LORENTZ GROUP EQUIVARIANT AUTOENCODERS

Zichun Hao, Raghav Kansal, & Javier Duarte
Department of Physics
University of California San Diego
La Jolla, CA 92093, USA
{zihao,rkansal,jduarte}@ucsd.edu

Nadya Chernyavskaya
European Center for Nuclear Research (CERN)
1211 Geneva 23, Switzerland
nadezda.chernyavskaya@cern.ch

ABSTRACT

We develop the Lorentz group autoencoder (LGAE), an autoencoder that is equivariant with respect to the proper, orthochronous Lorentz group $\mathrm{SO}^+(3,1)$, with a latent space living in the representations of the group. We present our architecture and several experimental results on data at the Large Hadron Collider and find it outperforms a graph neural network baseline model on several compression, reconstruction, and anomaly detection tasks. The PyTorch code for our models is provided in Hao et al. (2022a).

1 INTRODUCTION

Deep neural networks (DNNs) are increasingly applied to the exabyte-scale volume of data produced in high energy physics (HEP), such as at the Large Hadron Collider (LHC). They have a variety of applications, ranging from classification and regression to anomaly detection and generative modeling (Guest et al., 2018; Radovic et al., 2018; Carleo et al., 2019; HEP ML Community, 2021). Recently, the most successful have been those incorporating key inductive biases of HEP data, such as infrared and colinear (IRC) safety in energy flow networks (Komiske et al., 2019) and permutation symmetry and sparsity of jet constituents via GNNs (Thais et al., 2022; Qu & Gouskos, 2020; Kansal et al., 2021). In addition to potentially higher performance, incorporating such biases can enable training models with less data and more interpretable models.

In this work, we explore incorporating another fundamental symmetry in HEP, equivariance to Lorentz transformations, into a novel autoencoder (AE) model for data compression and anomaly detection. Data compression is of utmost importance to store and process the exabytes of data to be produced in the coming decade at the LHC, and detecting anomalous collisions in a model-independent manner could prove to be a powerful tool to uncover new physics. AEs can perform data compression by learning an encoding of input data into a lower dimensional latent space. Moreover, by training them on typical "background" data alone, we can exploit the poor reconstruction of out-of-distribution signal to identify them as anomalies Govorkova et al. (2022); Pol et al. (2020). While there has been recent success in developing Lorentz-equivariant models for classifying collimated sprays of particles prevalent at the LHC, known as jets (Bogatskiy et al., 2020; Gong et al., 2022; Butter et al., 2018), to our knowledge, this work represents the first attempt to exploit Lorentz equivariance in compressing and detecting anomalous jets.

2 RELATED WORK

A neural network $\mathrm{NN} : V \to W$ is *equivariant* with respect to a group G if

$$\forall g \in G, v \in V : \mathrm{NN}(\rho_V(g) \cdot v) = \rho_W(g) \cdot \mathrm{NN}(v), \tag{1}$$

where $\rho_V : G \to \mathrm{GL}(V)$ and $\rho_W : G \to \mathrm{GL}(W)$ are representations of G in spaces V and W respectively. Equivariance has long been built into a number of successful DNN architectures, such as translation equivariance in convolutional neural networks, and permutation equivariance in graph neural networks (GNNs) (Bronstein et al., 2021). This has also been extended to a broader set of symmetries like the group of 3D rotations and translations $\mathrm{E}(3)$ (Thomas et al., 2018; Batzner et al., 2022), and recently, even transformations defined by the special orthochronous Lorentz group

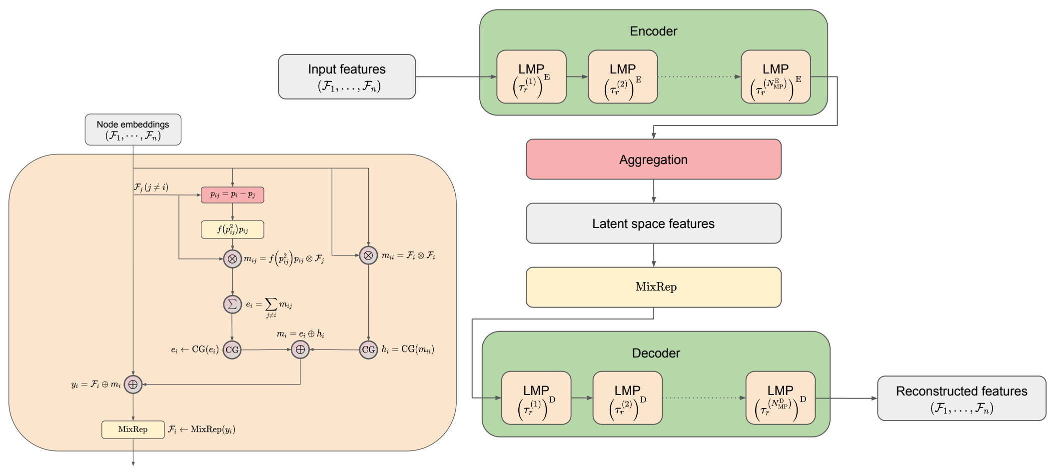

Figure 1: Individual Lorentz group equivariant message passing (LMP) layers are shown on the left, and the LGAE architecture is built out of LMPs on the right. Here, MixRep denotes the node-level operator that upsamples features in each representation space to τ_r channels, where r is the label for different components of the irreps; it appears as W in Eq. (3).

$\mathrm{SO}^+(3,1)$, which is relevant in high energy particle collisions. This group is generated by 3D spatial rotations and Lorentz boosts (transformations from one reference frame to another that is moving at a constant velocity relative to the former). This group expresses the fundamental symmetry of space and time and is respected by all known fundamental laws of nature, including the standard model of particle physics.

The Lorentz group network (LGN) (Bogatskiy et al., 2020) was the first DNN architecture developed to be covariant to the $\mathrm{SO}^+(3,1)$ group, with an architecture similar to that of a GNN, but operating entirely in Fourier space on objects in irreducible representations (irreps) of the Lorentz group. It uses tensor products between irreps and Clebsch–Gordan (CG) decompositions to introduce nonlinearities in the network. More recently, LorentzNet (Gong et al., 2022) uses a similar GNN framework for equivariance, with additional edge features—Minkowski inner products between node features—but restricted to only scalar and vector representations of the group. Both networks have been successful in jet classification, achieving state-of-the-art (SOTA) results.

3 LGAE Architecture

The LGAE (Figure 1), inspired by the LGN framework, is built out of Lorentz-group-equivariant message passing (LMP) layers comprising its encoder and decoder networks. Its input data is represented as a point cloud of N particles, or a "particle cloud", each associated with a 4-momentum vector (consisting of the particle's energy and 3D spatial momentum) and an arbitrary number of scalars representing physical features such as mass, charge, and spin.

The $(t+1)$-th LMP layer operation consists of message-passing between each pair of nodes in the particle cloud, with a message $m_{ij}^{(t)}$ to node i from node j (where $j \neq i$) and a self-interaction term m_{ii} defined as

$$m_{ij}^{(t)} = f\left(\left(p_{ij}^{(t)}\right)^2\right) p_{ij}^{(t)} \otimes \mathcal{F}_j^{(t)}, \quad m_{ii}^{(t)} = \mathcal{F}_i^{(t)} \otimes \mathcal{F}_i^{(t)}, \tag{2}$$

where $\mathcal{F}_i^{(t)}$ are the node features before the $(t+1)$-th layer, $p_{ij} = p_i - p_j$ is the difference between node four-vectors, p_{ij}^2 is the squared Minkowski norm of p_{ij}, and f is a learnable, differentiable function acting on Lorentz scalars. A CG decomposition, reducing the features to direct sums of irreps of $\mathrm{SO}^+(3,1)$, is performed on both terms before concatenating them to produce the aggregated message m_i for node $i \in \{1, \ldots, N_{\mathrm{particle}}\}$, which is then used to update the node's features:

$$m_i^{(t)} = \mathrm{CG}\left[m_{ii}^{(t)}\right] \oplus \mathrm{CG}\left[\sum_{j \neq i} m_{ij}^{(t)}\right], \quad \mathcal{F}_i^{(t+1)} = W\left(\mathcal{F}_i^{(t)} \oplus m_i^{(t)}\right), \tag{3}$$

where W is a node-wise operator with learnable parameters which linearly mixes features in the same representation space to the desired multiplicity.

The encoder first mixes each isotypic component in the input cloud via learned weights, to a chosen multiplicity. The resultant cloud is then processed through $N_{\mathrm{MP}}^{\mathrm{E}}$ LMP layers, after which node features are aggregated to the latent space by a component-wise minimum (min), maximum (max)—on the respective Lorentz invariants—or mean. We also find, empirically, interesting performance by simply concatenating isotypic components across each particle and linearly "mixing" them via a learned matrix as in Eq. (3), which thereby breaks the permutation symmetry.

The decoder recovers the N-particle cloud by acting on the latent space with N independent, learned linear operators, which mix components living in the same representations (the MixRep block in Fig. 1), an operation similar to Eq. 3, right. This cloud passes through $N_{\mathrm{MP}}^{\mathrm{D}}$ LMP layers, after which node features are mixed back to the input representation space by applying a linear mixing layer and truncating other isotypic components. We refer the interested reader to Bogatskiy et al. (2020); Hao et al. (2022b) for more detail on the architecture and theoretical background.

4 RESULTS

We experiment with and evaluate the performance of the LGAE on reconstruction and anomaly detection for simulated high-momentum jets. LGAE model results are presented using both the min-max (LGAE-Min-Max) and "mix" (LGAE-Mix) aggregation schemes for the latent space, which consists of varying numbers of complex Lorentz vectors—corresponding to different compression rate, where the compression rate is considered to be the total number of features in the latent space relative to the number of input features.

We compare the LGAE to a baseline GNN autoencoder model (referred to as "GNNAE") composed of fully-connected MPNNs adapted from Kansal et al. (2021). We experiment with two types of encodings: (1) particle-level (GNNAE-PL), as in the PGAE (Tsan et al., 2021) model, which compresses the features per node in the graph but retains the graph structure in the latent space, and (2) jet-level (GNNAE-JL), which averages the features across each node to form the latent space, as in the LGAE. Particle-level encodings produce better performance overall for the GNNAE, but the jet-level provides a more fair comparison with the LGAE, which uses jet-level encoding to achieve a high level of compression of the features.

4.1 RECONSTRUCTION

We evaluate the performance of the LGAE and GNNAE models, with the different aggregation schemes discussed, on the reconstruction of the particle and jet features of jets produced by quarks and gluons, or QCD jets, from the JETNET dataset Kansal et al. (2022). We consider relative transverse momentum $p_{\mathrm{T}}^{\mathrm{rel}} = {p_{\mathrm{T}}}^{\mathrm{particle}}/{p_{\mathrm{T}}}^{\mathrm{jet}}$ and relative angular coordinates $\eta^{\mathrm{rel}} = \eta^{\mathrm{particle}} - \eta^{\mathrm{jet}}$ and $\phi^{\mathrm{rel}} = \phi^{\mathrm{particle}} - \phi^{\mathrm{jet}} \pmod{2\pi}$ as each particle's features, and total jet mass, p_{T} and η as jet features. Each input sample x contains 30 particles, each with the three particle features, hence $x \in \mathbb{R}^{90}$.

The reconstructed particle-level feature distributions from each model, along with the target, are shown in Fig. 2. We observe that the out of the two permutation invariant jet-level compression models, LGAE-Min-Max outperforms GNNAE-JL, while LGAE-Mix is the best performing overall.

4.2 ANOMALY DETECTION

We test the performance of these models as unsupervised anomaly detection algorithms by pre-training them solely on QCD jets and then using the Chamfer distance between the input and in-

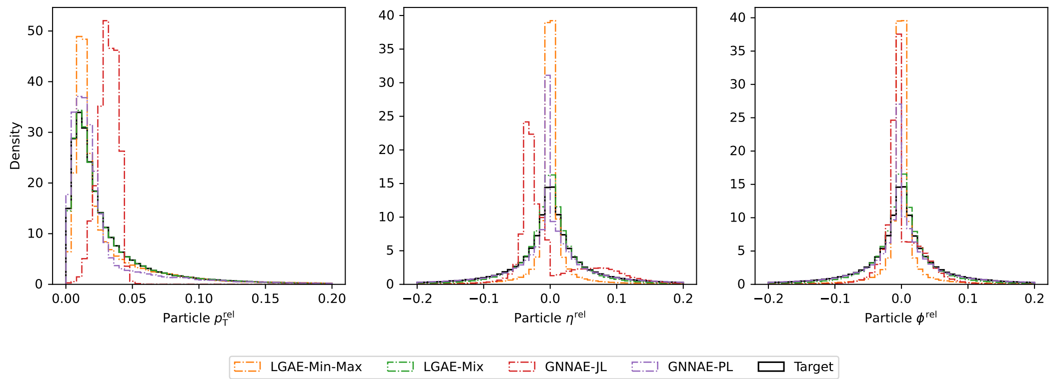

Figure 2: The reconstructed particle-level feature $(p_{\mathrm{T}}^{\mathrm{rel}}, \eta^{\mathrm{rel}}, \phi^{\mathrm{rel}})$ distribution by each model.

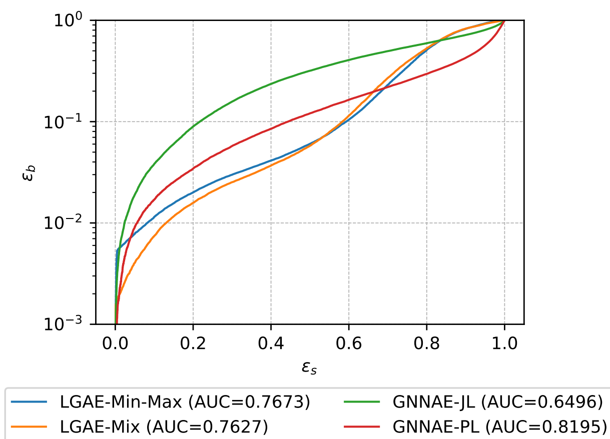

Figure 3: ROC curves and AUC per model trained on QCD jets, for a signal composed of top quark, W boson, and Z boson jets.

ferenced output as a discriminating variable. We consider top quark, W boson, and Z boson jets as potential signals and QCD as the "background".

Receiver operating characteristic (ROC) curves showing the signal efficiencies (ε_s) versus background efficiencies (ε_b) for individual and combined signals are shown in Figure 3. We see that in general LGAE models have significantly higher signal efficiencies than GNNAEs for all signals when rejecting $\varepsilon_b > 90\%$ of the background (which is the minimum level we typically require in HEP), and LGAE-Mix consistently performs better than LGAE-Min-Max.

5 Conclusion

We develop the Lorentz group autoencoder (LGAE), an autoencoder model equivariant to Lorentz transformations. We argue that incorporating this key inductive bias of high energy physics (HEP) data can have a significant impact on the performance, efficiency, and interpretability of machine learning models in HEP. We apply the LGAE to tasks of compression and reconstruction of input quantum chromodynamics (QCD) jets, and of identifying out-of-training-distribution anomalous top quark, W boson, and Z boson jets. We report excellent performance in comparison to a baseline graph neural network autoencoder (GNNAE) model, with the LGAE outperforming the GNNAE on several key metrics. The LGAE opens many promising avenues in terms of both performance and model interpretability, with the exploration of new datasets, higher-order Lorentz group representations, analysis of the LGAE's latent space, and analysis of the interplay between reconstruction and anomaly detection performance all exciting possibilities for future work.

Acknowledgments

Z. H. thanks the UC San Diego Faculty Mentor Program for supporting this research. R. K. was partially supported by the LHC Physics Center at Fermi National Accelerator Laboratory, managed and operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy (DOE). J. D. is supported by the DOE, Office of Science, Office of High Energy Physics Early Career Research program under Award No. DE-SC0021187, the DOE, Office of Advanced Scientific Computing Research under Award No. DE-SC0021396 (FAIR4HEP), and the NSF HDR Institute for Accelerating AI Algorithms for Data Driven Discovery (A3D3) under Cooperative Agreement OAC-2117997. N. C. was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement No. 772369). This work was performed using the Pacific Research Platform Nautilus HyperCluster supported by NSF awards CNS-1730158, ACI-1540112, ACI-1541349, OAC-1826967, the University of California Office of the President, and the University of California San Diego's California Institute for Telecommunications and Information Technology/Qualcomm Institute. Thanks to CENIC for the 100 Gpbs networks.

References

Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P. Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E. Smidt, and Boris Kozinsky. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13 (1):2453, 2022. doi: 10.1038/s41467-022-29939-5. URL `https://doi.org/10.1038/s41467-022-29939-5`.

Alexander Bogatskiy, Brandon Anderson, Jan Offermann, Marwah Roussi, David Miller, and Risi Kondor. Lorentz group equivariant neural network for particle physics. 119:992, 2020. URL `https://proceedings.mlr.press/v119/bogatskiy20a.html`.

Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. 2021.

Anja Butter, Gregor Kasieczka, Tilman Plehn, and Michael Russell. Deep-learned Top Tagging with a Lorentz Layer. *SciPost Phys.*, 5:028, 2018. doi: 10.21468/SciPostPhys.5.3.028. URL `https://doi.org/10.21468/SciPostPhys.5.3.028`.

Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Rev. Mod. Phys.*, 91(4):045002, 2019. doi: 10.1103/RevModPhys.91.045002. URL `https://doi.org/10.1103/RevModPhys.91.045002`.

Shiqi Gong, Qi Meng, Jue Zhang, Huilin Qu, Congqiao Li, Sitian Qian, Weitao Du, Zhi-Ming Ma, and Tie-Yan Liu. An efficient Lorentz equivariant graph neural network for jet tagging. *JHEP*, 07:030, 2022. doi: 10.1007/JHEP07(2022)030. URL `https://doi.org/10.1007/JHEP07(2022)030`.

Ekaterina Govorkova et al. Autoencoders on field-programmable gate arrays for real-time, unsupervised new physics detection at 40 MHz at the Large Hadron Collider. *Nat. Mach. Intell.*, 4:154, 2022. doi: 10.1038/s42256-022-00441-3. URL `https://doi.org/10.1038/s42256-022-00441-3`.

Dan Guest, Kyle Cranmer, and Daniel Whiteson. Deep Learning and its Application to LHC Physics. *Ann. Rev. Nucl. Part. Sci.*, 68:161, 2018. doi: 10.1146/annurev-nucl-101917-021019. URL `https://doi.org/10.1146/annurev-nucl-101917-021019`.

Zichun Hao, Raghav Kansal, Javier Duarte, and Nadezda Chernyavskaya. Lorentz group equivariance autoencoders. Dec 2022a. doi: 10.5281/zenodo.7434838. URL `https://doi.org/10.5281/zenodo.7434838`.

Zichun Hao, Raghav Kansal, Javier Duarte, and Nadezda Chernyavskaya. Lorentz group equivariant autoencoders, 2022b. URL `https://arxiv.org/abs/2212.07347`.

HEP ML Community. A living review of machine learning for particle physics, 2021. URL https://iml-wg.github.io/HEPML-LivingReview/.

Raghav Kansal, Javier Duarte, Hao Su, Breno Orzari, Thiago Tomei, Maurizio Pierini, Mary Touranakou, Jean-Roch Vlimant, and Dimitrios Gunopulos. Particle cloud generation with message passing generative adversarial networks. In M. Ranzato et al. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 23858. Curran Associates, Inc., 2021. URL https://papers.nips.cc/paper/2021/hash/c8512d142a2d849725f31a9a7a361ab9-Abstract.html.

Raghav Kansal, Javier Duarte, Hao Su, Breno Orzari, Thiago Tomei, Maurizio Pierini, Mary Touranakou, Jean-Roch Vlimant, and Dimitrios Gunopulos. JetNet dataset, 2022. URL https://doi.org/10.5281/zenodo.6975118.

Patrick T. Komiske, Eric M. Metodiev, and Jesse Thaler. Energy Flow Networks: Deep Sets for Particle Jets. *JHEP*, 01:121, 2019. doi: 10.1007/JHEP01(2019)121. URL https://doi.org/10.1007/JHEP01(2019)121.

Adrian Alan Pol, Victor Berger, Gianluca Cerminara, Cecile Germain, and Maurizio Pierini. Anomaly Detection With Conditional Variational Autoencoders. In *18th International Conference on Machine Learning and Applications*, 2020. URL https://arxiv.org/abs/2203.128522010.05531.

Huilin Qu and Loukas Gouskos. ParticleNet: Jet Tagging via Particle Clouds. *Phys. Rev. D*, 101:056019, 2020. doi: 10.1103/PhysRevD.101.056019. URL https://doi.org/10.1103/PhysRevD.101.056019.

Alexander Radovic, Mike Williams, David Rousseau, Michael Kagan, Daniele Bonacorsi, Alexander Himmel, Adam Aurisano, Kazuhiro Terao, and Taritree Wongjirad. Machine learning at the energy and intensity frontiers of particle physics. *Nature*, 560(7716):41, 2018. doi: 10.1038/s41586-018-0361-2. URL https://doi.org/10.1038/s41586-018-0361-2.

Savannah Thais, Paolo Calafiura, Grigorios Chachamis, Gage DeZoort, Javier Duarte, Sanmay Ganguly, Michael Kagan, Daniel Murnane, Mark S. Neubauer, and Kazuhiro Terao. Graph Neural Networks in Particle Physics: Implementations, Innovations, and Challenges. In *2022 Snowmass Summer Study*, 3 2022. URL https://arxiv.org/abs/2203.12852.

Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds. 2018. URL https://arxiv.org/abs/1802.08219.

Steven Tsan, Raghav Kansal, Anthony Aportela, Daniel Diaz, Javier Duarte, Sukanya Krishna, Farouk Mokhtar, Jean-Roch Vlimant, and Maurizio Pierini. Particle Graph Autoencoders and Differentiable, Learned Energy Mover's Distance. In *4th Machine Learning and the Physical Sciences Workshop at the 35th Conference on Neural Information Processing Systems*, 2021. URL https://ml4physicalsciences.github.io/2021/files/NeurIPS_ML4PS_2021_98.pdf.