



DATASET ARTICLE • **OPEN ACCESS**

Aspen Open Jets: unlocking LHC data for foundation models in particle physics

To cite this article: Oz Amram *et al* 2025 *Mach. Learn.: Sci. Technol.* **6** 030601

View the [article online](#) for updates and enhancements.

You may also like

- [OmniJet: the first cross-task foundation model for particle physics](#)
Joschka Birk, Anna Hallin and Gregor Kasieczka
- [LHC hadronic jet generation using convolutional variational autoencoders with normalizing flows](#)
Breno Orzari, Nadezda Chernyavskaya, Raphael Cobe et al.
- [Mass conditioned and constrained normalising flows for jet generation](#)
B Käch, D Krücker, I Melzer-Pellmann et al.



DATASET ARTICLE

OPEN ACCESS

RECEIVED
22 January 2025REVISED
13 May 2025ACCEPTED FOR PUBLICATION
17 June 2025PUBLISHED
7 July 2025

Original Content from
this work may be used
under the terms of the
Creative Commons
Attribution 4.0 licence.

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Aspen Open Jets: unlocking LHC data for foundation models in particle physics

Oz Amram¹ , Luca Anzalone^{2,3} , Joschka Birk⁴ , Darius A Faroughy⁵ , Anna Hallin^{4,*} ,
Gregor Kasieczka^{4,6} , Michael Krämer⁷ , Ian Pang⁵ , Humberto Reyes-Gonzalez⁷ and David Shih⁵

¹ Fermi National Accelerator Laboratory, Batavia, IL 60510, United States of America

² Department of Physics and Astronomy (DIFA), University of Bologna, 40127 Bologna, Italy

³ Istituto Nazionale di Fisica Nucleare (INFN), 40127 Bologna, Italy

⁴ Institut für Experimentalphysik, Universität Hamburg, 22761 Hamburg, Germany

⁵ NHETC, Dept. of Physics and Astronomy, Rutgers University, Piscataway, NJ 08854, United States of America

⁶ Center for Data and Computing in Natural Sciences (CDCS), 22607 Hamburg, Germany

⁷ Institut für Theoretische Teilchenphysik und Kosmologie, RWTH Aachen University, 52074 Aachen, Germany

* Author to whom any correspondence should be addressed.

E-mail: anna.hallin@uni-hamburg.de, oamram@fnal.gov, luca.anzalone2@unibo.it, joschka.birk@uni-hamburg.de,
darius.faroughy@rutgers.edu, gregor.kasieczka@uni-hamburg.de, mkraemer@physik.rwth-aachen.de,
ian.pang@physics.rutgers.edu, humberto.reyes@rwth-aachen.de and shih@physics.rutgers.edu

Keywords: foundation models, particle physics, transfer learning, generation, jet physics, transformer, open data

Abstract

Foundation models are deep learning models pre-trained on large amounts of data which are capable of generalizing to multiple datasets and/or downstream tasks. This work demonstrates how data collected by the CMS experiment at the Large Hadron Collider can be useful in pre-training foundation models for HEP. Specifically, we introduce the ASPENOPENJETS (AOJs) dataset, consisting of approximately 178 M high p_T jets derived from CMS 2016 Open Data. We show how pre-training the OMNIJET- α foundation model on AOJs improves performance on generative tasks with significant domain shift: generating boosted top and QCD jets from the simulated JetClass dataset. In addition to demonstrating the power of pre-training of a jet-based foundation model on actual proton–proton collision data, we provide the ML-ready derived AOJs dataset for further public use.

1. Introduction

While particle physics has long used machine learning techniques and is leading the way in adopting modern deep learning methods, the development and application of powerful foundation models—which have transformed natural language processing and computer vision—is still in its early stages. For the purposes of this work, a foundation model will be a model that has been pre-trained on a large amount of data for a specific task, and that can be used for different tasks downstream [1]. The pre-training thus serves as a foundation, upon which the downstream tasks can be built. Well-known foundation models include large language models such as GPT-3 [2], BERT [3] and LLaMA [4], image models such as Imagen [5], and mixed-modality models such as CLIP [6] and ALIGN [7].

There are several advantages to the use of foundation models. One example is the potential to boost the performance on small datasets. Having access to a model pre-trained on a larger, similar dataset provides a ‘head start’ for the downstream task, allowing the model to focus on the intricacies of the smaller dataset rather than having to rediscover the general structure of the data. In addition, a foundation model can save both human and computational resources. While pre-training may be a resource-intensive task, the downstream models would require less training, less data, and less time spent on optimization compared to what would be required if these models were built from scratch. Another interesting aspect is that a trained foundation model could enable implicit sharing of data and resources between different parts of the research community, a sharing that would not otherwise be possible.

Foundation models are poised to play a central role in particle physics, especially in the context of experiments at the Large Hadron Collider (LHC). They are able to process the huge amounts of complex and

diverse data generated by the LHC and have the potential to integrate multiple modalities—such as images, high-level event features or particle four-vectors—allowing more comprehensive data analyses. Current work on foundation models for particle physics has mainly focused on different event generation and classification tasks [8–13], investigating pre-training strategies, encoding schemes and different architectures, and quantifying the transfer learning and generalization capabilities of these models.

In this paper, we explore the potential of using large quantities of Open Data from the LHC for pre-training foundation models. The LHC experiments have made significant amounts of data publicly available [14–18]. Foundation models can leverage the LHC open data for training and fine-tuning, potentially integrating different data modalities and data from different experiments. Furthermore, training foundation models on real data will result in a more accurate representation of particle physics processes than training on simulations. Finally, foundation models trained on a large variety of actual LHC data will help to promote transparency and collaboration in particle physics research.

Here we present the ASPENOPENJETS (AOJs) dataset and study the transfer learning of foundation models obtained with this dataset. The AOJ dataset has been constructed from the CMS 2016 ‘JetHT’ open data [16, 17] and consists of 178 M jets, making it the largest ML-ready dataset for LHC applications. Jets are important objects at the LHC, consisting of collimated sprays of particles. Previous open data studies of CMS jets have explored the substructure and energy flow of the jets [19–21] and the usefulness of open data for quark/gluon tagging [22], but this is the first study to prepare and make public a form of the CMS open data that allows easy use of the dataset for phenomenological studies, especially for training machine learning models.

As an initial proof-of-concept study, we pre-train the OMNIJET- α [11] foundation model on the AOJ dataset, which is expected to mostly contain QCD jets, and then fine-tune it to generate jets of different types from the JetClass [23] dataset of simulated jets. We do two such fine-tunings: first for JetClass QCD jets, which occupy a different kinematic region compared to the jets in AOJ, and then for top jets, which differ significantly in several features. We show that pre-training on AOJ indeed benefits the downstream task, making it easier for the model to generate these types of jets with fewer training examples compared to training from scratch.

The utility of CMS open data for pre-training has been previously explored in [8], using 1 M single-lepton events at the object level which were collected by CMS in 2015. Here we enlarge the scope and complexity by several orders of magnitude to encompass 178 M jet events at the full constituent level, that is, including not only the high-level features of the jets themselves, but also the features of all the particles inside the jet. Furthermore, the authors of [8] used event selections to shape the open data into a similar topology to the intended target simulated data. In our work, no such shaping is needed.

This paper is organized as follows. Section 2 describes the CMS open data and the method used to derive AOJ. It also includes a brief description of JetClass, which will be used for the downstream task. Our application example is described in section 3, detailing the method and results. Finally, we present our conclusions in section 4. Links to the codebase and the AOJ dataset itself can be found after the conclusions.

2. Data

2.1. CMS open data

The CMS experiment [24] is a large general-purpose detector at the LHC. The CMS collaboration has recently released 16.4 fb^{-1} of data from their 2016 runs (G and H) [16, 17] and a corresponding set of simulations. In contrast to previous data releases [14, 15, 25], this data has been provided in the simplified MINIAOD format [26], as well as the even further simplified NANOAOD format [27]. In these simplified analysis formats, the central CMS reconstruction steps have been applied, discarding much of the low-level content of the event in favor of a simplified description of the high-level objects (jets, electrons, muons, etc) commonly used in data analysis. These formats are therefore significantly easier to use for those outside of the collaboration.

The inner layers of CMS consist of a silicon pixel and silicon strip tracker used to measure charged particle trajectories. These are followed by a lead tungstate crystal electromagnetic calorimeter and a brass and scintillator hadronic calorimeter. These detector elements are enclosed in a large solenoid, which produces a magnetic field of 3.8 T. Muon stations are located outside the iron return yoke of the magnet. Events of interest are selected using a two-stage trigger system. The first stage (L1) is implemented in custom hardware to select potentially interesting events within a fixed latency of about $4 \mu\text{s}$ [28] using information from the calorimeters and muon stations. Events accepted by the L1 system are sent to the high-level trigger (HLT), which consists of a dedicated computing cluster which runs a speed-optimized version of the full reconstruction. The HLT uses this more accurate reconstruction of the event to select events to be saved at a rate of approximately 1 kHz [29].

Table 1. Breakdown of the expected number of jet types in the AOJ dataset from MC simulations.

Jet type	Approx. Number
QCD	1.8×10^8
W	6.3×10^5
Z	1.8×10^5
Top	1.3×10^5
$H \rightarrow b\bar{b}$	500

CMS uses a particle-flow algorithm [30], which combines the information from the different detector systems to reconstruct and identify each individual particle in an event. Particles are classified as either muons, electrons, photons, charged hadrons or neutral hadrons. Jets are clustered from the particle-flow candidates in an event using the anti- k_t jet finding algorithm [31, 32]. The typical large-radius jet size used in CMS for jet substructure searches and measurements is $R = 0.8$, commonly referred to as AK8 jets. For AK8 jets the pileup-per-particle identification algorithm (PUPPI) [33, 34] is used to mitigate the effect of pileup. The PUPPI algorithm assigns a weight to each particle based on the probability that it originates from the primary vertex or from pileup. The momentum of AK8 jets is determined from the weighted vector sum of all particles inside the jet. Additional corrections are applied based on simulation and data-driven measurements to better calibrate the measured jet momentum to the true momentum of the particle from the hard process [35].

2.2. Extracting AOJs

The NANOAOJ format does not natively store information of the jet constituents. However, an extended version (PFNANO), which does include this information, can be easily produced from MINIAOD using tools provided by CMS [36]. We therefore, for all data and simulations samples used, start with the MINIAOD format, process it to PFNANO and then apply our selections to select the jets of interest for our dataset.


In order to construct the largest sample of jets possible, we consider all events from the ‘JetHT’ [16, 17] Run-2 dataset released by CMS to this date. The ‘JetHT’ dataset includes all events that have been selected by one of the HLT triggers related to jet momenta or total event hadronic activity (HT). Events are required to either have a high scalar sum of jet transverse momenta, or contain at least one or two jets with high energy. For the exact details of all different trigger paths used, see [16, 17]. Because we do not specify a specific trigger, the sample consists of events from an inclusive ‘or’ of any of the available triggers of this dataset. We note that this choice works well for building a large dataset for the purposes of foundation model training, but may introduce subtle kinematic biases resulting from the combined thresholds of the different triggers, which would complicate its use in a physics-based analysis.

Events are required to pass several data-quality filters commonly used within CMS to remove events that are likely to cause problems in the reconstruction. AK8 jets are required to have transverse momentum greater than 300 GeV⁸. The jet p_T spectrum is steeply falling, and the 300 GeV threshold was chosen to balance having a large dataset of jets for the pre-training task with the consideration of matching the high-momentum kinematic regime typical of typical jet substructure studies. Jets are also required to have an absolute value of pseudorapidity (η) less than 2.5, so that they fall within the range of the CMS tracker. The AK8 jets are also required to pass standard CMS quality criteria used to reject leptons reconstructed as jets and jets that may be poorly reconstructed. We construct our sample out of all AK8 jets that pass these criteria, with no limitation on the number of selected jets per event.

Applying these selections to the CMS open data yields a dataset of 178 million jets, which we call the AOJs. Due to the large cross section of QCD processes, the vast majority of these jets are expected to originate from light quarks and gluons. A small fraction of the jets are expected to originate from the decays of boosted heavy resonances. A rough estimate of the number of boosted fully merged jets produced by W , Z , top, and Higgs bosons in the $b\bar{b}$ decay mode are estimated with Monte Carlo (MC) simulations provided by CMS [37], with the remainder assumed to originate from QCD. The numbers are reported in table 1. More details on the MC samples used are given in appendix.

For each jet we store its transverse momentum (p_T), pseudorapidity (η), and azimuthal angular coordinate (ϕ). We also store its mass, groomed with the softdrop algorithm [38] as computed within the CMS reconstruction. Up to 150 constituents of the jet are stored. For each constituent, its 4-momentum is stored in the format (p_x, p_y, p_z, E) . We additionally store its transverse impact parameter (d_0) and

⁸ The jet transverse momentum (p_T) distribution in figure 1 extends down to 250 GeV as it is computed directly from vector sum of the jet constituents. Whereas, 300 GeV threshold is applied to the jet after additional corrections have been applied to the jet energy.



longitudinal impact parameter (d_z) with their uncertainties, the charge of the candidate, its particle-ID (PID) in the PDG format⁹ [39] and its weight from the PUPPI algorithm. We also include additional jet substructure quantities computed within the CMS reconstruction, including the number of constituents in the jet, N -subjettiness variables [40], various jet-tagging observables from the CMS implementation of ParticleNet [41, 42] and a regression of the jet mass from ParticleNet [42].

In figure 1 we show the distributions (red histograms) of several particle and jet-level features, based on a sub-sample of 700 K jets. Interestingly, the effects of detector granularity are seen as spikes in the particle η and ϕ distributions. In figure 2 we plot a 2D histogram in the $\eta - \phi$ plane. A closer look reveals a grid-like structure corresponding to the granularity of the detector. Furthermore, there are vertical features indicating the endcaps at $|\eta| \sim 1.6$ and several ‘dead cells’ throughout the detector. The distribution of the subjettiness ratio τ_{32} is consistent with the expectation that the AOJ is predominantly composed of QCD jets. The last panel in figure 1 shows the average number of constituents per jet for each PID, accompanied by the corresponding 1σ standard deviation bands.

2.3. JetClass

The JetClass [23] dataset is based on simulations and was first introduced in [9]. It contains both particle-level and jet-level information for 125 M simulated jets, distributed over 10 different jet types and divided into a training (100 M), validation (5 M) and test (20 M) set. The simulation is performed with MADGRAPH5_AMC@NLO [43] for the hard process, followed by PYTHIA [44] for parton showering and hadronization, and DELPHES [45] with the CMS detector [24] card for the detector response. Jets are required to have a transverse momentum of $500 \text{ GeV} < p_T < 1000 \text{ GeV}$ and a pseudorapidity $\eta < 2$, which is more restrictive than the criteria used for AOJ. We overlay the JetClass QCD particle- and jet-level distributions (blue histograms) in figure 1 to facilitate easy comparison with the corresponding AOJ distributions.

⁹ Note that neutral hadrons (h^0) are assigned the PID = 130 of the neutral kaon K_L^0 while positively/negatively charged hadrons ($h^\pm$) are assigned PID = ± 211 of the charged pion.




Figure 2. The particle-level positional distribution of the ASPENOPENJETS in the $\eta - \phi$ plane, illustrating the granularity of the CMS detector. Note the appearance of vertical features indicating the endcaps at $|\eta| \sim 1.6$ and multiple ‘dead cells’ throughout the detector.

3. Experiment

In this section, we demonstrate the usefulness of the AOJ dataset in ML-related applications. In particular, we pre-train a foundation model on the full AOJ dataset.

3.1. Method

The foundation model chosen for this work is OMNIJET- α [11], a GPT-style [46] model that predicts the next token in a sequence. The features of the jet constituents¹⁰ are first converted into integer tokens using a VQ-VAE [10, 47–49] model, representing jets as sequences of tokens. The backbone of OMNIJET- α is trained to predict the next token in these sequences. After pre-training, the model can generate new sequences autoregressively, which are then decoded by the frozen VQ-VAE back into physical space. Additionally, the pre-trained backbone weights can be loaded and used for other tasks, such as jet classification. In this work, we focus on dataset transfer learning for jet generation and leave the exploration of task switching for future studies.

The features used for training are the constituent kinematics $(p_T, \Delta\eta, \Delta\phi)$, where $\Delta\eta = \eta - \eta_{\text{jet}}$ and $\Delta\phi = \phi - \phi_{\text{jet}}$ represent the relative coordinates with respect to the jet axis. Before tokenization, these features were preprocessed using standardization¹¹. The AOJ dataset was tokenized using a codebook size of 8192 tokens, employing VQ-VAE hyperparameters as outlined in [11], with two modifications: the number of attention heads was reduced from 8 to 4 without sacrificing performance, and the AdamW optimizer was replaced with the Ranger optimizer [50]. A larger codebook with ~ 32 K tokens was also tested but produced results comparable to the 8 K configuration. Therefore, we only present results for the 8 K tokenization. The backbone model was trained on the full 178 M AOJ dataset with hyperparameters as specified in [11]. The model was trained for 900 K steps on 12 GPUs with a batch size of 256 per GPU, and the final state at the end of the training was selected as the backbone model.

To evaluate the transfer learning capabilities of OMNIJET- α , we devised a downstream task to test its ability to generate jets of a different type than those it was pre-trained on, and to adapt to datasets with different kinematic cuts. The fine-tuning process involved loading the pre-trained backbone weights and retraining it using jets from the JetClass [23] dataset. To systematically evaluate the performance of the model, we trained it on six datasets with $10^2, 10^3, 10^4, 10^5, 10^6, 10^7$ jets to study the effect of dataset size D , and fine-tuned the model separately on QCD jets (q/g) and top-quark jets ($t \rightarrow bqq'$) to evaluate its adaptability to different jet substructure types. For each configuration, the model with the lowest validation loss during fine-tuning was selected for jet generation. To provide a baseline for comparison, we used a separate VQ-VAE tokenizer and trained generative models with the same backbone architecture from scratch using randomly initialized weights. This VQ-VAE is the one from [11] that was trained on all 10 classes of the JetClass dataset. The generative models were trained on the tokenized QCD and top-quark samples. This approach allowed us to quantify the performance improvement achieved by pre-training on the AOJ dataset.

¹⁰ The jet constituents are p_T ordered and up to the first 128 constituents are considered when training the models.

¹¹ Standardization is a common technique used to bring feature values into a manageable range for neural networks. By subtracting the mean and dividing by the standard deviation, the data is transformed to having a mean of zero and standard deviation of one.

Both training strategies used a batch size of 256¹² per GPU¹³ and a patience of 10 epochs. For training samples with sizes up to 100 K jets, the validation set was the same size as the training set. For larger samples, we used validation set sizes of 200 K jets for 1 M training samples and 500 K jets for 10 M training samples.

3.2. Performance metrics

We consider two complementary evaluation metrics that compare the 1D distribution of a high-level observable $\mathcal{O}$ between the generated jets and the target JetClass reference data:

- Kullback–Leibler divergence:

$$\text{KL}^{\mathcal{O}}(P||Q) = \sum_x P(x) \log \left(\frac{P(x)}{Q(x)} \right). \quad (1)$$

The probability densities $P(x)$ and $Q(x)$ for the observable $\mathcal{O}$ are obtained from the normalized histogram bin counts and $\text{KL}^{\mathcal{O}}(P||Q)$ is computed using `scipy.stats.entropy` [51]. Note that $\text{KL}^{\mathcal{O}}(P||Q)$ is anti-symmetric. In our results, we take P to be the generated distribution and Q to be the reference JetClass distribution.

- Wasserstein-1 distance:

$$W_1^{\mathcal{O}}(P, Q) = \min_{\pi \in \Pi} \sum_{x, y} |x - y| \pi(x, y). \quad (2)$$

Here $\Pi(P, Q)$ is the space of all joint ‘coupling’ distributions whose marginals are the observable $\mathcal{O}$ distributions P and Q . This quantity is computed with `scipy.stats.wasserstein_distance` [51]. Unlike the $\text{KL}^{\mathcal{O}}(P||Q)$ divergence, $W_1^{\mathcal{O}}(P, Q)$ is a symmetric similarity metric based on optimal transport, i.e. the optimal cost plan that morphs P into Q .

While further detailed performance evaluations such as classifier tests [52–54] could also be envisioned, this paper is first and foremost demonstrating a use case for the AOJ dataset in the context of foundation models. As such, we are rather concerned with the overall qualitative behavior, for which the 1D distributions suffice.

3.3. Results

A comparison of the two training strategies –‘fine-tuned’ and ‘from scratch’– is presented in figure 3 for QCD and in figure 4 for top-quark jets. The rows show various high-level distributions, including jet p_T , jet mass, subjeiness ratios τ_{21} and τ_{32} , and particle multiplicity N (i.e. the number of constituents in a jet). The left-most column in each figure displays the distributions generated by the fine-tuned models, while the second column shows the distributions generated by the models trained from scratch. Each plot also includes the reference JetClass distribution (200 K reference jets) as a filled histogram and, for additional context, the AOJ distributions (dotted lines) in the left-most column.

Overall, we find that the pre-trained models are able to achieve a performance gain when using a fraction of the total fine-tuning data. This can be seen in the third and last columns of figures 3 and 4, where the fine-tuned models (blue lines) typically achieve better¹⁴ metric scores than the models trained from scratch (red lines). This is apparent for the KL divergence and Wasserstein-1 metrics across different training sample sizes below 1 M jets. However, as the available training data exceeds 1 M jets, we find that the fine-tuning begins to saturate, and the fine-tuned and from-scratch models perform similarly within uncertainties. In some features, the fine-tuned models even exhibit a slight decrease in performance compared to those trained entirely from scratch, indicating a diminishing return from pre-training at larger dataset scales. The reasons for this are unclear and would require further study. One possibility is that the AOJ tokenization is suboptimal for describing JetClass and this is exposed with enough JetClass training data. Another possibility is that the pre-training itself can harm model performance by locking in the weights to a suboptimal part of parameter space, a phenomenon that has been observed previously in the ML literature [55] and was termed ‘ossification’.

¹² For training samples of 100 jets, the batch size was reduced to 100.

¹³ We used a single GPU to train the models with training sample sizes less than or equal to 100 K, and multiple GPUs for models trained on more data.

¹⁴ However, there are instances where models trained from scratch perform equally well on a single metric within uncertainties. For example, this occurs with KL^{p_T} for training sample size $D = 10^3$. This behavior can be attributed to the definition of KL divergence, which penalizes high-probability regions in the generated data where the true probability distribution is low. Therefore, it is crucial to evaluate performance using multiple metrics to gain a comprehensive understanding.




Figure 3. A comparison of the generation quality of models trained on QCD jets from JetClass, for different training sample sizes D . The performance of a foundation model pre-trained on AOJ (far left) is compared to a model trained from scratch (center left). The generation quality is compared across several high level features of the jets. Two quantitative metrics, the Kullback–Leibler divergence (center right) and Wasserstein-1 distance (far right), are computed as a function of the training sample size D to compare how well the generated jets from each model matches the target distribution from JetClass. We also report the mean value and the envelopes over 5 trainings with different random seeds. When applicable, we also show power-law fits $\propto D^{-\gamma}$ to the metrics.

We now examine in more detail the transfer learning from AOJ to JetClass for QCD jets, as shown in figure 3. The primary challenge in fine-tuning arises from the differences in the domains of the jet p_T distributions and particle multiplicities between the two datasets. JetClass imposes a narrower p_T window between 500 GeV and 1000 GeV and tends to have around ~ 30 particles per jet on average, while AOJ only applies a lower cut around 250 GeV and has around ~ 50 particles per jet on average. When fine-tuning, the pre-trained model must adapt to these differences and learn to generate jets with fewer particles and the new kinematic cuts. The results for the p_T distributions and particle multiplicities show how the fine-tuned model smoothly interpolates between the AOJ and JetClass as a function of the training sample size D . Furthermore, the generation quality improves with increasing training sample size, achieving almost optimal results with a reduced sample size of 100 K jets for the p_T distribution and 10 K for the multiplicities. In contrast, the randomly initialized models trained from scratch require a significantly larger dataset, approximately an order of magnitude more jets, to achieve optimal performance. For the jet mass and the substructure observable τ_{21} , the fine-tuning performance saturates with a remarkably small number of training samples, resulting in an almost flat dependence of the metrics on D . Specifically, fine-tuning requires only about 1000 jets to reach optimal performance, whereas training from scratch requires significantly more data to achieve comparable results. This is not too surprising, as the transfer from AOJ to JetClass for these features does not require much effort, given the similarity in the shapes of these distributions.

Next, we examine the results in figure 4, which correspond to top-quark jet generation. Since the AOJ dataset is completely dominated by QCD jets, the transfer learning task between AOJ and JetClass becomes more challenging due to the change in jet type. For instance, fine-tuning must map the low-mass Sudakov




Figure 4. A comparison of the generation quality of models trained on top-quark jets from JetClass, for different training sample sizes D . The performance of a foundation model pre-trained on AJO (far left) is compared to a model trained from scratch (center left). The generation quality is compared across several high level features of the jets. Two quantitative metrics, the Kullback–Leibler divergence (center right) and Wasserstein-1 distance (far right), are computed as a function of the training sample size D to compare how well the generated jets from each model matches the target distribution from JetClass. We also report the mean value and the envelopes over 5 trainings with different random seeds. When applicable, we also show power-law fits $\propto D^{-\gamma}$ to the metrics.

peak characteristic of QCD jets to the top-quark resonance peak at $m_t = 175$ GeV. Additionally, the model must transition from generating 1-prong QCD jets to 3-pronged top-quark jets, as reflected in the horizontal gap between the peaks in the τ_{32} substructure distributions. Despite these obstacles, we find that transfer learning consistently outperforms training from scratch when using smaller datasets. For example, the top-quark mass, represented by the mode in the jet mass distribution, is learned with good precision by the fine-tuned model using a training dataset of only 10 K jets. In contrast, the model trained from scratch requires an order of magnitude more data to achieve comparable performance. A similar trend is observed for the τ_{32} distribution, where transfer learning morphs the AJO jets into JetClass top-quark jets, achieving competitive results with only 100 K training jets. In comparison, models trained from scratch fail to generate jets with a reliable substructure for training sample sizes below 1 M jets.

3.4. Power-law scaling

Finally, we observe that the performance metrics for the fine-tuned models, when not in the saturated regime, seem to follow a power-law scaling of the form $D^{-\gamma}$ with respect to the training sample size D . Neural scaling laws, such as these, have attracted broad interest and have been analyzed extensively in various settings in recent years, see [56] for a nice overview of the ML literature and original references. In particular, they have been observed in the context of jet classification [57]. We find that the exponents for a wide range of high-level observables are consistently within the range $0.2 < \gamma < 0.45$ (see dashed lines in figures 3 and 4). This scaling behavior holds for both the KL and Wasserstein-1 metrics and applies to both QCD and top-quark, suggesting a form of *universality*. The only exception to this scaling is the particle multiplicity distribution, which shows a more irregular dependence on the training sample size. Previous scaling law

behavior has been observed as a function of the loss of the model, not on such hand-crafted metrics. However, due to the different tokenizers used in the from-scratch and the AOJ fine-tuned model, the losses for the generation objective are not directly comparable, making statements of true scaling comparisons difficult. We leave further investigation of this apparent scaling behavior to future work.

In contrast, the randomly initialized models trained from scratch are not seen to follow such universal scaling across most physical observables. The reason behind this may be that these operate in the small-data regime [58], where their performance is limited by the lack of sufficient training data. In this regime, models can only perform marginally better than random guessing, as each additional sample provides limited new information for learning. The breaking of scaling is particularly evident for the jet substructure features, such as τ_{21} for QCD and τ_{32} for tops, and to a lesser extent for the jet mass, where the performance metrics exhibit irregular and dataset-dependent scaling as the training sample size increases. Only after the model has seen sufficient training data, approximately around 1 M jets, does the training stabilize, resulting in reliable jet substructure and jet mass modeling. While further studies into these scaling laws need to be performed, the observed discrepancy highlights the possible advantage of pre-training in achieving a more robust and systematic improvement in generative performance.

4. Conclusion

We are releasing the AOJs dataset, which comprises 178 million jets from CMS open data, in a standard machine learning-friendly format. This dataset provides an excellent resource for exploring foundation models and other machine learning techniques that use large, unlabeled datasets for pre-training.

In addition, we present the first jet-based foundation model in high-energy physics trained on real collider data from the AOJ dataset. By using this rich dataset to pre-train a foundation model, we show that fine-tuning for specific tasks—such as generating jets of different types in different kinematic regions—yields significant performance improvements over models trained from scratch, and boosts the achievable downstream performance on small labeled datasets. In particular, the generative performance of our foundation model scales predictably with training sample size, facilitating effective transfer learning. This demonstrates the power of foundation models to provide improved performance even over significant domain shifts.

We hope that this study will encourage further engagement with open data from LHC experiments, and promote advances that bridge phenomenological studies and deployment in actual experiments.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://doi.org/10.25592/uhhfdm.16505>.

Acknowledgments

This work was initiated at the Aspen Center for Physics, supported by National Science Foundation Grant PHY-2210452. We would like to thank Alexander Mück for discussions. The research of MK and HR-G is supported by the Deutsche Forschungsgemeinschaft DFG under Grant 396021762 – TRR 257: Particle physics phenomenology after the Higgs discovery. JB, AH and GK are supported by the DFG under the German Excellence Initiative—EXC 2121 Quantum Universe – 390833306, and by PUNCH4NFDI—project number 460248186. DAF, IP, and DS are supported by DOE Grant DOE-SC0010008. OA is supported by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. LA is supported by the University of Bologna. Additionally, we acknowledge support from the Maxwell computational resources at Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany, and computing resources provided by RWTH Aachen University under project rwth0934. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC award HEP-ERCAP0027491.

Code

The code that was used to create ASPENOPENJETS from CMS open data can be found at <https://github.com/OzAmram/AOJProcessing>, and the code used for the OMNIJET- α model and its training can be found at https://github.com/uhh-pd-ml/omnijet_alpha.

Appendix. CMS samples

MC samples provided by CMS were used to estimate the composition of the AOJ sample. CMS uses a variety of generators for the simulation different physics processes, including MADGRAPH5_AMC@NLO [43], PYTHIA [44], and POWHEG [59–61]. Each sample uses PYTHIA with the underlying event tune CP5 [62] to simulate the parton shower and hadronization. All samples use the NNLO NNPDF 3.1 parton distribution functions [63–65]. Some physics processes are split into several samples covering different kinematic regions. Each sample is normalized to the appropriate cross section.

To estimate the number of fully merged W , Z , top, and $H \rightarrow b\bar{b}$ jets the corresponding MC samples were used. A matching criteria was applied based on truth-level information from the generator to ensure the estimate reflected the number of fully-merged jets of each type. For each heavy resonance, generator-level information was used to check that all the quarks from the heavy resonance decay were within $\Delta R < 0.8$ of the reconstructed AK8 jet found by our selection. The remaining 99.5% of jets not estimated to be heavy resonance decays are assumed to be QCD. The resulting numbers should be considered rough estimates, as the MC modeling of these signals is imperfect, and the estimates do not incorporate corrections to the simulated trigger and reconstruction efficiency typically employed in physics analyses by CMS.

Table 2 lists the CMS data and MC samples used and their corresponding cross sections.

Table 2. A list of the datasets used to construct AOJ and the Monte Carlo samples used to estimate its composition.

Process	CMS dataset name	Cross section (pb)
Data	/JetHT/Run2016G-UL2016_MiniAODv2-v2/MINIAOD	—
	/JetHT/Run2016H-UL2016_MiniAODv2-v2/MINIAOD	—
QCD	/QCD_Pt_300to470_TuneCP5_13TeV_pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v1/MINIAODSIM	7823
	/QCD_Pt_470to600_TuneCP5_13TeV_pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v1/MINIAODSIM	648.2
	/QCD_Pt_600to800_TuneCP5_13TeV_pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v1/MINIAODSIM	186.9
	/QCD_Pt_800to1000_TuneCP5_13TeV_pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v1/MINIAODSIM	32.3
$t\bar{t}$	/TTToHadronic_TuneCP5_13TeV-powheg-pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v1/MINIAODSIM	378.5
	/TTToSemiLeptonic_TuneCP5_13TeV-powheg-pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v1/MINIAODSIM	365.2
$W + \text{jets}$	/WJetsToQQ_HT-400to600_TuneCP5_13TeV-madgraphMLM-pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v2/MINIAODSIM	315.6
	/WJetsToQQ_HT-600to800_TuneCP5_13TeV-madgraphMLM-pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v2/MINIAODSIM	68.6
	/WJetsToQQ_HT-800toInf_TuneCP5_13TeV-madgraphMLM-pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v2/MINIAODSIM	34.9
$Z + \text{jets}$	/ZJetsToQQ_HT-400to600_TuneCP5_13TeV-madgraphMLM-pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v2/MINIAODSIM	145.4
	/ZJetsToQQ_HT-600to800_TuneCP5_13TeV-madgraphMLM-pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v2/MINIAODSIM	34.0
	/ZJetsToQQ_HT-800toInf_TuneCP5_13TeV-madgraphMLM-pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v2/MINIAODSIM	18.7
$H \rightarrow b\bar{b}$	/GluGluHTtoBB_Pt-200toInf_M-125_TuneCP5_MINLO_13TeV-powheg-pythia8/RunIISummer20UL16MiniAODv2-106X_mcRun2_asymptotic_v17-v2/MINIAODSIM	0.27

References

- [1] Bommasani R et al 2021 (available at: <https://crfm.stanford.edu/assets/report.pdf>)
- [2] Brown T et al 2020 *Advances in Neural Information Processing Systems* vol 33 (Curran Associates, Inc.) pp 1877–901
- [3] Devlin J, Chang M-W, Lee K and Toutanova K 2019, arXiv:1810.04805 [cs.CL]
- [4] Touvron H et al 2023, arXiv:2302.13971 [cs.CL]
- [5] Saharia C et al (Curran Associates Inc.) 2022 arXiv:2205.11487
- [6] Radford A et al 2021 arXiv:2103.00020 [cs.CV]
- [7] Jia C, Yang Y, Xia Y, Chen Y, Parekh Z, Pham H, Le Q V, Sung Y, Li Z and Duerig T 2021, arXiv:2102.05918
- [8] Kishimoto T, Morinaga M, Saito M and Tanaka J 2023 *37th Conf. on Neural Information Processing Systems*
- [9] Qu H, Li C and Qian S 2022, arXiv:2202.03772 [hep-ph]
- [10] Heinrich L, Golling T, Kagan M, Klein S, Leigh M, Osadchy M and Raine J A 2024 arXiv:2401.13537 [hep-ph]
- [11] Birk J, Hallin A and Kasieczka G 2024 *Mach. Learn. Sci. Tech.* **5** 035031
- [12] Harris P, Kagan M, Krupa J, Maier B and Woodward N 2024 arXiv:2403.07066 [hep-ph]
- [13] Mikuni V and Nachman B 2024 arXiv:2404.16091 [hep-ph]
- [14] CMS Collaboration 2022 Multijet primary dataset in AOD format from RunA of 2011 (/MultiJet/Run2012A-22Jan2013-v1/AOD) (available at: <http://dx.doi.org/10.7483/OPENDATA.CMS.WING.7QKV>)
- [15] CMS Collaboration 2022 Multijet primary dataset in AOD format from RunA of 2012 (/MultiJet/Run2012A-22Jan2013-v1/AOD) (available at: <http://dx.doi.org/10.7483/OPENDATA.CMS.WING.7QKV>)
- [16] CMS Collaboration 2024 JetHT primary dataset in MINIAOD format from RunG of 2016 (/JetHT/Run2016G-UL2016_MiniAODv2-v2/MINIAOD) (available at: <http://dx.doi.org/10.7483/OPENDATA.CMS.1KTG.X0W4>)
- [17] CMS Collaboration 2024 JetHT primary dataset in MINIAOD format from RunH of 2016 (/JetHT/Run2016H-UL2016_MiniAODv2-v2/MINIAOD) (available at: <http://dx.doi.org/10.7483/OPENDATA.CMS.LT9E.T7RQ>)
- [18] ATLAS Collaboration 2024 DAOD_PHYSLITE format 2015-2016 Open Data for Research from the ATLAS experiment (available at: <http://dx.doi.org/10.7483/OPENDATA.ATLAS.9HK7.P5SI>)
- [19] Larkoski A, Marzani S, Thaler J, Tripathi A and Xue W 2017 *Phys. Rev. Lett.* **119** 132003
- [20] Tripathi A, Xue W, Larkoski A, Marzani S and Thaler J 2017 *Phys. Rev. D* **96** 074003
- [21] Komiske P T, Mastandrea R, Metodiev E M, Naik P and Thaler J 2020 *Phys. Rev. D* **101** 034009
- [22] Dolan M J, Gargalionis J and Ore A 2023 arXiv:2312.03434 [hep-ph]
- [23] Qu H, Li C and Qian S 2022 JetClass: a large-scale dataset for deep learning in jet physics (available at: <http://dx.doi.org/10.5281/zenodo.6619768>)
- [24] Chatrchyan S et al (CMS) 2008 *JINST* **3** S08004
- [25] CMS Collaboration 2014 Jet primary dataset in AOD format from RunB of 2010 (/Jet/Run2010B-Apr21Reco-v1/AOD) (available at: <http://dx.doi.org/10.7483/OPENDATA.CMS.3S7F.2E9W>)
- [26] Petrucci G, Rizzi A and Vuosalo C (CMS) 2015 *J. Phys.: Conf. Ser.* **664** 7
- [27] Peruzzi M, Petrucci G, Rizzi A and (the CMS Collaboration) 2020 *J. Phys.: Conf. Ser.* **1525** 012038
- [28] Sirunyan A M et al (CMS) 2020 *JINST* **15** 10017
- [29] Khachatryan V et al (CMS) 2017 *JINST* **12** 01020
- [30] Sirunyan A M et al (CMS) 2017 *JINST* **12** 10003
- [31] Cacciari M, Salam G P and Soyez G 2008 *J. High Energy Phys.* **JHEP04(2008)063**
- [32] Cacciari M, Salam G P and Soyez G 2012 *Eur. Phys. J. C* **72** 1896
- [33] Sirunyan A M et al (CMS) 2020 *JINST* **15** 09018
- [34] Bertolini D, Harris P, Low M and Tran N 2014 *J. High Energy Phys.* **JHEP10(2014)059**
- [35] Khachatryan V et al (CMS) 2017 *JINST* **12** 02014
- [36] CMS Collaboration 2024 PFNano producer tool for CMS 2016 (available at: <https://opendata.cern.ch/record/12504>)
- [37] CMS Collaboration, 2024 Monte Carlo simulations of 2016 data (available at: <https://opendata.cern.ch/search?q=&f=experiment%3ACMS&f=type%3ADataset%2Bsubtype%3ASimulated&f=year%3A2016>)
- [38] Larkoski A J, Marzani S, Soyez G and Thaler J 2014 *J. High Energy Phys.* **JHEP05(2014)146**
- [39] Navas S et al (Particle Data Group) 2024 *Phys. Rev. D* **110** 030001
- [40] Thaler J and Van Tilburg K 2011 *J. High Energy Phys.* **JHEP03(2011)015**
- [41] Qu H and Gouskos L 2020 *Phys. Rev. D* **101** 056019
- [42] CMS 2020 (available at: <https://cds.cern.ch/record/2707946>)
- [43] Allwall J, Frederix R, Frixione S, Hirschi V, Maltoni F, Mattelaer O, Shao H S, Stelzer T, Torrielli P and Zaro M 2014 *J. High Energy Phys.* **JHEP07(2014)079**
- [44] Sjöstrand T, Ask S, Christiansen J R, Corke R, Desai N, Ilten P, Mrenna S, Prestel S, Rasmussen C O and Skands P Z 2015 *Comput. Phys. Commun.* **191** 159
- [45] de Favereau J, Delaere C, Demin P, Giammanco A, Lemaitre V, Mertens A and Selvaggi M (DELPHES 3) 2014 *J. High Energy Phys.* **JHEP02(2014)057**
- [46] Radford A, Narasimhan K, Salimans T and Sutskever I 2018 (available at: https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf)
- [47] van den Oord A, Vinyals O and Kavukcuoglu K 2018 arXiv:1711.00937 [cs.LG]
- [48] Bao H, Dong L, Piao S and Wei F 2022 arXiv:2106.08254 [cs.CV]
- [49] Huh M, Cheung B, Agrawal P and Isola P 2023 arXiv:2305.08842 [cs.LG]
- [50] Wright L 2019 GitHub repository
- [51] Virtanen P et al 2020 *Nat. Methods* **17** 261
- [52] Krause C and Shih D 2023 *Phys. Rev. D* **107** 113003
- [53] Kansal R, Li A, Duarte J, Chernyavskaya N, Pierini M, Orzari B and Tomei T 2023 *Phys. Rev. D* **107** 076017
- [54] Das R, Favaro L, Heimerl T, Krause C, Plehn T and Shih D 2024 *SciPost Phys.* **16** 031
- [55] Hernandez D, Kaplan J, Henighan T and McCandlish S 2021 arXiv:2102.01293
- [56] Villalobos P 2023 (available at: <https://epoch.ai/blog/scaling-laws-literature-review>) (Accessed: 13 December 2024)
- [57] Batson J and Kahn Y 2023 arXiv:2312.02264 [hep-ph]
- [58] Hestness J, Narang S, Ardalani N, Diamos G, Jun H, Kianinejad H, Patwary M M A, Yang Y and Zhou Y 2017 arXiv:1712.00409

- [59] Nason P 2004 *J. High Energy Phys.* [JHEP 0411:040](#)
- [60] Frixione S, Nason P and Ridolfi G 2007 *J. High Energy Phys.* [JHEP09\(2007\)126](#)
- [61] Alioli S, Nason P, Oleari C and Re E 2010 *J. High Energy Phys.* [JHEP06\(2010\)043](#)
- [62] Sirunyan A M *et al* (CMS) 2020 *Eur. Phys. J. C* [80 4](#)
- [63] Ball R D, Del Debbio L, Forte S, Guffanti A, Latorre J I, Rojo J and Ubiali M 2010 *Nucl. Phys. B* [838 136](#)
- [64] Ball R D *et al* (NNPDF) 2015 *J. High Energy Phys.* [JHEP04\(2015\)040](#)
- [65] Ball R D *et al* (NNPDF) 2017 *Eur. Phys. J. C* [77 663](#)