fermions, Phys. Rev. X 7, 031038 (2017).
- [14] P. Broecker, J. Carrasquilla, R. G. Melko, and S. Trebst, Machine learning quantum phases of matter beyond the fermion sign problem, Sci. Rep. 7, 8823 (2017).

SAMSON AG Frankfurt (K. Z.), DFG Emmy Noether Programme No. EN 1064/2-1 (G. E.), NSF JETSCAPE No. ACI-1550228 (L.-G. P.), and endowment through the Judah M. Eisenberg Laureatus Chair at Goethe University and the Walter Greiner Gesellschaft at Frankfurt (H. St.). K. Z. gratefully acknowledges support from the NVIDIA Corporation with the donation of Nvidia TitanXp GPU cards for the research. G. E. acknowledges inspiring discussions with Gábor Varga.

- [15] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, Nat. Phys. 13, 431 (2017).
- [16] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, Science 355, 602 (2017).
- [17] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, Self-learning Monte Carlo method, Phys. Rev. B 95, 041101 (2017).
- [18] G. Torlai and R. G. Melko, Learning thermodynamics with Boltzmann machines, Phys. Rev. B 94, 165134 (2016).
- [19] L. Huang and L. Wang, Accelerated Monte Carlo simulations with restricted Boltzmann machines, Phys. Rev. B 95, 035105 (2017).
- [20] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, Unsupervised Generative Modeling Using Matrix Product States, Phys. Rev. X 8, 031012 (2018).
- [21] Y. Mori, K. Kashiwa, and A. Ohnishi, Application of a neural network to the sign problem via the path optimization method, Prog. Theor. Exp. Phys. 2018, 023B04 (2018).
- [22] P. E. Shanahan, D. Trewartha, and W. Detmold, Machine learning action parameters in lattice quantum chromodynamics, Phys. Rev. D 97, 094506 (2018).
- [23] A. Tanaka and A. Tomiya, Towards reduction of autocorrelation in HMC by machine learning, arXiv:1712.03893.
- [24] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative adversarial nets, in *Advances in Neural Information Processing Systems 27*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Curran Associates, Inc., 2014), pp. 2672–2680, http://papers.nips.cc/paper/5423-generative-adversarial-nets .pdf.
- [25] T. D. Cohen, Functional Integrals for QCD at Nonzero Chemical Potential and Zero Density, Phys. Rev. Lett. 91, 222001 (2003).
- [26] C. Gattringer and T. Kloiber, Lattice study of the Silver Blaze phenomenon for a charged scalar φ^4 field, Nucl. Phys. **B869**, 56 (2013).
- [27] O. Orasch and C. Gattringer, Canonical simulations with worldlines: an exploratory study in φ_2^4 lattice field theory, Int. J. Mod. Phys. A **33**, 1850010 (2018).
- [28] C. Gattringer, M. Giuliani, and O. Orasch (private communication).

- [29] D. Nie, R. Trullo, C. Petitjean, S. Ruan, and D. Shen, Medical image synthesis with context-aware generative adversarial networks, arXiv:1612.05362.
- [30] D. Mahapatra, B. Bozorgtabar, S. Hewavitharanage, and R. Garnavi, Image super resolution using generative adversarial networks and local saliency maps for retinal image analysis, in *MICCAI* (Springer, Cham, 2017), http://dx.doi .org/10.1007/978-3-319-66179-7_44.
- [31] L. de Oliveira, M. Paganini, and B. Nachman, Learning particle physics by example: Location-aware generative adversarial networks for physics synthesis, Comput. Software Big Sci. 1, 4 (2017).
- [32] M. Paganini, L. de Oliveira, and B. Nachman, CaloGAN: Simulating 3D high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks, Phys. Rev. D 97, 014021 (2018).
- [33] M. Paganini, L. de Oliveira, and B. Nachman, Accelerating Science with Generative Adversarial Networks: An Application to 3D Particle Showers in Multilayer Calorimeters, Phys. Rev. Lett. **120**, 042003 (2018).
- [34] S. Ravanbakhsh, F. Lanusse, R. Mandelbaum, J. Schneider, and B. Poczos, Enabling Dark Energy Science with Deep Generative Models of Galaxy Images, arXiv:1609.05796.

- [35] K. Schawinski, C. Zhang, H. Zhang, L. Fowler, and G. K. Santhanam, Generative adversarial networks recover features in astrophysical images of galaxies beyond the deconvolution limit, Mon. Not. R. Astron. Soc. 467, L110 (2017).
- [36] M. Mustafa, D. Bard, W. Bhimji, Z. Lukić, R. Al-Rfou, and J. Kratochvil, CosmoGAN: creating high-fidelity weak lensing convergence maps using Generative Adversarial Networks, arXiv:1706.02390.
- [37] L. Mosser, O. Dubrule, and M. J. Blunt, Reconstruction of three-dimensional porous media using generative adversarial neural networks, Phys. Rev. E 96, 043309 (2017).
- [38] K. Mills and I. Tamblyn, Deep neural networks for direct, featureless learning through observation: The case of two-dimensional spin models, Phys. Rev. E 97, 032119 (2018).
- [39] M. Arjovsky, S. Chintala, and L. Bottou, Wasserstein GAN, arXiv:1701.07875.
- [40] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, Improved Training of Wasserstein GANs, arXiv:1704.00028.
- [41] M. Mirza and S. Osindero, Conditional Generative Adversarial Nets, arXiv:1411.1784.