Constructing gauge-invariant neural networks for scientific applications

Emmanouil Theodosis [1] **Demba Ba** [1] **Nima Dehmamy** [2]

Abstract

Our current models for fundamental forces in nature are "gauge theories". These models are suitable for systems where interactions are local and where the local choice of coordinates does not affect physical quantities. While recent works have introduced gauge equivariant neural networks, these models focus on tangent bundles or quotient space and are not applicable to most gauge theories appearing in physics. We propose an architecture for learning general gauge invariant quantities. Our framework fills a gap in the existing literature, providing a general recipe for gauge invariance without restrictions on the spaces of the measurement vectors. We evaluate our method on a classical physical system, the XY model, that is invariant to the choice of local gauges. We make our code publicly available at https://github.com/manosth/gauge-net/.

1. Introduction

Gauge theories (Baez & Muniain, 1994; Nakahara, 2018) are one of the most consequential paradigms in modern physics, providing a framework for understanding the fundamental interactions in the universe. They describe how fields, like the electromagnetic and the gravitational fields, interact with matter, forming the foundation of the Standard Model of particle physics and general relativity. The key point in systems described by gauge theories is that different symmetry transformations can be applied at points in space, without impacting the quantities we measure.

In recent years, the principles of gauge theories have inspired advances in neural network architectures, particularly through the introduction of gauge (Cohen et al., 2019b; de Haan et al., 2021; Luo et al., 2022) and group equivariant neural networks (Cohen & Welling, 2016; Dieleman et al., 2016; Cohen & Welling, 2017; Weiler & Cesa, 2020; Sosnovik et al., 2020). Models such as Gauge Equivariant Mesh CNNs (de Haan et al., 2021) and others have been central to applying group equivariance to structured data, such as images and meshes, demonstrating significant improvements in tasks like image classification and 3D shape analysis.

However, existing approaches often restrict themselves to simpler geometric settings, such as homogeneous spaces or quotient spaces (Kondor & Trivedi, 2018; Cohen et al., 2019a), which are only special cases of the more complex and varied structures found in theoretical physics. These simplifications compromise the ability of models to handle the more general and intricate gauge symmetries seen in

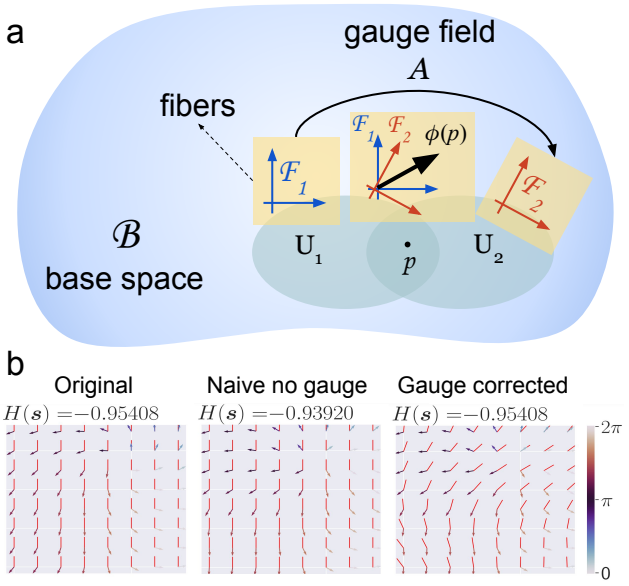

Figure 1: **Fiber bundles and gauge invariance.** (a) Overlapping neighboring open sets $U_1, U_2 \subset \mathcal{B}$ can have different bases in the fibers $\mathcal{F}_1$, $\mathcal{F}_2$ to represent *the same* feature $\phi(p)$ at point $p \in U_1 \cap U_2$. A gauge field A accounts for such local (gauge) basis transformations in the fibers, ensuring measured quantities are invariant under them. Examples of gauge invariance: (b) The energy of locally interacting magnets (XY model) changes little by small perturbation to spins and is completely invariant under local basis changes.

[1]School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA [2]IBM-Research, Cambridge, MA, USA. Correspondence to: Emmanouil Theodosis <etheodosis@g.harvard.edu>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024.

physical theories, where the base space can be any manifold and the fibers are not necessarily linked by the symmetries of the base space itself. Part of the reason for these limitations may be that existing work often begins with the framework of group equivariance (Cohen & Welling, 2016; Kondor & Trivedi, 2018) and steerable networks (Cohen & Welling, 2017; Weiler et al., 2018), which generally assume a direct correspondence between the group acting on the base space and the actions within the fibers.

In this paper, we aim to bridge this gap by proposing an architecture that does not rely on group equivariance on the base space. Instead, we focus on local invariance under the fiber symmetry group. By focusing on infinitesimal transformations, we utilize the "connection" or "gauge field"–a fundamental concept in gauge theories that describes how to compare local fiber elements across different points in the base space. This approach opens new avenues for applying neural networks to problems grounded in physics and other fields where local symmetries are essential.

2. Related work

Here we briefly review recent work on symmetry learning, equivariant networks, and gauge equivariance.

Equivariant frameworks. Certain frameworks have enforced equivariance with arbitrary groups, tackled groups with non-standard product structures, and have reconsidered the enforcement of strict equivariance. EMLP (Finzi et al., 2021) provides a framework to create equivarint multilayer perceptrons for arbitrary groups by solving a constrained optimization problem. Given a matrix group an MLP can be constructed that is equivariant to that group. However, it is unclear how this framework can be extended to work with local gauges, and can't handle large data like fields as we consider in this work[1]. Bogatskiy et al. (2020) introduced an equivariant neural network for the Lorentz group; extending that line of though from the Minkowski metric to arbitrary metrics (Ruhe et al., 2023b), Ruhe et al. (2023a) introduced Clifford equivariant networks and their steerable versions (Zhdanov et al., 2024). However, these works require explicit constructions for each metric, in the case of the Lorentz networks, and ignore the action of groups on both the base space and the fibers, in the case of the Clifford networks. Wang et al. (2022) consider a setting where the equivariance is approximate and satisfies $\|f(\rho_X(g)(\boldsymbol{x}) - \rho_Y f(\boldsymbol{x})\|_2 \leq \epsilon$ to account for imperfections in measuring devices or experimental tools. The framework can better model real world data, however requires explicit knowledge of the desired symmetry.

[1]These shortcomings are discussed in the official reprepository of the paper.

Symmetry learning. The efficiency and generalization properties of equivariant methods is limited in certain application domains where we lack a model, the model is not exact, or the symmetries are not immediately apparent. Still, efficiency and generalization make learning the acting symmetry desirable in these settings, and in many cases the abundance of data makes it possible. Dehmamy et al. (2021) introduce an architecture to learn continuous groups through their Lie algebras and follow up works (Yang et al., 2023b;a) make it more computationally efficient through the introduction of GANs and consider latent symmetries, but do not consider vector-valued fields. A different thrust of works relies on sequential data to learn transformations acting on the data; these include approaches focusing on the bispectrum (Sanborn et al., 2023), topographic organization (Keller & Welling, 2021), or disentanglement (Quessard et al., 2020). However, in all of these instances sequential data are required, with the added constraint that they need to be orbit-separated: each sequence must be a product of the same group action, leading to a, at least partially, supervised flavor.

Gauge equivariance. Early works in the field tackled equivariant learning on spaces with non-trivial curvature, such as the sphere (Cohen et al., 2018; Esteves et al., 2018; 2020). However, a global coordinate system, or frame, always exists in these works, unlike physical theories. Importantly, Cohen et al. (2019a) introduced a mathematical framework for equivariant learning on fiber bundles, though their work is limited to homogeneous spaces and their quotient groups. More recent works account for different local frames for their respective applications. Cohen et al. (2019b) consider gauge-equivariant learning on icosahedral approximations of spheres by engineering the transition maps between gauges. A growing body of work has considered gauge equivariance on mesh grids (de Haan et al., 2021; He et al., 2021; Basu et al., 2022; Park et al., 2023). In these works, local gauges are assigned on discrete grid points and the proposed architectures align the representations of different grid points, with possibly the inclusion of attention mechanisms (He et al., 2021; Basu et al., 2022; Park et al., 2023). In their entirety, these works consider exclusively tangent bundles, limiting the expressivity of the fibers, or restrict the fibers to be quotient spaces, which doesn't allow for complex fiber bundles found in physics, such as $\mathrm{SU}(2) \times \mathrm{U}(1)$ fibers for the electroweak theory, or $\mathrm{SU}(3)$ fibers for quantum chromodynamics. Finally, recent work has studied gauge-invariant networks in simulated physical applications involving quantum lattice models (Luo et al., 2023; Chen et al., 2022).

3. Background

In both physics and applied mathematics, the concept of a "field" is fundamental to understanding how properties vary across space. In simple terms, a field is a state variable that depends on spatial coordinates. For instance, in an image, each pixel $p \in \mathbb{Z}^2$ has a color $f(p) \in \mathcal{F} = [0,1]^3$ with values in the feature space $\mathcal{F}$. This construction where to each neighborhood of a base space (e.g. $\mathbb{Z}^2$) we attach a feature space (called "fibers") is an example of a trivial "fiber bundle" (Baez & Muniain, 1994) which we will formally define later. This notion can be extended to more complex data structures, like temperature or wind direction across a region, electromagnetic fields in physics, or any other variable distributed over a spatial domain.

Intuition for gauge invariance. In the example of images, the basis for the fibers needs to be the same for all of the fibers to ensure colors are defined consistently. However, in many problems in physical sciences, the quantities of interest may be invariant under local basis transformations, called "gauge transformation". Gauge theory is a paradigm for modelling such systems. A simple, analogy can be drawn from the field of image processing. For instance, consider the task of digit recognition in MNIST. When these images undergo small local distortions–small diffeomorphisms–the identity of the digits remains unchanged.

Example: XY model. In physics, a clear illustration of gauge theory principles can be seen in the XY model (Nagaosa, 2013), a system of 2D magnets (spins) interacting on a lattice (Figure 1c)[2]. Let $\boldsymbol{s}_j = \exp[i\theta_j] \in \mathbb{C}$ (with $\theta_j \in [0, 2\pi]$) denote the continuous spin state at node $j \in \mathcal{B} = \{1, \dots, N\}$ of the lattice. Here, the lattice $\mathcal{B}$ is the "base space" and the $\mathbb{C}$ at each j where the $\boldsymbol{s}_j$ take values are the "fibers". Given the overall spin configuration $\boldsymbol{s} = \{\boldsymbol{s}_i\}_{i=1}^N$, the energy is defined as

$$H(\boldsymbol{s}) = -\sum_{ij} J_{ij} \boldsymbol{s}_i^T \boldsymbol{s}_j = \frac{1}{2} \sum_{ij} J_{ij} \|\boldsymbol{s}_i - \boldsymbol{s}_j\|^2 - N \quad (1)$$

where J is the coupling matrix (the adjacency matrix of the lattice in this case). Clearly, the XY system's energy is invariant under *global* rotation in the fiber spaces $\boldsymbol{s}_j \to e^{i\alpha} \boldsymbol{s}_j$ (the same transformation on all the fibers). But more importantly, it also is invariant under *local* rotations, where in different neighborhoods of the lattice, the neighboring fibers are rotated together in such a way that $\boldsymbol{s}_i^T \boldsymbol{s}_j$ is not changed. This local symmetry is the "gauge invariance" of the XY model. But to allow for the fiber basis at each point to transform independently, we need to introduce a "gauge field" A which cancels the difference in gauge (basis) at neighboring points. To do so, we define the energy as

$$H(\boldsymbol{s}, A) = \frac{1}{2} \sum_{ij} J_{ij} \|(1 + A_{ij}) \boldsymbol{s}_i - \boldsymbol{s}_j\|^2 - N \quad (2)$$

and define the local transformations of $\boldsymbol{s}$ and A as (see Appendix A)

$$\begin{aligned} \boldsymbol{s}_j &\to e^{i\alpha_j} \boldsymbol{s}_j, \\ A_{ij} &\to e^{i(\alpha_j - \alpha_i)}(1 + A_{ij}) - 1. \end{aligned} \quad (3)$$

This way, the energy of the system remains invariant under local rotations of the angle defining the spins. This observation is essential for designing neural network architectures that can recognize and respect local symmetries and transformations.

Continuous XY model. Now consider the case where instead of a discrete lattice the spins are defined as a continuous complex field $S : \mathbb{R}^2 \to \mathbb{C}$, where the base space is $\mathcal{B} \sim \mathbb{R}^2$ and the fibers are the same as before. The position index i has become x and its neighbors j become $x + \delta x$, with infinitesimal δx. Expanding to first order in δx, $\boldsymbol{s}_i - \boldsymbol{s}_j = -\delta x \cdot \nabla S(x)$. Similarly, A_{ij} becomes a function $\delta x A(x)$, which determines the fiber basis transformation from x to $x + \delta x$. From this, the energy becomes $H = \int dx^2 \|(\nabla + A) S(x)\|^2$. Under a gauge transformation

$$S(x) \to e^{i\alpha(x)} S(x), \qquad A \to A - i\nabla\alpha \quad (4)$$

which keeps the overall energy unchanged. This concept illustrates how gauge fields absorb local changes in the fiber space to ensure that global properties like energy remain consistent.

3.1. Gauge symmetry

Differential geometry Informally, a manifold is a space that locally resembles $\mathbb{R}^n$. We define an *n-dimensional manifold* $\mathcal{B}$ as a topological space equipped with charts (local mappings) $f_\alpha : U_\alpha \to \mathbb{R}^n$, where $U_\alpha \subset \mathcal{B}$ are open sets that cover $\mathcal{B}$. These charts are such that the transition function $f_\alpha \circ f_\beta^{-1}$ is smooth where it is defined.

Definition 3.1 (Fiber bundle)**.** A fiber bundle is a triple $\mathcal{E} = (\mathcal{B}, \mathcal{F}, \pi)$ consisting of a total space $\mathcal{E}$, a base space $\mathcal{B}$, and a projection map $\pi : \mathcal{E} \to \mathcal{B}$. For each point $p \in \mathcal{B}$, the set

$$\mathcal{F}_p = \{q \in \mathcal{E} : \pi(q) = p\} \quad (5)$$

is called the fiber over p. A "*section*" of the fiber bundle is a map $\phi : \mathcal{B} \to \mathcal{E}$ such that for any point $p \in \mathcal{B}$, it holds that $\pi(\phi(p)) = p$.

In physics, sections are called *fields*. A **trivial fiber bundle** is a particular case where the total space $\mathcal{E}$ is simply the

[2] In Appendix B.1 we provide further context as to how Figure 1b is generated.

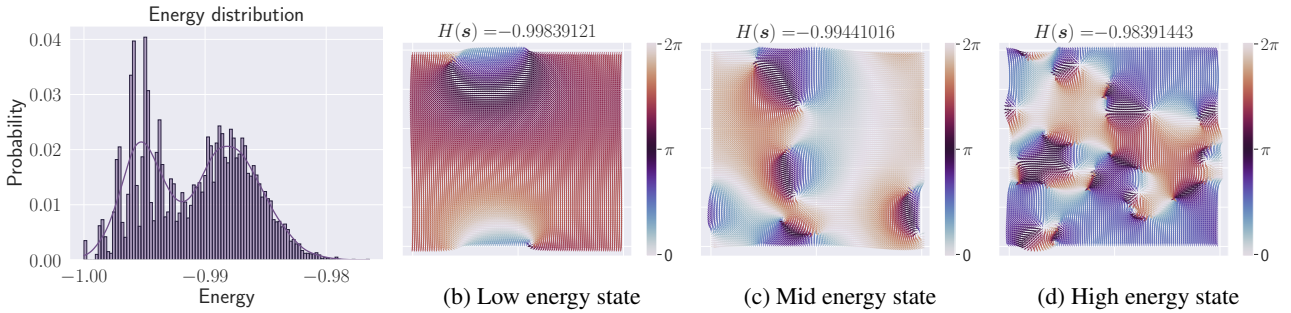

(a) Bimodally distributed

(b) Low energy state

(c) Mid energy state

(d) High energy state

Figure 2: Energy distribution (a) and samples of final configurations (b-d) for the XY model. Energy is correlated with vorticity: low, mid, and high energy states have low, mid, and high levels of vorticity, respectively.

product of the base space $\mathcal{B}$ and the fiber $\mathcal{F}$ (isomorphic to the product $\mathcal{B} \times \mathcal{F}$). For instance, an image is a section of a trivial fiber bundle. In the problems we are interested in, the fiber space has a symmetry group which acts on it. A **principal bundle** is a special type of fiber bundle particularly significant in gauge theories. In a principal bundle, each fiber is a homogeneous space—an orbit of a group G acting smoothly, freely, and transitively. Gauge theories primarily deal with principal bundles, as they are crucial for defining connections (gauge fields). These connections describe how to "transport" elements within the fiber from one point in the base space to another.

Connections in Fiber Bundles A connection [3] on a fiber bundle is a mathematical tool that allows us to define how to "transport" elements along paths within the base space $\mathcal{B}$, while respecting the fiber structure. As illustrated in Figure 1, different local bases in the fibers leads to conflicting representations of fields. A connection enables us to compare elements of the fiber over different points in $\mathcal{B}$, effectively defining a "covariant derivative". The key idea is that the standard derivative ∂_μ does not account for the difference in the basis of fiber $\mathcal{F}_p$ and $\mathcal{F}_{p+\delta x}$. A gauge field A modifies the standard derivative ∂_μ to a covariant derivative D_μ to maintain consistency between the local frames.

$$D_\mu = \partial_\mu + A_\mu. \tag{6}$$

Gauge equivariance Let G denote the the symmetry group of all $\mathcal{F}_p$, and let $g : \mathcal{B} \to G$ be set of group transformations which can act on a section (field) $\phi : \mathcal{B} \to \mathcal{E}$ to yield $\phi'(p) = g(p)\phi(p)$. To achieve gauge equivariance, the covariant derivative needs to be gauge (or group) equivariant, i.e., $D'(g \cdot f) = g \cdot Df$. This implies a constraint on the vector potentials of the form

$$A' = gAg^{-1} - (\nabla g)g^{-1}. \tag{7}$$

This concept is fundamental in physics and differential geometry, as it enables us to compare elements of the fiber over different points in $\mathcal{B}$, defining a "covariant derivative".

In a more formal setting, A_μ can be seen as a one-form on the base manifold that takes values in the Lie algebra of the gauge group G. This Lie algebra-valued one-form is defining how the gauge transformations adjust the fields and their derivatives, ensuring all physical predictions remain invariant under local transformations of the gauge group.

4. Experiments

As we discussed in Section 3, the XY model can be expressed as a gauge field and its energy is invariant to local (and global) transformations. We propose an architecture that constructs features that are, by design, gauge-invariant. Using the gauge-invariant features, we make local energy predictions, which are then aggregated to solve the problem of energy estimation in the XY model.

Embedded XY Some processes, such as embedding, can induce a natural gauge on the fiber bundle. For example, assume that we have the XY model and we have identified the fibers with the tangent spaces, meaning $\mathcal{F}_p \simeq T\mathcal{B}_p \simeq \mathbb{R}^2$. Now, suppose we map the XY system to a sphere in $\mathbb{R}^3$. This naturally results in a nontrivial gauge. To see this, choose the basis for $T\mathcal{B}_p$ to be the tangents in spherical coordinates (ϕ, θ), where $\phi \in [0, 2\pi)$ is the azimuthal angle and $\theta \in (0, \pi)$ is the polar angle. The tangent vectors $\partial_{(\phi,\theta)(p)} \equiv (\partial/\partial\phi(p), \partial/\partial\theta(p))$ form the local basis for $T\mathcal{B}_p$. To see how $\partial_{(\phi,\theta)}$ varies from p to $p + \delta p$, we can use the embedding coordinates $\vec{p} = (x, y, z)$ and express $\partial_{(\phi,\theta)}$

[3] The term "connection" is used slightly differently among mathematicians, who use it for the covariant derivative D, compared to physicists, using it for the gauge field A. To avoid confusion we will mostly use the terms covariant derivative for D and gauge field for A.

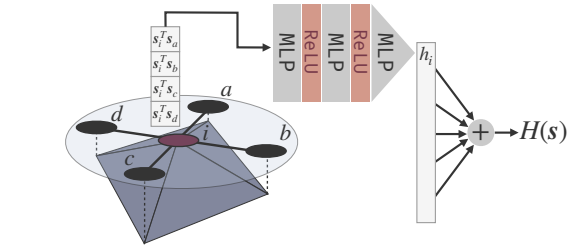

(a)

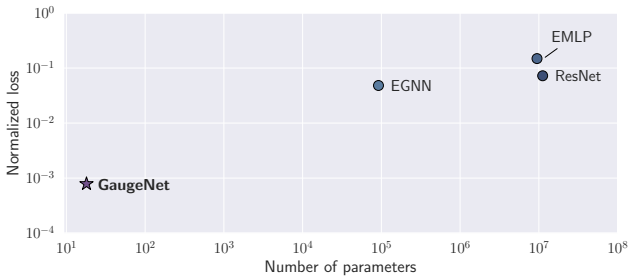

(b)

Figure 3: (a) Gauge-invariant architecture. Each neighborhood computes a gauge-corrected local estimate for the energy, which are then aggregated for the total energy of the configuration $H(\boldsymbol{s})$. (b) Performance of neural network architectures on energy regression for the XY model. Our method, GaugeNet, is shown in bold.

in terms of $\partial_{\vec{p}} = (\partial_x, \partial_y, \partial_z)$.We have

$$\partial_{(\phi,\theta)} = \frac{\partial(x,y,z)}{\partial(\phi,\theta)} \partial_{(x,y,z)} = J \partial_{(x,y,z)} \tag{8}$$

where J is the Jacobian of the transformation.

Dataset We generated $10,000$ samples, each representing a field configuration $\boldsymbol{s} \in \mathbb{R}^{100\times100}$. Original spin configurations were i.i.d. samples from a uniform distribution and iterated over $10,000$ gradient steps to reach each final configuration. The dataset consists of data pairs $\{(\boldsymbol{s}_i, y_i)\}_{i=1}^{10,000}$ with $y_i = H(\boldsymbol{s}_i) \in \mathbb{R}$ being the energy of each final configuration, which we consider as the target variable for a regression task. The XY model has multiple energy basins that the system can converge to. In Figure 2 we plot the overall distribution of the energies and samples of low-, mid-, and high- energy configurations. For details about the data-generating procedure see Appendix B.2 and for more samples from the dataset see Appendix B.3.

Architecture To construct an architecture that is inherently invariant to the choice of gauge, we opted to construct gauge invariant features. Inner products are preserved under gauge transformations, assuming i, j are expressed in the same local gauge. The challenge, then, is to ensure that quantities are expressed in the same gauge when computing the inner products. To that end, in light of Equation (2), we introduce a small, one hidden layer MLP that estimates the misalignment in gauges from the node positions.

After correcting for the difference in gauge, for every local neighborhood at a node i, we compute the inner products between the central node and its neighbors. To model the local contribution of each neighborhood of the XY model configuration to the overall energy, we adopt a simple two hidden layer MLP to estimate the local energy. Then the local energies are summed, yielding the overall estimate for the configuration $H(\boldsymbol{s})$. An illustration of the architecture, which we call **GaugeNet**, is given in Figure 3a.

Baselines We compare against high-performing general architectures and equivariant neural networks. As a general architecture, we use ResNet18 (He et al., 2016) as a baseline. For the equivariant baselines, we compare against EMLP (Finzi et al., 2021) equipped with $\mathrm{SO}(2)$-invariance.

Results In Figure 3b, we show the performance of the architectures relative to their parameters. We see that GaugeNet provides significant performance gains over baselines, reducing the loss by almost a factor of 100. At the same time, our model is parameter efficient, requiring almost $10,000\times$ less parameters than the next smallest model. EMLP struggles to provide accurate predictions as it models global, instead of local invariances and suffers from large parameter counts due to its full connectivity. While great general purpose architectures, ResNets do not have any local invariance properties to allow them to perform favorably. Finally, EGNN is more parameter efficient than the other baselines due to its graph structure and performs all computations locally. However, despite its equivariance properties, it does not exploit the problem structure adequately to accurately predict the configuration's energy, and still requires orders of magnitude more parameters than GaugeNet.

5. Conclusion and discussions

In this work we presented GaugeNet to solve the energy estimation problem in the XY model, utilizing the problem's structure to construct gauge-invariant features. As inner products are preserved under gauge transformations, at each neighborhood we compute the inner products between the central node and its neighbors, which we use as features. As the energy is location-agnostic, our model makes local energy estimations at every neighborhood, which are aggregated for the final prediction. GaugeNet provides significant performance gains over baselines which include both equivariant and general purpose architectures, for only a small fraction of the parameters.

References

Baez, J. and Muniain, J. *Gauge fields, knots and gravity*. World Scientific Publishing Company, 1994.

Basu, S., Gallego-Posada, J., Viganó, F., Rowbottom, J., and Cohen, T. Equivariant mesh attention networks. *Transactions on Machine Learning Research*, 2022.

Bogatskiy, A., Anderson, B., Offermann, J., Roussi, M., Miller, D., and Kondor, R. Lorentz group equivariant neural network for particle physics. In *International Conference on Machine Learning*, 2020.

Chen, Z., Luo, D., Hu, K., and Clark, B. Simulating $2+1d$ lattice quantum electrodynamics at finite density with neural flow wavefunctions. In *arXiv*, 2022.

Cohen, T. and Welling, M. Group equivariant convolutional networks. In *International Conference on Machine Learning*, 2016.

Cohen, T. and Welling, M. Steerable cnns. In *International Conference on Learning Representations*, 2017.

Cohen, T., Geiger, M., Köhler, J., and Welling, M. Spherical cnns. In *International Conference on Learning Representations*, 2018.

Cohen, T., Geiger, M., and Weiler, M. A general theory of equivariant cnns on homogeneous spaces. In *Advances in Neural Information Processing Systems*, 2019a.

Cohen, T., Weiler, M., Kicanaoglu, B., and Welling, M. Gauge equivariant convolutional networks and the icosahedral cnn. In *International Conference on Machine Learning*, 2019b.

de Haan, P., Weiler, M., Cohen, T., and Welling, M. Gauge equivariant mesh cnns: Anisotropic convolutions on geometric graphs. In *International Conference on Learning Representations*, 2021.

Dehmamy, N., Walters, R., Liu, Y., Wang, D., and Yu, R. Automatic symmetry discovery with lie algebra convolutional network. In *Advances in Neural Information Processing SystemsNeural*, 2021.

Dieleman, S., De Fauw, J., and Kavukcuoglu, K. Exploiting cyclic symmetry in convolutional neural networks. In *International Conference on Machine Learning*, 2016.

Esteves, C., Allen-Blanchette, C., Makadia, A., and Daniilidis, K. Learning $\mathrm{SO}(3)$ equivariant representations with spherical cnns. In *European Conference on Computer Vision*, 2018.

Esteves, C., Makadia, A., and Daniilidis, K. Spin-weighted spherical cnns. In *Advances in Neural Information Processing Systems*, 2020.

Finzi, M., Welling, M., and Gordon Wilson, A. A practical method for constructing equivariant multilayer perceptrons for arbitrary matrix groups. In *International Conference on Machine Learning*, 2021.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

He, L., Dong, Y., Wang, Y., Tao, D., and Lin, Z. Gauge equivariant transformer. In *Advances in Neural Information Processing Systems*, 2021.

Keller, T. A. and Welling, M. Topographic vaes learn equivariant capsules. In *Advances in Neural Information Processing Systems*, 2021.

Kondor, R. and Trivedi, S. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *International Conference on Machine Learning*, 2018.

Luo, D., Yuan, S., Stokes, J., and Clark, B. Gauge equivariant neural networks for $2+1du(1)$ gauge theory simulations in hamiltonian formulation. In *NeurIPS Workshop on AI for Science*, 2022.

Luo, D., Chen, Z., Hu, K., Zhao, Z., Mikyoung, V., and Clark, B. Gauge-invariant and anyonic-symmetric autoregressive neural network for quantum lattice models. *Physical Review Research*, 5:013216, 2023.

Nagaosa, N. *Quantum field theory in condensed matter physics*. Springer Science & Business Media, 2013.

Nakahara, M. *Geometry, topology and physics*. CRC press, 2018.

Park, J. Y., Wong, L., and Walters, R. Modeling dynamics over meshes with gauge equivariant nonlinear message passing. In *Advances in Neural Information Processing Systems*, 2023.

Quessard, R., Barrett, T., and Clements, W. Learning disentangled representations and group structure of dynamical environments. In *Advances in Neural Information Processing Systems*, 2020.

Ruhe, D., Brandstetter, J., and Forré, P. Clifford group equivariant neural networks. In *Advances in Neural Information Processing Systems*, 2023a.

Ruhe, D., Gupta, J., Keninck, S. d., Welling, M., and Brandstetter, J. Geometric clifford algebra networks. In *International Conference on Machine Learning*, 2023b.

Sanborn, S., Shewmake, C., Olshausen, B., and Hillar, C. Bispectral neural networks. In *International Conference on Learning Representations*, 2023.

Sosnovik, I., Szmaja, M., and Smeulders, A. Scale-equivariant steerable networks. In *International Conference on Learning Representations*, 2020.

Wang, R., Walters, R., and Yu, R. Approximately equivariant networks for imperfectly symmetric dynamics. In *International Conference on Machine Learning*, 2022.

Weiler, M. and Cesa, G. General $e(2)$-equivariant steerable cnns. In *Advances in Neural Information Processing Systems*, 2020.

Weiler, M., Geiger, M., Welling, M., Boomsma, W., and Cohen, T. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. In *Advances in Neural Information Processing Systems*, 2018.

Yang, J., Dehmamy, N., Walters, R., and Yu, R. Latent space symmetry discovery. In *arXiv*, 2023a.

Yang, J., Walters, R., Dehmamy, N., and Yu, R. Generative adversarial symmetry discovery. In *International Conference on Machine Learning*, 2023b.

Zhdanov, M., Ruhe, D., Weiler, M., Lucic, A., Brandstetter, J., and Forré, P. Clifford-steerable convolutional neural networks. In *arXiv*, 2024.

A. XY Model Gauge Symmetry

In the XY model on a discrete lattice we defined the energy as

$$H(\boldsymbol{s}, A) = \frac{1}{2}\sum_{ij} J_{ij}\|(1+A_{ij})\boldsymbol{s}_i - \boldsymbol{s}_j\|^2 - N \tag{9}$$

and define the local transformations of $\boldsymbol{s}$ as

$$\boldsymbol{s}'_j = e^{i\alpha_j}\boldsymbol{s}_j. \tag{10}$$

Here, we will derive how A should transform such that the energy remains invariant. Thus, we want

$$H(\boldsymbol{s}, A) = H(\boldsymbol{s}', A') \tag{11}$$

We will match every term in the sum, such that for every pair i, j we have

$$\begin{aligned}
\|(1+A_{ij})\boldsymbol{s}_i - \boldsymbol{s}_j\|^2 &= \|(1+A'_{ij})\boldsymbol{s}'_i - \boldsymbol{s}'_j\|^2 \\
&= \|(1+A'_{ij})e^{i\alpha_i}\boldsymbol{s}_i - e^{i\alpha_j}\boldsymbol{s}_j\|^2 \\
&= \|e^{i\alpha_j}\left(e^{-i\alpha_j}(1+A'_{ij})e^{i\alpha_i}\boldsymbol{s}_i - \boldsymbol{s}_j\right)\|^2 \\
&= \|e^{-i\alpha_j}(1+A'_{ij})e^{i\alpha_i}\boldsymbol{s}_i - \boldsymbol{s}_j\|^2
\end{aligned} \tag{12}$$

From this we can equate the gauge parts and obtain

$$\begin{aligned}
1+A_{ij} &= e^{-i\alpha_j}(1+A'_{ij})e^{i\alpha_i} \\
1+A'_{ij} &= e^{i\alpha_j}(1+A_{ij})e^{-i\alpha_i} \\
A'_{ij} &= e^{i\alpha_j}(1+A_{ij})e^{-i\alpha_i} - 1
\end{aligned} \tag{13}$$

which is shows how the gauge field transforms. We can readily verify that when all $\alpha_i = 0$, we recover $A' = A$. Also, note that since A has to take values in the Lie algebra of the gauge group, $A_{ij} = M_{ij}L$, where $L = i = \sqrt{-1}$ is the basis of $\mathfrak{u}(1)$, the Lie algebra of $U(1)$. Since $U(1)$ is Abelian, $e^{i\alpha}A = Ae^{i\alpha}$ and so

$$A'_{ij} = e^{i(\alpha_j - \alpha_i)}(1+A_{ij}) - 1. \tag{14}$$

For global transformations, where $\alpha_i = \alpha$, we again recover $A' = A$. This only happens for Abelian gauge symmetries.

B. Implementation details

B.1. Effect of gauge transformations

In Section 3 we discussed that the spin configuration in the XY model changes as $\boldsymbol{s}_i \to e^{i\alpha_i}\boldsymbol{s}_i$ under a gauge transformation, where the gauge field α_i is smoothly changing. To generate the transformed XY configuration in Figure 1 we considered a simple gauge field that changes with a low frequency across the x and y dimensions. Specifically, let $\alpha_{i_x}, \alpha_{i_y}$ be the x, y components of α_i

$$\begin{aligned}
\alpha_{i_x} &= 0.25 \cdot \cos\left(2\pi\tfrac{x}{n}\right), \\
\alpha_{i_y} &= 0.25 \cdot \cos\left(2\pi\tfrac{2y}{n}\right),
\end{aligned}$$

where n is the size of the grid. In this example, we assume the two components act independently on the spin configuration, and thus the transformed angle of each vector is $\theta_g = \theta + \alpha_{i_x} + \alpha_{i_y}$. The two components of α_i alongside the effective gauge field are shown in Figure 4.

B.2. Data generation

Each data sample consists of a pair $(\boldsymbol{x}_i, y_i)$ with $\boldsymbol{x}_i \in \mathbb{R}^{100\times 100}$ and $y_i \in \mathbb{R}$ denoting the spin configuration at each grid point and the energy of the configuration. Dropping the sample index for clarity, each $x_{kl} \equiv \theta_{kl}$ represents the angle of the

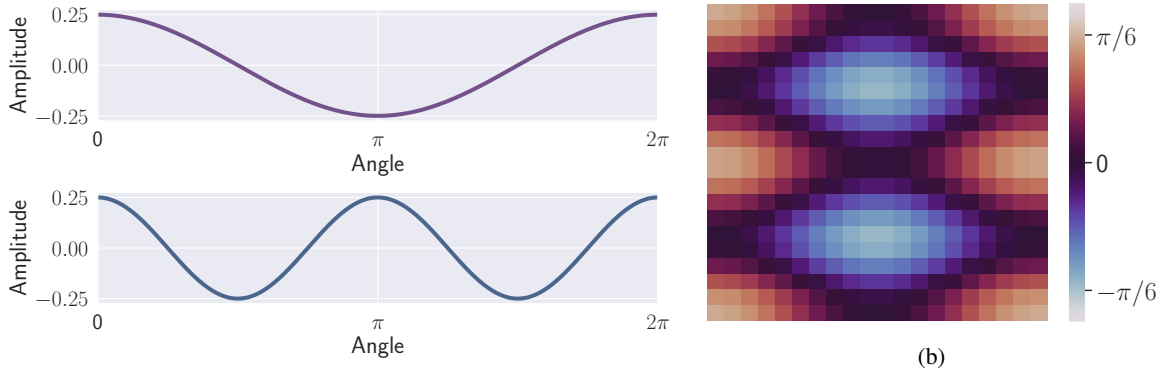

(a) (b)

Figure 4: (a) Low-frequency components of the gauge field along the x (top) and y (bottom) directions. (b) Effective gauge field, comprising of an angular offset $\theta_g = \theta + \alpha_{i_x} + \alpha_{i_y}$.

spin vector at that grid point and $\boldsymbol{s}_{kl} = \begin{bmatrix} \cos(\theta_{kl}) \\ \sin(\theta_{kl}) \end{bmatrix}$. Each angle is originally sampled as random variable $\theta_{kl} \sim 2\pi U$ with $U \sim \text{Unif}$. We then run $\eta = 10,000$ gradient steps optimizing the energy of the configuration $H(\boldsymbol{s}) = -\sum_{ij} J_{ij} \boldsymbol{s}_i^T \boldsymbol{s}_j$.

The XY model has many low energy states and convergence to these states depends on the temperature of the system. In order to get a diverse distribution of energies suitable for a learning task, we need to sample different system temperatures. In order simulate different temperatures, we varied the learning rate of the optimizer $lr \in \{100, 10, 1, 0.1, 0.01\}$. Higher learning rates correspond to higher temperatures where the system is more likely to reach the ground state (where all charges are aligned). To ensure convergence of the descent procedure for all learning rates, we choose a target learning rate that showed consistent converging behavior ($lr_{\text{final}} = 0.01$) and designed a linear learning rate schedule where the target is reached after $\kappa = 100$ learning rate updates. Specifically, we set $\gamma = \left(\frac{lr}{lr_{\text{final}}}\right)^{\frac{1}{\kappa}}$ and we updated the learning rate every $\frac{\eta}{\kappa}$ gradient steps. We used the Adam optimizer and uniformly at random choose a learning rate for each data sample. The process to generate $N = 10,000$ samples took $\sim$10 hours on a NVIDIA GeForce RTX 3090 Ti.

B.3. XY model samples

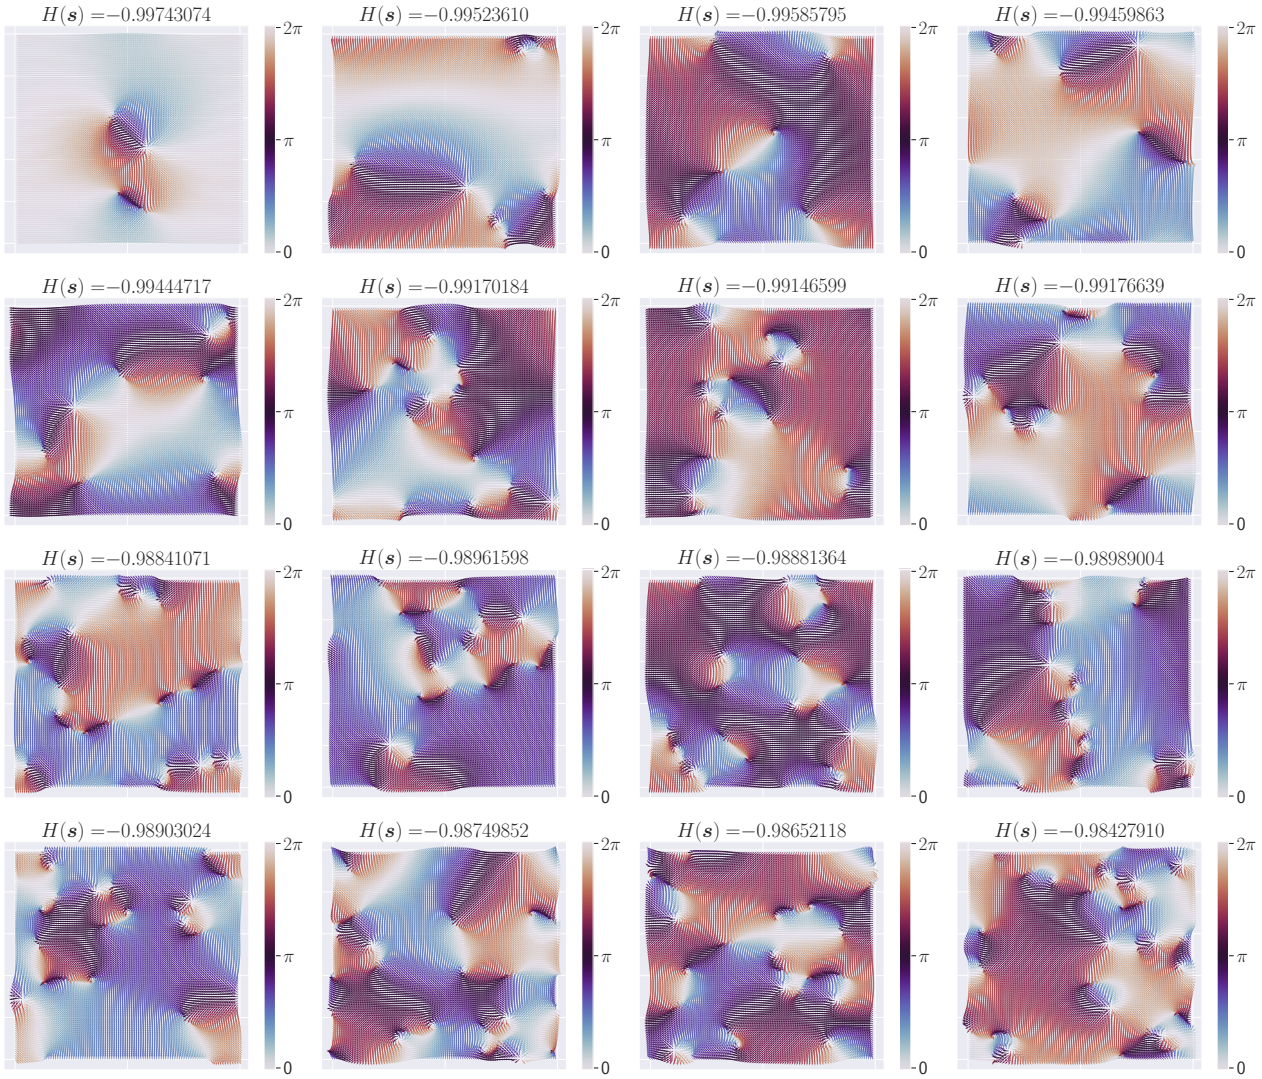

Figure 5: Final configurations in the XY model dataset alongside their energies (as the plot titles). Samples were picked at random from the full dataset ($N = 10,000$).