Granularity Beyond Hardware: Super-Resolution for Enhanced Particle Reconstruction in Calorimeters

Nilotpal Kakati, Etienne Dreyer, Eilam Gross
Weizmann Institute of Science
Rehovot, Israel
{nilotpal.kakati, etienne.dreyer, eilam.gross}@weizmann.ac.il

Abstract

The spatial resolution of calorimeters is a crucial parameter in particle detector design which is often constrained by cost and construction complexity. We propose machine learning based super-resolution as a software technique to increase effective calorimeter granularity, enhancing a detector's performance with zero changes to hardware. Upsampling is performed with a transformer-based continuous normalizing flow conditioned on low-granularity calorimeter data. We showcase the impact of our approach on particle reconstruction using a generic particle flow algorithm based on machine learning. Our results demonstrate that super-resolution can be readily applied at current and future collider experiments.

1 Introduction

Calorimeters play a central role in modern particle detectors by converting particles into a cascade of energy deposits. In dense collision environments such as the LHC [1], reconstruction techniques like particle flow (PFlow) algorithms [2–4] rely on calorimeter granularity to resolve nearby particles. In practice, calorimeter granularity is subject to budget and engineering constraints, leading to a compromise in overall detector performance. As an alternative to enhanced hardware, we propose a fully software-based Super-resolution (SR) framework that virtually upsamples calorimeter granularity from low-resolution (LR) to high-resolution (HR) as shown in Figure 1. By training a neural network to predict the HR calorimeter readout from the LR inputs, our technique offers a cost-effective path to upgrade overall detector performance. We demonstrate the first integration of SR directly into a PFlow algorithm, directly enhancing the quality of reconstructed particles.

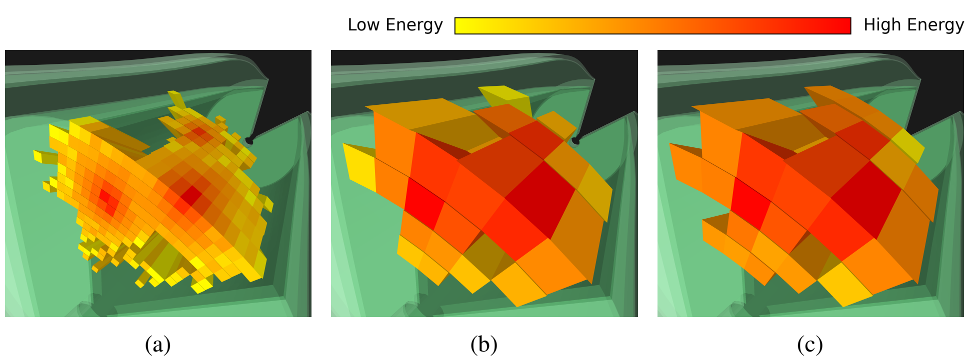

Figure 1: Event display of one electron and two photons in COCOA, illustrating the dataset creation procedure. Here we see (a) the truth high-resolution configuration, (b) the truth low-resolution configuration (c) the measured low-resolution configuration which also contains noise.

Machine Learning and the Physical Sciences Workshop, NeurIPS 2025.

Related Work SR techniques have been extensively studied in the field of image processing [5–11]. A proof-of-concept for calorimeter SR was first explored in [12] using a simplified detector and physics setup. Since then, related techniques have been applied to specific tasks in particle physics, such as enhancing jet images with generative models [13], using normalizing flows to upsample calorimeter showers for fast simulation [14], and improving reconstruction in neutrino telescopes [15].

Our Contributions We present the first end-to-end pipeline that uses SR to directly enhance particle reconstruction. Our key technical contributions are the following: (1) A generative model trained with a dual objective to simultaneously denoise and increase the spatial resolution of calorimeter energy deposits. (2) A novel graph-based architecture for this model that is inherently suited for the sparse signals and irregular geometries of modern calorimeters.

2 Datasets

To generate the necessary training data, we use the COCOA [16], a generic, GEANT4-based[17–19] calorimeter simulation for LHC-like collisions. The core of our dataset creation, illustrated in Figure 1, is a procedure to record the same underlying particle shower at two different calorimeter granularities. First, particles are simulated in a high-resolution (HR) detector configuration with no electronic noise. The energies in these HR cells are then summed to create a corresponding low-resolution (LR) representation, to which realistic electronic noise is subsequently added. For this study, we focus exclusively on electromagnetic showers, which are contained within the first three calorimeter layers.

We employ two distinct datasets. The first is a simplified set of single electrons with transverse momentum (p_{T}) of 50 GeV, where the resolution is downsampled by a factor of two in the two lateral coordinates, pseudorapidity (η) and azimuthal angle (ϕ). The second, more challenging dataset features multi-particle signatures containing an electron and up to three photons, specifically designed to have small angular separations on the scale of the calorimeter cells. For this multi-particle set, the resolution is downsampled by a more aggressive factor of four, posing a greater challenge to the SR task.

3 Methods

To achieve super-resolution, we employ a generative model based on Continuous Normalizing Flow (CNF) [20]. The model architecture, illustrated in Figure 2, operates on graph-based representations of the calorimeter data. Each LR cell is upsampled into a finer grid of HR cells (without noise),

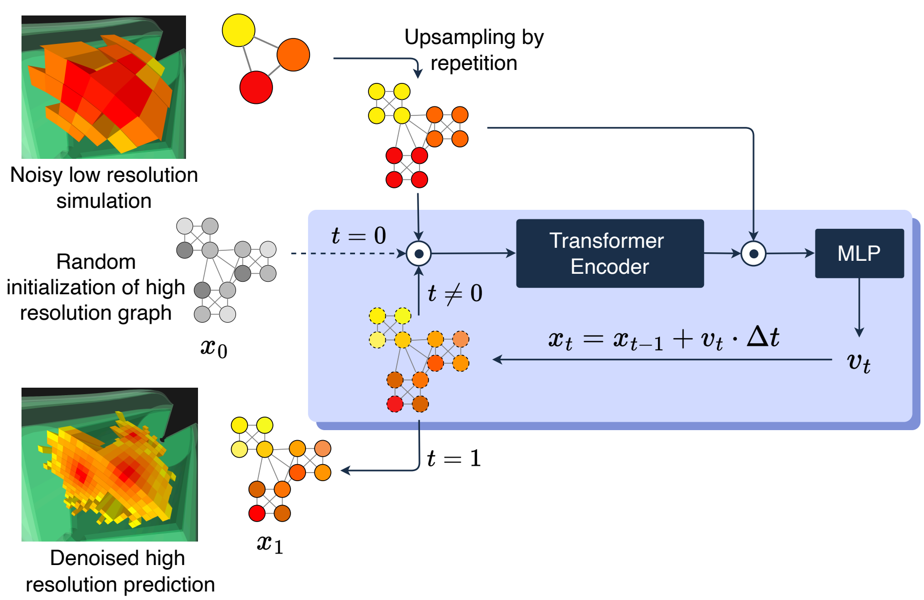

Figure 2: The super-resolution architecture. $\odot$ indicates concatenation.

with the original LR energy (including noise) serving as a conditional input to guide the generative process. A transformer network employing DiT [21] is then trained with a conditional flow matching objective [22], evolving a Gaussian base distribution into the final HR energy distribution. Since the CNF-based generation is inherently stochastic, we use ensemble sampling – averaging the outputs from multiple passes – to obtain a stable and precise prediction.

To evaluate the impact of our approach, we utilize a novel PFlow model inspired by HGPflow [23]. This model uses a supervised attention matrix to relate each calorimeter cell to each particle candidate in terms of energy fractions. The energy and direction of reconstructed particles is then computed as an attention-weighted sum over cell kinematic features. By comparing the performance of this PFlow model when trained on the original LR data and on our generated HR data, we investigate the impact of SR on reconstruction quality.

4 Results

Single electron Our super-resolution approach yields a significant improvement in energy resolution. For single electron events, the width of the relative residual energy distribution, characterized by its standard deviation, improved by approximately 40% using our predicted HR cells compared to the initial LR energies. Moreover, we investigated whether the network learns details about the shower substructure by computing the energy correlator observables C_2, C_3 and D_2 [24, 25] at the cell level. As shown in Figure 3, the distributions for the predicted HR cells show remarkable agreement with the ground truth compared to the LR results.

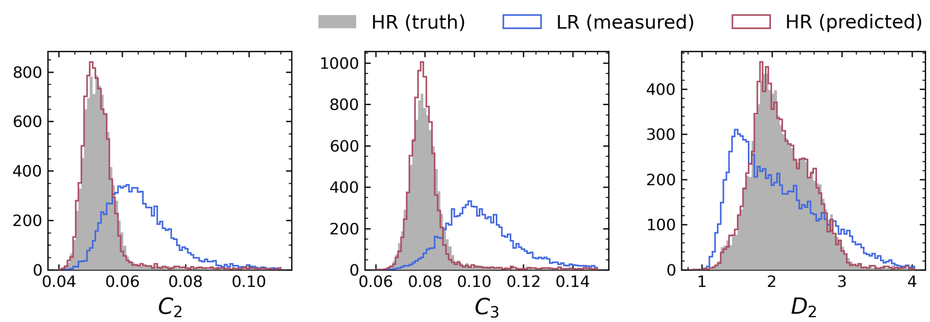

Figure 3: Substructure distributions for the truth HR, measured LR and predicted HR cell energies.

Multiple particles The primary benefit of our SR technique is its impact on PFlow reconstruction. A PFlow network, trained on truth HR cells and performing inference on our predicted HR cells, identifies the number of particles with up to 4% higher accuracy than an identically configured PFlow network, trained and inferred on the original LR data. This enhancement extends to kinematic reconstruction, as shown in Figure 4. The particle energy relative residual distribution for the model using predicted HR cells for inference is substantially narrower and more accurately centered at zero.

The event display in Figure 5 shows the improvement qualitatively. The PFlow model using our HR predictions successfully identifies three distinct particles with accurate energies, while the LR-based model incorrectly merges two of the particles.

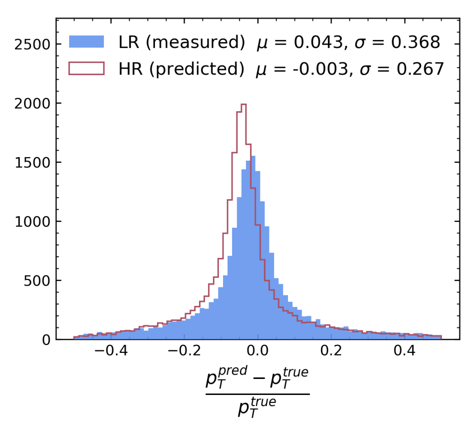

Figure 4: Relative p_{T} residual for PFlow models inferred on LR (original) vs. HR (predicted) cells.

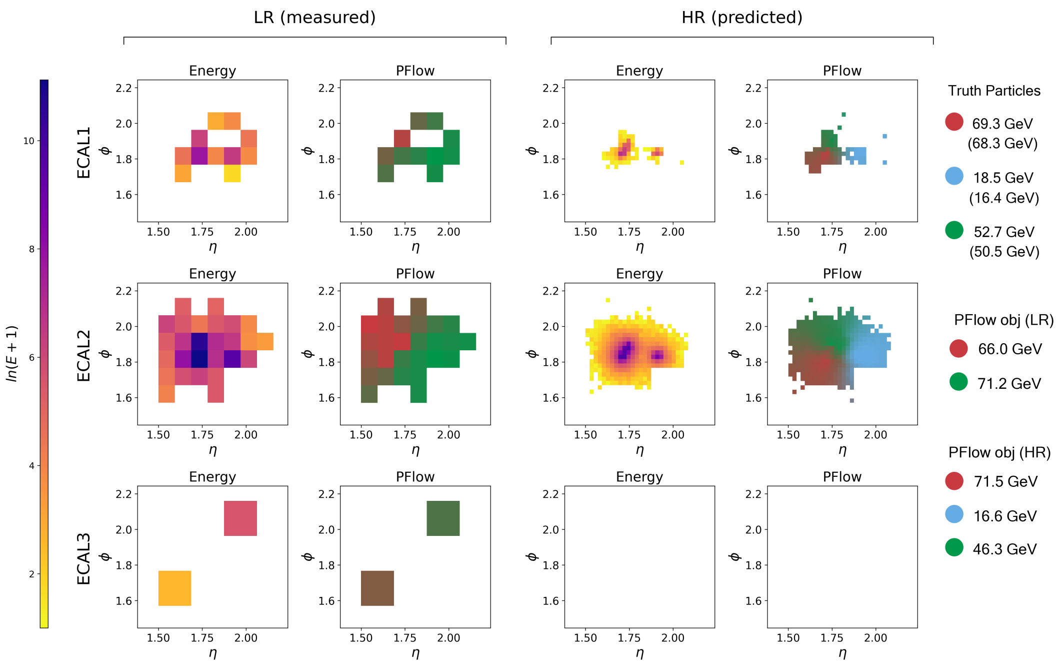

Figure 5: Event display illustrating the application of super-resolution on particle reconstruction, as well as the interpretability of the particle flow predictions. Each cell is assigned a fractional weight based on the predicted attention score, which determine the cell colors. The legend shows the truth and predicted particles with their energies. The total deposited energy of each particle is also shown in parentheses for the truth particles. The particle flow model trained on the predicted high-resolution cells correctly identifies all three particles with very accurate energy estimations, while the model trained on low-resolution measured cells fails to disentangle the event into three particles and predicts only two particles.

5 Discussion

In this work, we have demonstrated for the first time that SR can be readily integrated into a standard LHC-like reconstruction pipeline to yield significant performance gains. Our approach successfully replicates HR cell distributions, leading to improved accuracy in shower substructure and downstream PFlow reconstruction.

Ensuring that HR features are physically meaningful and not mere "hallucinations" will be important in real-world applications. During training, the model learns the patterns characterizing electromagnetic showers by modeling the conditional probability distribution $p\,(HR \,|\, LR)$. This behavior is verifiable in simulation but not in real experimental data where no ground truth exists. Therefore, thorough calibration and uncertainty quantification will be important for deploying robust SR in experiments.

Furthermore, SR should be viewed not just as a preprocessing step but as a powerful auxiliary objective. Training the model to reconstruct fine-grained details forces it to extract physically salient features, which explains why the downstream PFlow algorithm benefits so significantly. Future work can apply SR in the longitudinal direction and extend this framework to more complex scenarios, including hadronic showers and full collision events. We envision that SR techniques will bring significant improvements to current experiments while guiding the design and optimization of future detectors.

Code and data availability

To facilitate reproducibility, the code used for the analyses presented in this paper is available on GitHub: https://github.com/nilotpal09/SuperResolutionHEP. The dataset generated and analysed during this study can be accessed through Zenodo (DOI: https://doi.org/10.5281/zenodo.15582324)

Acknowledgments

We would like to thank Dmitrii Kobilianskii for fruitful discussions on flow matching implementation, and Kyle Cranmer and Yihui Ren for their valuable insights. EG also acknowledges support from ISF Grant 2871/19, the BSF-NSF Grant 2020780 and the Weizmann Institute of Science and and Mohamed bin Zayed University of Artificial Intelligence collaboration grant. We also thank the Benoziyo center for High Energy Physics for supporting our research.

References

[1] Lyndon Evans and Philip Bryant. LHC Machine. *JINST*, 3:S08001, 2008.

[2] ATLAS collaboration. Jet reconstruction and performance using particle flow with the ATLAS Detector. *Eur. Phys. J. C*, 77(7):466, 2017.

[3] CMS collaboration. Particle-flow reconstruction and global event description with the CMS detector. *JINST*, 12(10):P10003, 2017.

[4] M. A. Thomson. Particle Flow Calorimetry and the PandoraPFA Algorithm. *Nucl. Instrum. Meth. A*, 611:25–40, 2009.

[5] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):295–307, 2016.

[6] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1646–1654, 2016.

[7] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1132–1140, 2017.

[8] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 105–114, 2017.

[9] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 294–310, Cham, 2018. Springer International Publishing.

[10] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In Laura Leal-Taixé and Stefan Roth, editors, *Computer Vision – ECCV 2018 Workshops*, pages 63–79, Cham, 2019. Springer International Publishing.

[11] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4713–4726, 2023.

[12] Francesco Armando Di Bello, Sanmay Ganguly, Eilam Gross, Marumi Kado, Michael Pitt, Lorenzo Santi, and Jonathan Shlomi. Towards a computer vision particle flow. *The European Physical Journal C*, 81(2):107, Feb 2021.

[13] Pierre Baldi, Lukas Blecher, Anja Butter, Julian Collado, Jessica N. Howard, Fabian Keilbach, Tilman Plehn, Gregor Kasieczka, and Daniel Whiteson. How to GAN Higher Jet Resolution. *SciPost Phys.*, 13:064, 2022.

[14] Ian Pang, David Shih, and John Andrew Raine. Calorimeter shower superresolution. *Phys. Rev. D*, 109:092009, May 2024.

[15] Felix J. Yu, Nicholas Kamp, and Carlos A. Argüelles. Enhancing events in neutrino telescopes through deep learning-driven super-resolution, 2024.

[16] F. A. Di Bello, A. Charkin-Gorbulin, K. Cranmer, E. Dreyer, S. Ganguly, et al. Configurable calorimeter simulation for AI applications. *Machine Learning: Science and Technology*, 4(3):035042, 2023.

[17] S. Agostinelli et al. GEANT4–a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250–303, 2003.

[18] John Allison et al. Geant4 developments and applications. *IEEE Trans. Nucl. Sci.*, 53:270, 2006.

[19] J. Allison et al. Recent developments in Geant4. *Nucl. Instrum. Meth. A*, 835:186–225, 2016.

[20] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

[21] William S. Peebles and Saining Xie. Scalable diffusion models with transformers. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4172–4182, 2022.

[22] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.

[23] Nilotpal Kakati, Etienne Dreyer, Anna Ivina, Francesco Armando Di Bello, Lukas Heinrich, Marumi Kado, and Eilam Gross. HGPflow: extending hypergraph particle flow to collider event reconstruction. *Eur. Phys. J. C*, 85(8):847, 2025.

[24] Andrew J. Larkoski, Gavin P. Salam, and Jesse Thaler. Energy Correlation Functions for Jet Substructure. *JHEP*, 06:108, 2013.

[25] Andrew J. Larkoski, Ian Moult, and Duff Neill. Power Counting to Better Jet Observables. *JHEP*, 12:009, 2014.