




Performance of a geometric deep learning pipeline for HL-LHC particle tracking

Xiangyang Ju^{1,a} , Daniel Murnane¹, Paolo Calafiura¹, Nicholas Choma¹, Sean Conlon¹, Steven Farrell¹, Yaoyuan Xu¹, Maria Spiropulu², Jean-Roch Vlimant², Adam Aurisano³, Jeremy Hewes³, Giuseppe Cerati⁴, Lindsey Gray⁴, Thomas Klijsma⁴, Jim Kowalkowski⁴, Markus Atkinson⁵, Mark Neubauer⁵, Gage DeZoort⁶, Savannah Thais⁶, Aditi Chauhan⁷, Alex Schuy⁷, Shih-Chieh Hsu⁷, Alex Ballow⁸, Alina Lazar⁸

¹ Lawrence Berkeley National Laboratory, Berkeley, CA, USA

² California Institute of Technology, Pasadena, CA, USA

³ University of Cincinnati, Cincinnati, OH, USA

⁴ Fermi National Accelerator Laboratory, Batavia, IL, USA

⁵ University of Illinois at Urbana-Champaign, Urbana, IL, USA

⁶ Princeton University, Princeton, NJ, USA

⁷ University of Washington, Seattle, WA, USA

⁸ Youngstown State University, Youngstown, OH, USA

Received: 3 May 2021 / Accepted: 21 September 2021

© The Author(s) 2021

Abstract The Exa.TrkX project has applied geometric learning concepts such as metric learning and graph neural networks to HEP particle tracking. Exa.TrkX's tracking pipeline groups detector measurements to form track candidates and filters them. The pipeline, originally developed using the TrackML dataset (a simulation of an LHC-inspired tracking detector), has been demonstrated on other detectors, including DUNE Liquid Argon TPC and CMS High-Granularity Calorimeter. This paper documents new developments needed to study the physics and computing performance of the Exa.TrkX pipeline on the full TrackML dataset, a first step towards validating the pipeline using ATLAS and CMS data. The pipeline achieves tracking efficiency and purity similar to production tracking algorithms. Crucially for future HEP applications, the pipeline benefits significantly from GPU acceleration, and its computational requirements scale close to linearly with the number of particles in the event.

1 Introduction

Charged particle tracking plays an essential role in High-Energy Physics (HEP), including particle identification and kinematics, vertex finding, lepton reconstruction, and flavor jet tagging. At the core of particle tracking there is a pattern recognition algorithm that must associate a list of 2D or 3D

position measurements from a tracking detector (known as *hits* or *spacepoints* in literature) to a list of particle track candidates (or *tracks*. A *track* is defined as a list of spacepoints associated by the pattern recognition to a charged particle).

The number of particle track candidates varies significantly from one experiment setup to another. For example, in a High-Luminosity LHC (HL-LHC) [1] collision *event*, due to the *pile-up* of multiple proton–proton collision per bunch crossing, there are typically 5000 charged particles and 100,000 spacepoints, about 50% of which are associated to particles of interest.

A typical HEP offline tracking algorithm [3–5] has four stages: spacepoint formation, track seeding, track following, and track fitting. The spacepoint formation stage combines the detector readout cell raw data in clusters from which the spacepoint 3D coordinates, and their uncertainties, are determined. Track seeding combines spacepoints in *doublet* or *triplet seeds*. Each seed provides an initial track direction, origin, and possibly a curvature, with associated uncertainties. The track following stage adds more spacepoints to the seed by looking for matching spacepoints along the extrapolated trajectory. Finally a track fitting stage, which may be combined with the track following, fits a trajectory through the track spacepoints to assess the track quality and measure the particle's physical and kinematic properties (charge, momentum, origin, etc). To avoid biasing physics results, each stage of the algorithm must have high *efficiency*, meaning it must identify e.g. > 90% of the charged particles within a fiducial region (e.g. $p_T > 1$ GeV, $|\eta| < 4$) as track candi-

^a e-mail: xju@lbl.gov (corresponding author)

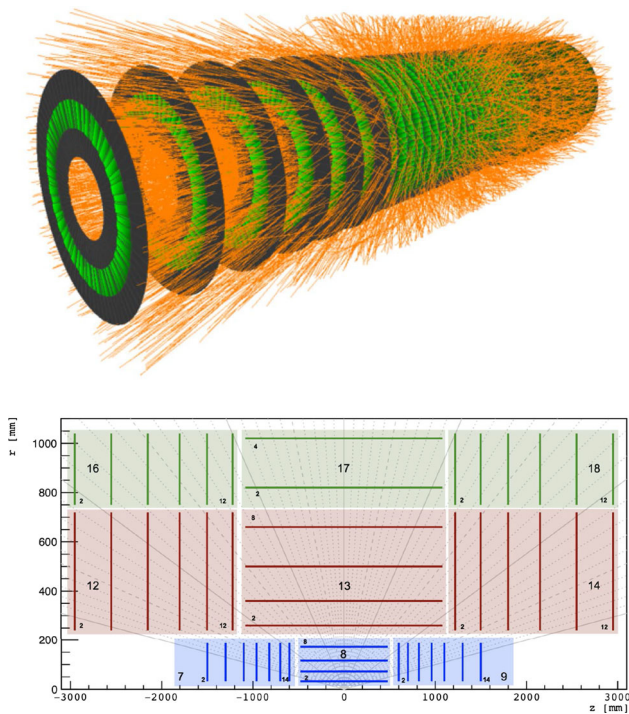


Fig. 1 A simulated HL-LHC collision event (top) as seen by the TrackML tracking detector [2]. The detector schematic (bottom) shows the top half of the detector projected on the r-z plane. The z-axis is along the beam direction

dates. Track seeding and track filtering must also have high *purity*, meaning that e.g. > 60% of the track seeds and track candidates must correspond to charged particles. High purity allows to keep the number of track candidates, and the associated computational costs, under control.

Online tracking algorithms may use different pattern recognition algorithms¹ to create and filter track seeds and candidates, but share the same high efficiency requirements. Online application also have stringent computing requirements (e.g. latency $O(10)$ μ s for LHC triggers).

The computational cost of current tracking algorithms grows worse than linearly with beam intensity and detector occupancy, as demonstrated in Fig. 2. Given the order-of-magnitude increase for beam intensity at HL-LHC, charged particle pattern recognition algorithms might well limit the discovery potential of HL-LHC experiments.

Over the last two decades, tracking computational challenges arising from the increased number of combinations have been addressed by tightening fiducial regions for charged particles, developing highly optimized tracking algorithms [4,5], and even optimizing the geometry of tracking detectors. These optimizations brought order-of-magnitude gains in tracking computational performance with limited impact on physics. While these efforts continue [12],

¹ Including Hough transforms [6,7] and cellular automata [8,9].

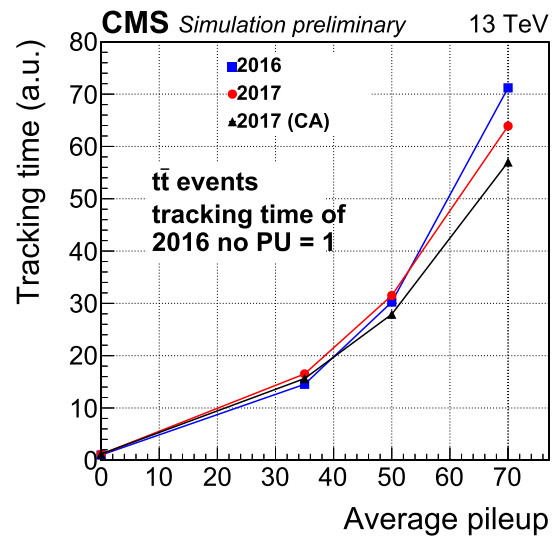
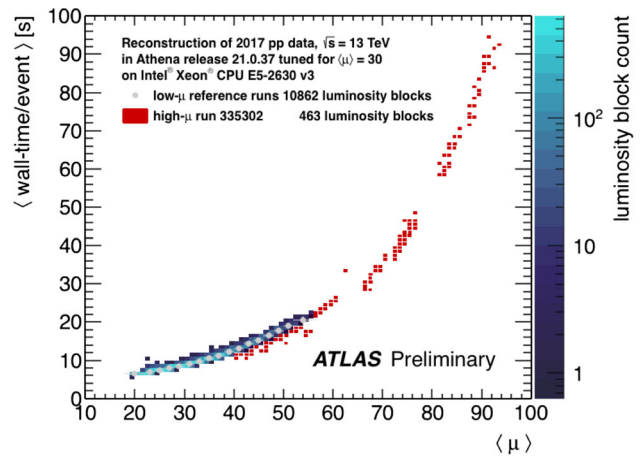


Fig. 2 Reconstruction wall time per event as a function of the average number of interactions per bunch crossing $\langle \mu \rangle$. Top: ATLAS Run 2 Inner Detector reconstruction with default configurations [10]. Bottom: CMS time spent in tracking sequence for 2016 tracking, 2017 tracking with conventional seeding, and 2017 tracking with Cellular Automaton (CA) seeding [11]

it is unlikely that another order of magnitude can be gained through incremental optimization without impacting physics performance. Furthermore, given the computational complexity and iterative nature of current track following and filtering algorithms, it is challenging to run them efficiently on data parallel architectures like GPUs.

The TrackML challenge [2] jump-started the application of deep learning pattern recognition methods applied to HEP tracking. The HEP.TrkX pilot project [13] proposed the use of graph networks to filter track doublet and triplet seeds [14]. Building on that work, the Exa.TrkX project [15] has demonstrated the applicability of *Geometric Deep Learning* (GDL) methods [16] – specifically metric learning and Graph Neural Networks (GNN) – to particle tracking [17]. GDL is concerned with learning representations of data that have com-

plex geometrical relationships and no natural ordering, like detector spacepoints. GDL models are computationally regular, naturally parallel and therefore well-suited to run on hardware accelerators.

This work describes new developments that enabled the first study of the computing and physics performance of the Exa.TrkX pipeline on the entire TrackML detector at HL-LHC design luminosity, a step towards the validation of the pipeline on ATLAS and CMS data.

2 Related work

Early on, the Hep.TrkX pilot project attempted to assign and regress track parameters to single spacepoints using image processing models. Subsequent attempts at estimating track parameters using image processing and recurrent networks showed promising results [18] in a simplified environment. A similar realization of the method is reported in [19] where a model processing image from successive pixel detector layers is used to produce tracklets, seeds to classical pattern recognition. The method yields superior seeding efficiency for tracks within jets in dense environments. The concept of using LSTM [20] to supplement the Kalman Filter method for track following developed by HEP.TrkX [14, 18, 21] was later found in one of the promising solutions of the accuracy phase [22] of the TrackML challenge. The task of particle tracking was addressed with a hit-to-track assignment method using gated recurrent unit [23] (GRU), producing promising result in sparse environments [21]. This approach was constrained computationally due to the use of recurrent models.

Reference [24] applies the track finding approach developed in Ref. [25] to the whole detector by exploiting a new data-driven graph construction method and large model support in Tensorflow [26]. Reference [27] applies a similar GNN model to the task of particle-flow reconstruction. The model has a classification objective, followed by a partial regression of generator-level particle candidate kinematics. The method performs at least as well as a classical particle-flow algorithm in HL-LHC-like collision conditions. As part of the Exa.TrkX project, graph networks are used for LArTPC track reconstruction [28]. Reference [29] explores the opportunity to implement Exa.TrkX-inspired graph networks on FPGAs. Starting from the input stage of the Exa.TrkX pipeline, Ref. [30] studies the impact of cluster shape information on track seeding performance. In Ref. [31], metric learning is used to improve the purity in spacepoints buckets formed using similarity hashing. With the advent of quantum computer of increasing size came the development of quantum machine learning techniques, also applied in particle physics [32]. In particular, inspired by the use of GNN for charged particle tracking of the Exa.TrkX

team, quantum graph networks have been tested on the same problem [33].

3 Methodology

3.1 Input data

This study is based on the TrackML dataset that uses a Monte Carlo simulation of top quark pair production from proton–proton collisions at the HL-LHC. To simulate the effect of event pileup and produce realistic detector occupancy, a Poisson random number (with $\mu = 200$) of QCD “minimum bias” events are overlaid on top of the $t\bar{t}$ collisions.

The TrackML detector is a set of concentric cylindrical layers of pixelated sensors (the *barrel*) complemented by a set of circular disks (the *endcaps*) to ensure nearly 4π coverage in solid angle, as pictured in Fig. 1. Figure 3 shows the spatial distribution of the spacepoints of a typical event. One notable feature of this dataset is the inclusion of “noise” spacepoints, added as a proxy for various low-momentum particle interactions and detector effects which would otherwise require more expensive and detailed simulations.

3.2 The Geometric Deep Learning Pipeline

This paper updates the methodology previously presented in Ref. [17] to a fully-learned pipeline, where both graph construction and graph classification are trained. This section describes the pipeline (represented schematically in Fig. 4) used to obtain the results in Sect. 4. Details of the latest model design, parameter choices, and technical optimizations are discussed in Sect. 5.

The pipeline currently used to reconstruct tracks from a pointcloud of spacepoints requires six discrete stages of processing and inference. These broadly consist of a preprocessing stage, three stages required to construct a spacepoint graph, and two stages required to classify the graph edges and partition them into track candidates. Each stage is trained independently (due to memory constraints) on the output of the previous stage’s inference.

First, the dataset is processed into a format suitable for model training. This includes calculating directional information and summary statistics from the charge deposited in each spacepoint, i.e. the *cell features* in Fig. 4. These values are appended to the cylindrical coordinates of each spacepoint to form an input feature vector to the pipeline. To apply a graph neural network to this set of data, it is necessary to arrange them into a graph. One can apply various geometric heuristics to define which spacepoints are likely to be connected by an edge (i.e. belong to the same track), but a useful technique is to train a model on the geometry of connected tracks. Thus, our second stage is to train an Embedding Net-

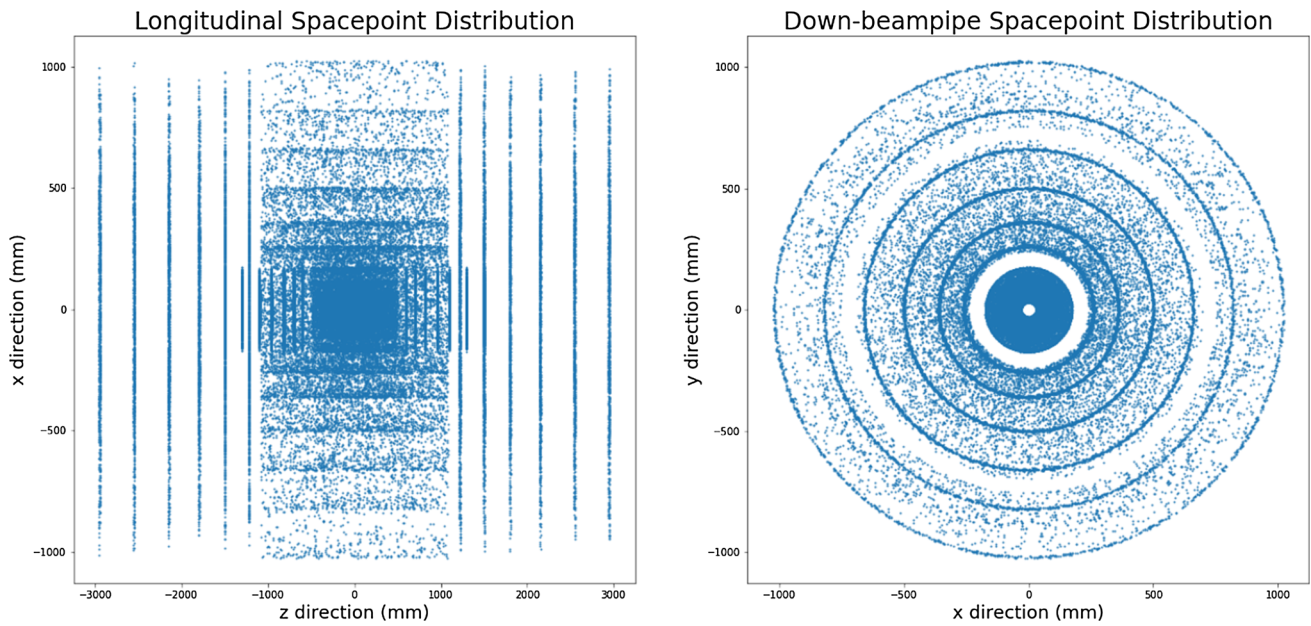


Fig. 3 A typical event distribution of spacepoints projected on the x - z plane, parallel to the beam direction (left), and the x - y plane, orthogonal to the beam direction (right)

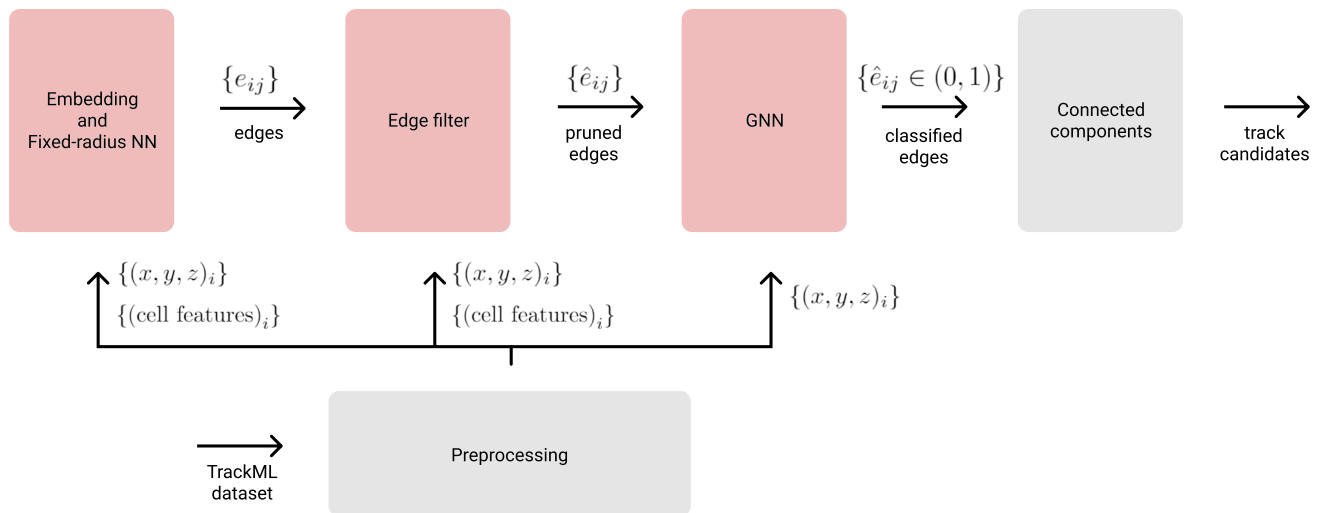


Fig. 4 Stages of the TrackML track formation inference pipeline. Light red boxes are trainable stages

work – a multi-layer perceptron (MLP) which embeds each spacepoint into an N -dimensional latent space. The graph is constructed by connecting neighboring spacepoints within a radius $r_{\text{embedding}}$, in the latent space. We train this embedding with a pairwise hinge loss, to encourage spacepoints that belong to the same track to be close in the embedded space, according to the Euclidean metric. This allows for a highly efficient edge construction, since we do not rely on any heuristics of the detector geometry that may lead to missed edges.

The edge selection at this stage is close to 100% efficient but $O(1)\%$ pure, with a graph size of $O(10^5)$ nodes

and $O(10^7)$ edges (the purity-efficiency trade-off can be tuned with the choice of $r_{\text{embedding}}$). Before running training or inference on the memory-intensive GNN, we filter these edges down with another MLP. The input to this third stage is the concatenated features on either side of each edge. That is, the Filter Network is a binary classifier applied to the set of edges. Constraining edge efficiency to remain high (above 96%) leads to much sparser graphs, of $O(10^6)$ edges.

The fourth stage of the pipeline is the training and inference of the graph neural network. The results presented in this work are predominantly obtained from the Interaction Network architecture, first proposed in Ref. [34]. This vari-

etal of GNN includes hidden features on both nodes and edges, which are propagated around the graph (called “message passing”) with consecutive concatenations along edges and aggregations of messages at receiving nodes. In the final layer of the network, a binary classification is obtained for each edge as true or fake, and trained on a cross-entropy loss.

The final stage of the TrackML pipeline involves task-specific post-processing. If our goal is track formation, we can place a threshold on the edge scores produced by the GNN and partition the graph into connected components. If our goal is track seeding, we can directly sample the classified edges for high likelihood combinations of connected triplets, or convert the entire graph to a *triplet graph* and train this on a second GNN to classify the triplets. A triplet graph is formed by taking all edges in the original (*doublet*) graph and assigning them as nodes in the new triplet graph. The nodes in this triplet graph are connected if they share a hit in the doublet graph. Applying a GNN to this structure produces highly pure sets of seeds as shown in Ref. [17].

Many of these techniques are common to other applications being explored in the Exa.TrkX collaboration. The pattern of nearest-neighbor graph-building and GNN edge classification has shown its potential for neutrino experiments [28] and CMS High Granularity Calorimeter [25]. Indeed, these applications build on the TrackML pipeline and extend it, for example by adding the particle type as an edge feature.

4 Results

4.1 Tracking performance of the TrackML pipeline

4.1.1 Tracking efficiency and purity

The performance of a tracking pipeline is mainly characterized by tracking efficiency and purity. For efficiency calculations, only charged particles that satisfy $|\eta| < 4.0$ and $p_T > 100$ MeV are considered. These *selected* particles, $N_{particles}(selected)$, are hereafter referred to as *particles*.

The overall tracking efficiency, known as *physics efficiency* ϵ_{phys} (Eq. 1), is defined as the fraction of particles that are *matched* to at least one reconstructed track. A particle is considered to be matched to a reconstructed track when (1) the majority of spacepoints in the reconstructed track belong to the same true track, and (2) the majority of spacepoints in the matched true particle track are found in the reconstructed track.²

To measure the efficiency of the tracking pipeline itself, we also define the *technical efficiency* ϵ_{tech} (Eq. 2) as the

fraction of *reconstructable* particles matching at least one reconstructed track. Reconstructable particles have a trajectory that leaves at least five spacepoints in the detector. Tracking purity (Eq. 3) is defined as the fraction of reconstructed tracks that match a selected particle.³

$$\epsilon_{phys} = \frac{N_{particles}(selected, matched)}{N_{particles}(selected)} \quad (1)$$

$$\epsilon_{tech} = \frac{N_{particles}(selected, reconstructable, matched)}{N_{particles}(selected, reconstructable)} \quad (2)$$

$$Purity = \frac{N_{tracks}(selected, matched)}{N_{tracks}(selected)} \quad (3)$$

Averaged over 50 testing events from the TrackML dataset, the physics efficiency for particles with $p_T > 500$ MeV is $88.7 \pm 0.3\%$ and the technical efficiency is $97.6 \pm 0.3\%$. Without any fiducial p_T cut, the physics efficiency becomes $67.2 \pm 0.1\%$ and the technical efficiency $91.3 \pm 0.2\%$. The tracking purity is $58.3 \pm 0.6\%$. Using the TrackML challenge scoring system and all tracks in the event, we obtained a score of 0.877 ± 0.005 .⁴ The errors quoted are statistical only.

Figure 5 shows the p_T distribution of particles as well as the tracking efficiency as a function of particle p_T . The physics efficiency for particles with p_T of [100, 300] MeV is 43%, therefore, is not displayed in the plot. The physics efficiency for particles with $p_T > 700$ MeV is above 88%. The technical efficiency is 82% for particles with p_T of [100, 300] MeV, and increases to above 97% for particles with $p_T > 700$ MeV. Figure 5 also shows the η distribution of particles with $p_T > 500$ MeV as well as the tracking efficiency as a function of the particle η . The physics efficiency is higher in the *barrel* region of the detector (volumes 8, 13, 17 in Fig. 1), while the technical efficiency is almost flat across the η range. In Fig. 5 the p_T and η of the matched truth particle were used, rather than the p_T and η of the reconstructed track. We leave a study of track quality and detector resolution effects for future work.

4.1.2 Systematic studies

Before using a tracking algorithm in production, it is necessary to measure its sensitivity to systematic effects, including pile-up, noise and digitization errors, and uncertainties in the measurement of detector properties (alignment, rotation, magnetic field map, etc.).

³ HEP tracking literature often quotes fake rate = 1 – purity.

⁴ We obtained a score of 0.914 ± 0.006 by training the pipeline with a dataset that includes noise hits, that we otherwise removed from our training dataset to facilitate the noise impact studies of section 4.1.2.

² This nomenclature and the associated definitions broadly follow [2, 35].

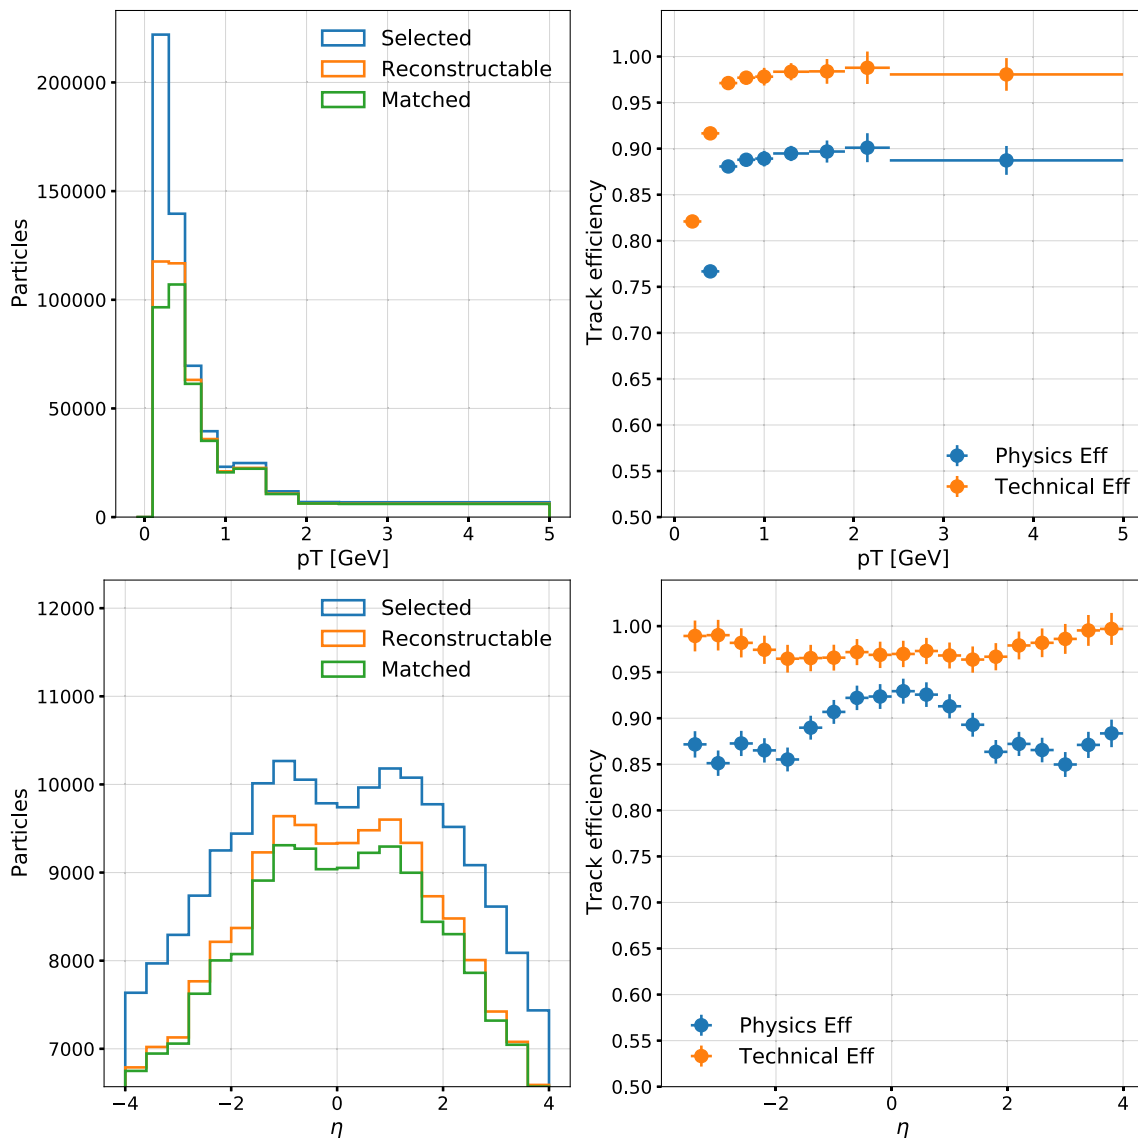


Fig. 5 Top row: selected, reconstructable, and matched particles (left) and tracking efficiency (right) as a function of p_T for particles with $|\eta| < 4$. Bottom row: selected, reconstructable, and matched particles

(left) and tracking efficiency (right) as a function of η for $p_T > 0.5$ GeV. The definition of “selected”, “reconstructable”, and “matched” can be found in Sect. 4.1.1

Measuring precisely the impact of pile-up collisions on tracking performance is beyond the scope of this work, but we can estimate pile-up’s impact on tracking performance by plotting efficiency and purity as a function of the number of spacepoints in the detector. Figure 6 shows that the effect of the increased detector occupancy is a smooth performance degradation $O(\%)$. In future work, we will study the origin of this degradation to achieve the stable performance of traditional algorithms [36].

The impact of noise spacepoints can be estimated using the TrackML dataset by studying the inference performance of the tracking pipeline, trained without any noise spacepoints, as a function of the fraction of noise spacepoints (up to a

maximum of 20% of the total). Table 1 shows the technical tracking efficiency and purity for different noise levels. The efficiency decreases by $\simeq 1.6\%$ and the purity by $\simeq 5.4\%$ when 20% of noise spacepoints are presented. The loss of efficiency happens primarily for particles with $p_T < 500$ MeV (Fig. 7).

Detector misalignment effects are approximated by shifting by up to 1 mm the x -axis of all spacepoints in the innermost TrackML barrel detector layer or the four innermost layers (volume 8 in Fig. 1). In both cases, the impact on the tracking efficiency is less than 0.1%. However, studying in depth misalignments, and other detector effects, requires access to experiment detailed detector simulation data. We

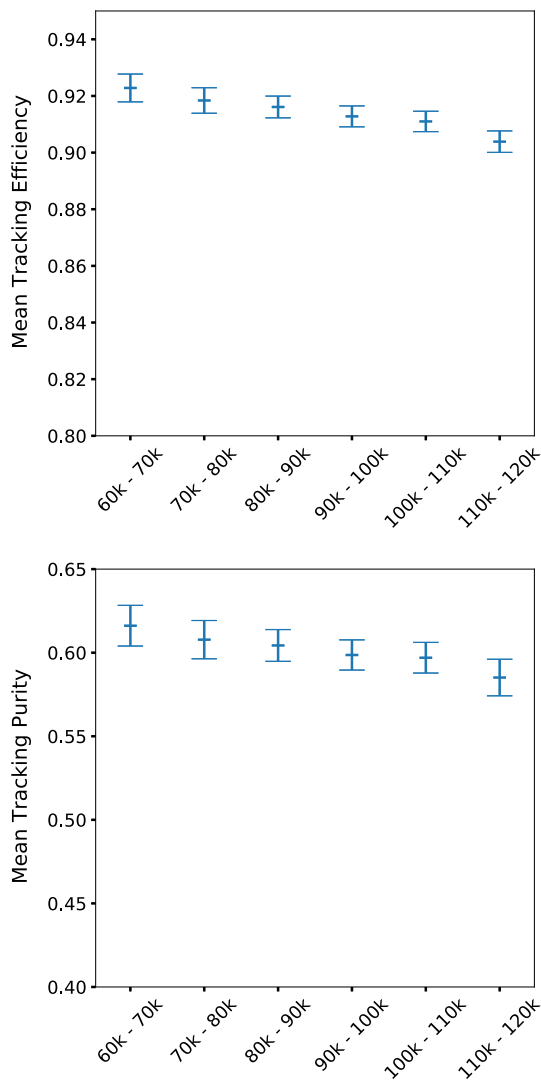


Fig. 6 Mean and standard deviation of the technical efficiency (top) and purity (bottom) as a function of the total number of spacepoints in an event

Table 1 Technical efficiency and purity for different noise fractions ($N_{spp}^{noise}/N_{spp} * 100\%$)

Noise	ϵ_{tech}	Purity
0	91.5	59.3
4%	91.5	59.3
8%	91.1	58.0
12%	90.9	56.8
16%	92.2	54.8
20%	89.9	53.9

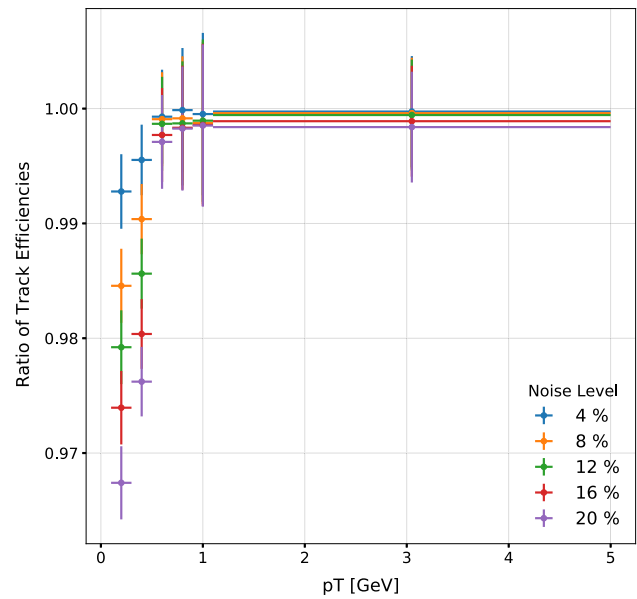


Fig. 7 Relative technical efficiency as a function of p_T . Each curve shows the ratio of $eff(noise = N\%)/eff(noise = 0)$

leave these studies as future work to be performed in collaboration with each experiment.

4.2 Distributed training performance

Our training sample consists of 7500 pileup events from the TrackML dataset. It takes about 1.5 days to train the Exa.TrkX pipeline on a Nvidia A100 GPU for a set of hyper-parameters. It is therefore desirable to use distributed training to parallelize model training and hyper-parameter optimization (HPO). This study relied on data parallel training [37] implemented using Horovod [38] and Tensorflow’s TF.DISTRIBUTED framework [39]. Horovod supports distributed training across multiple nodes, while TF.DISTRIBUTED allows to use the same code across CPUs, TPUs, and GPUs.

For this study, the TrackML pipeline is trained on up to 64 Nvidia V100 GPUs across eight NERSC Cori-GPU computing nodes. Using the Horovod framework (Fig. 8), training time is reduced from 22 min, with 1 GPU, to 0.5 min with 64 GPUs.⁵ The strong scaling efficiency⁶ is about 90% with 2 GPUs and 75% with 8 GPUs. This deviation from ideal scaling is due to the model setup time and data movement costs.

⁵ All measurements in this section were taken training on spacepoints from the barrel region of the TrackML detector. For comparison, training with spacepoints from the whole detector takes $\simeq 70$ minutes per epoch on one Nvidia A100 GPU.

⁶ Defined as $t_1/(N \times t_N) * 100\%$ where t_N is the time to train on a fixed total number of events across N GPUs.

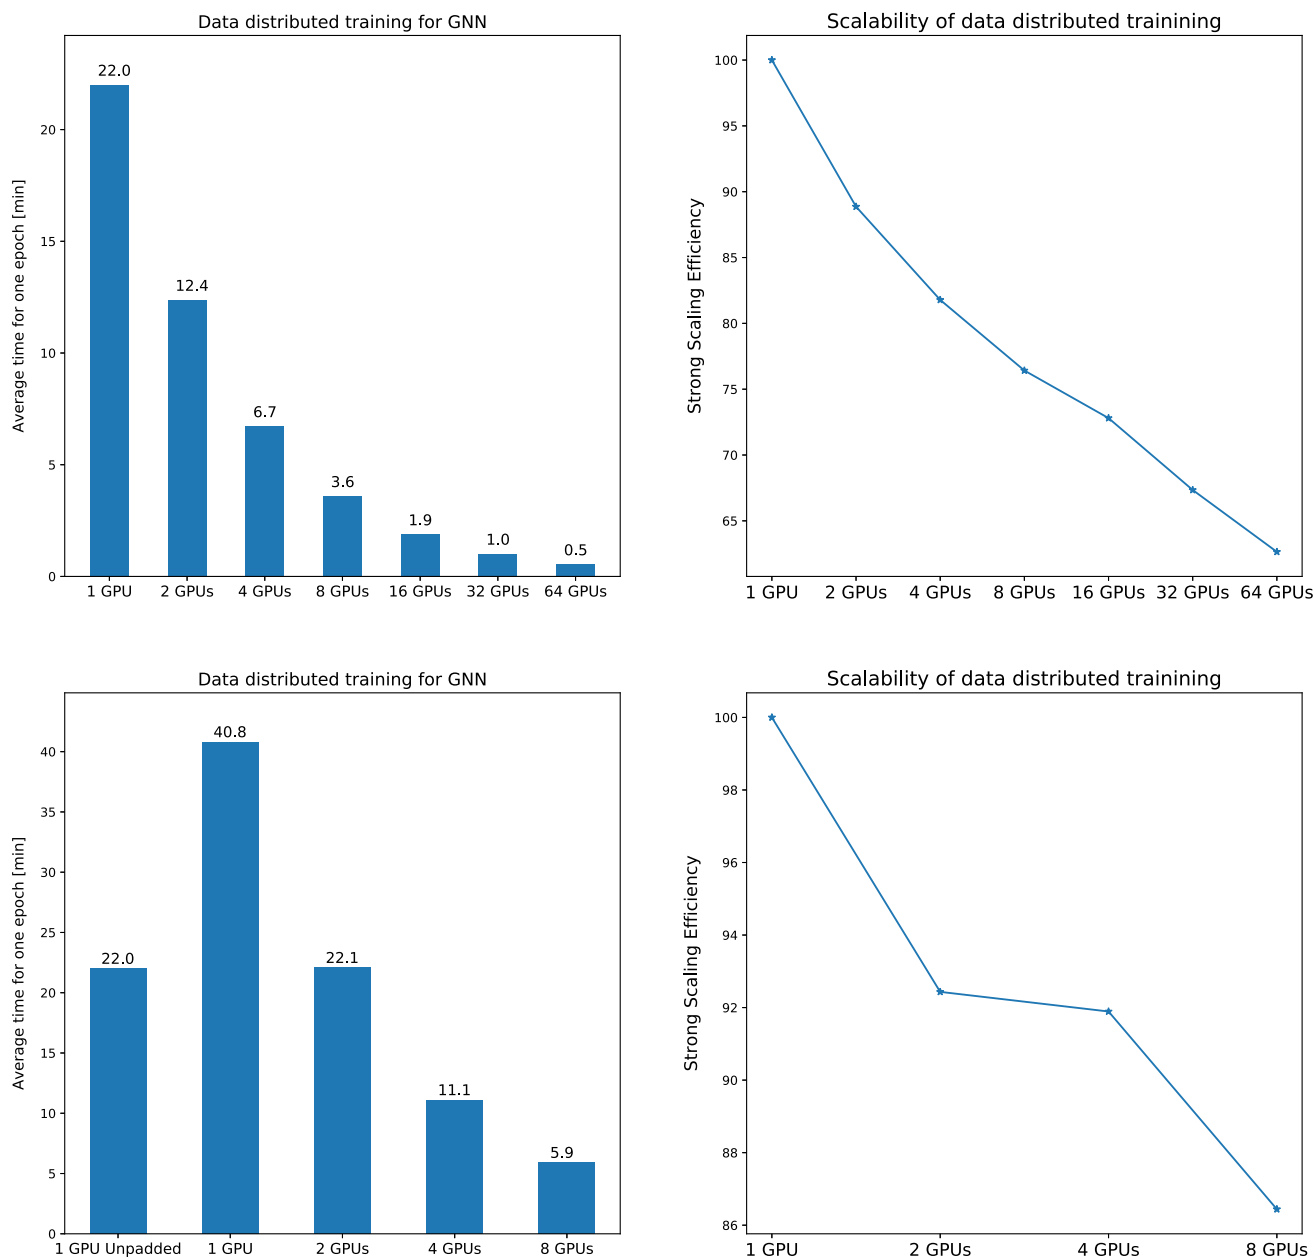


Fig. 8 Time per training epoch (left) and Strong scaling efficiency (right) for GNN’s distributed training. The top row refers to the Horovod implementation, the bottom row to the TF.DISTRIBUTED one. The first bin in the bottom left diagram refers to the serial case, in which the input graph is not padded

Figure 8 also shows the scaling behaviour of the TF.DISTRIBUTED implementation. Since this implementation requires all input data to be of the same size, we have to pad all input graphs to a fixed size. This essentially doubles the time needed to train one epoch, that increases from 22 minutes for dynamic input graph sizes to 41 min for constant graph sizes. Leaving aside this fixed overhead, TF.DISTRIBUTED appears to scale better than Horovod, achieving $\approx 85\%$ strong scaling efficiency with 8 GPUs.

4.3 Inference performance on CPU and GPU

It is crucial to characterize the computational cost of the end-to-end learned tracking algorithm. We rely on the PYTORCH and TENSORFLOW libraries to optimize our inference pipeline on CPU and GPU. The execution time for the inference pipeline has been measured on two hardware platforms: Nvidia V100 GPUs with 16 GB on-board memory, and Intel Xeon 6148s (Skylake) CPUs with 40 cores and 192 GB memory per node. The inputs to the filtering step do not fit into the

Table 2 Average inference time for synchronous execution of the TrackML pipeline benchmarked on CPUs and GPUs. For these step-by-step measurements, we force the pipeline to execute serially by calling

`torch.cuda.synchronize` after each step. The total inference time comprises all the steps including ones not listed in the table

	Wall time [s] on Xeon 6148s single core	Wall time [s] on Nvidia V100 synchronous
Data loading	0.0049 ± 0.0153	0.0023 ± 0.0003
Embedding	3.02 ± 0.39	0.024 ± 0.003
Build edge	66 ± 13	0.76 ± 0.10
Filtering	99 ± 19	1.57 ± 0.34
GNN	27 ± 2	0.45 ± 0.06
Labeling	3.23 ± 0.34	0.08 ± 0.01
Total (sync)	202 ± 35	3.3 ± 0.5

GPU memory. Therefore, edge filtering for one event is executed in mini-batches with a fixed batch size of 800k edges. Typically, the inputs to the filtering from one event are split into seven batches, leading to additional computational cost for moving data from host to GPU. The peak GPU memory consumption is about 15.7 GB as obtained from the Nvidia profiling tool.

Averaging over 500 events, it takes 2.2 ± 0.3 wall-clock seconds per event (as measured by measured by the python module *time*) to run the inference pipeline on the GPU and 202 ± 35 seconds to run it on a single CPU core. This total execution time includes every step of the calculation, and in particular the time needed to move data from host to GPU. Table 2 breaks down the wall-clock time for the most significant steps of the pipeline. The results show how the graph creation and filtering steps are the biggest targets for further optimization in order to surpass traditional algorithms in terms of inference time [40].

In addition, Fig. 9 shows how the total inference time depends almost linearly on the number of spacepoints in the event for both CPUs and GPUs. The step-like dispersion in the GPU case is due to the splitting of the inputs to the filtering step into mini-batches. A step-like jump indicates one more mini-batch is added.

Many optimizations were introduced to the pipeline in order to achieve these GPU timings, which before optimization took over 20 seconds per event. These improvements include porting all data processing to the GPU-accelerated CuPy library [41], writing custom sparse operations for graph processing (e.g. doublet-to-triplet conversion [42], graph intersection methods), using FAISS [43] for large-k NN graph construction, and performing track labelling with CuGraph's connected component algorithm on GPU [44]⁷. These improvements are specific to the inference stage; training optimizations will be discussed in the following section, and ongoing developments in Sect. 6. No CPU-specific optimization was performed in this work.

⁷ On CPU, track labeling uses the DBSCAN algorithm [45].

5 Discussion

The performance given above is the result of experimentation across various feature sets, architectures, model configurations and hyperparameters. It has also been necessary to overcome a variety of training hurdles in terms of memory and computational availability. We describe here training and inference details that should allow a reader to reproduce these results on the provided codebase.

5.1 Feature set

The input dataset includes both spatial coordinates and highly granular pixel cluster shape information. Graph construction (the second pipeline step in Fig. 4, that includes learned embedded space model and edge filter model) appears to benefit significantly from the cluster shape information, approximately doubling the purity for a held fixed high efficiency. The summary cluster shape statistics include the number of channels and the total charge deposited, as well as local and global representations of the cluster as a high-level feature vector. Details about the calculation of this feature vector as well as a thorough exploration of the effect of cluster shape information on seeding performance are provided in Ref. [30]. Cluster shape information does not appear to improve the performance of the GNN, and in fact seems to degrade it. This suggests that the width of the GNN hidden layers is not great enough to capture the functional relationship of cluster information between nodes. Scaling to a width that properly explores this question would require more memory than available on the Nvidia A100 GPUs used for this study.

Depending on the final goal of the pipeline, further features can be included in the loss calculation in order to bias the model towards desired regions. For example, if our aim is to maximize the TrackML score (described in Ref. [2]) – a weighting function s_i that places more importance on a spacepoint i from a longer and higher p_T track, and in the first and last sets of detector layers – we can weight-up true edges

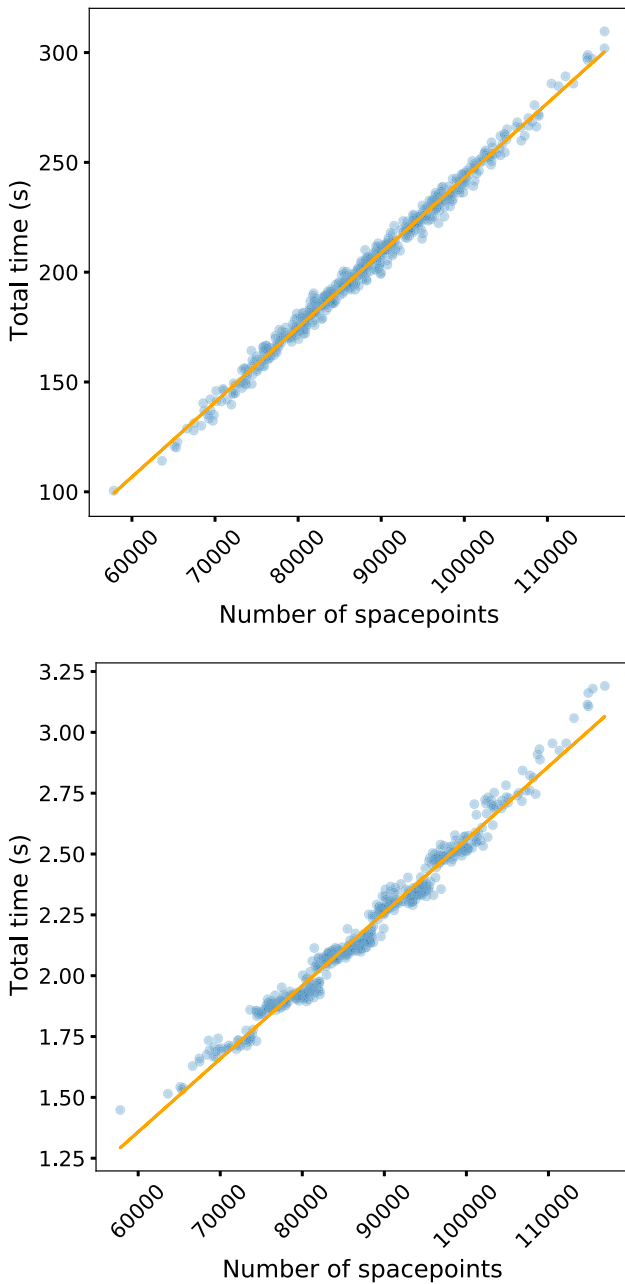


Fig. 9 Total inference time as a function of number of spacepoints in each event for CPUs (top) and GPUs (bottom)

by this function, normalized to have a mean of weight = 1. To measure the performance of models trained to this goal, we introduce a *weighted* purity measure. Weighted purity is defined as a function the TrackML weights w_{ij} and the truth $y_{ij} \in \{0, 1\}$ of each edge connecting spacepoint i and spacepoint j ,

$$\text{Purity}_{\text{weighted}} = \frac{\sum_{ij} w_{ji} y_{ij}}{\sum_{ij} w_{ij}}$$

$$w_{ij} = \begin{cases} \frac{1}{2}(s_i + s_j), & \text{if } y_{ij} = 1 \\ 1, & \text{if } y_{ij} = 0 \end{cases}$$

We see significant improvements in this metric when validating on the weighted model: the Embedding Network improves from a weighted purity of $1.7\% \pm 0.2\%$ to $2.0\% \pm 0.3\%$, while the Filter Network improves from a weighted purity of $8.4\% \pm 0.6\%$ to $11.7\% \pm 1.0\%$. Given this weighting, the model learns to prioritize higher p_T and longer tracks, while disregarding less informative tracks. Using this bias, we can achieve the same TrackML score with a constructed graph size reduced by approximately 25%. Using this technique to improve the TrackML score is an ongoing work.

5.2 Graph construction

Having chosen a feature set, to train the learned embedding space we use a training paradigm commonly referred to as a Siamese Network [46], where a particular spacepoint – called the *source* – is run through an MLP, here 6 layers each with 512 hidden channels, hyperbolic tan activations, and layer normalization. The final layer of the MLP takes the features to an 8-dimensional latent space. A different, comparison spacepoint – called the *target* – is also run through this same Embedding Network, and the L2 norm distance d in the latent space between the source and target enters a comparative hinge loss

$$\mathcal{L}_{\text{hinge}} = \begin{cases} d^p, & \text{if } y_{ij} = 1 \\ \max(0, 1 - d^p), & \text{if } y_{ij} = 0 \end{cases} \tag{4}$$

where p is a hyperparameter that we choose to be 2.

If the source i and target j spacepoints share an edge in the event’s truth graph,⁸ we designate them as neighbours with $y_{ij} = 1$, otherwise they are designated $y_{ij} = 0$. In this way, the hinge loss draws together truth graph neighbors and repels non-neighbors.

Training performance of the Embedding Network is highly dependent on choice of source-target example pairs. In early epochs, it is enough to choose random pairs. However, at some point, many random pairs will contribute no gradient to the loss, as they will be separated by a distance greater than the margin. At that point, it is useful to implement hard negative mining [47]. We run a GPU-optimised k-nearest-neighbor (KNN) algorithm⁹ to mine examples around each

⁸ One can also designate $y_{ij} = 1$ for source and target in the same track, rather than *immediate neighbors* in the track. This does lead to similar performance in later stages of the pipeline, but the more lax concept of truth leads to graphs around three times more dense than the strict track neighbor definition.

⁹ We use two high-performance libraries, FAISS [43] and Pytorch3D [48], depending on number of nearest neighbors k .

source vector, within the hinge margin $d = 1$. The computational overhead of the KNN step is significantly offset by the examples mined which all contribute to the loss.

A similar technique is used in the Filter Network, where the vast majority of the edges produced from the graph construction in the embedded space are easy to classify as fake. This is already a highly imbalanced dataset, with around 98.5% of edges fake. Again, within several epochs, the Filter Network is able to classify many of these as fake, so we balance each batch with all true edges, the same number of hard negatives (i.e. negatives the filter is unsure of) and the same number of easy negatives (to maintain performance on these edges). The Filter Network is a MLP that takes the 24-feature concatenated edge features and feeds forward through 3 layers of 1024 hidden channels, to a binary cross-entropy loss function.

5.3 GNN edge classification

In choosing the best GNN architecture, memory usage remains a significant constraint. The Interaction Network (IN) [34] presented in these results does appear to marginally attain the best performance against Attention Graph Neural Networks (AGNN) [14, 49] – the other class of GNN considered for the pipeline. However, both of these networks require gradients to be retained in memory for every graph edge. Indeed, this anisotropic treatment of edges (i.e. a node is able to receive the messages of each of its neighbors in a non-uniform way) is what allows these two architectures to be so expressive. Depending on hardware availability, we have found two solutions to the memory constraint. Access to next-generation Nvidia A100 GPUs allowed an IN to be trained with 8 steps of message passing, aggregating edge features at each node, and each node and concatenated edge features passing through two-layer MLPs of [128, 64] hidden features and ReLU activations [50]. Choice of aggregation function should be permutation invariant. In this work, we take it to be a summation.

For lower-memory GPUs, such as the Nvidia V100, we attained similar performance training the AGNN architecture, with [64, 64, 64]-channel MLPs applied to each edge and node. Adding residuals [51] across the 8 message passing steps greatly improved performance in this case. To fit full-event training on a single V100, it was necessary to employ various techniques, such as mixed precision training and gradient checkpointing. The latter stores only the input of each layer, not the gradients. On the backward pass, gradients are re-calculated on the fly, allowing for a $4\times$ reduction in memory usage for an 8-iteration GNN. Another technique explored is to split the events piecemeal and train on each

piece as a standalone batch. There is a noticeable impact on performance due to messages being interrupted at the graph edges. In future work, we will present ongoing efforts to parallelise these graph pieces across multiple GPUs, retaining the high performance that full-event training allows.

5.4 Physics-inspired data augmentation

Preliminary work on using coordinate transforms to augment the training data has been explored with varying degrees of success. In this study, focused on track seeding, only the innermost detector layers (volumes 7–9 in Fig. 1) were used.

One promising approach is to make a copy of each graph in the training set that has been reflected across the phi-axis [52]. The phi reflection creates the charge conjugate graph and helps to balance any asymmetry between positive and negatively charged particles within the training set. Using the phi-reflected graphs boosts efficiency by $\simeq 2\%$ and purity by $\simeq 1\%$ in the barrel. This performance boost comes at the cost of doubling the training time. In future work, we will investigate the opportunity of integrating charge conjugation symmetry into the network itself.

A second promising trick is to use a Hough Transform [6, 7] on the graph to create edge features. Using the Hough parameters as edge features boosts efficiency by $\simeq 2\%$ and purity by $\simeq 1\%$. A further efficiency boost of $\simeq 3\%$ (and $\simeq 2\%$ to purity) comes from using the Hough accumulator to extract an edge weight. This edge weight effectively pools information from every node, and therefore comes at a large computational cost (filling the accumulator in Hough space). On the other hand, the Hough parameters can be computed quickly from the two nodes that define the edge.

6 Conclusions and future work

This work shows how a tracking pipeline based on geometric deep learning can achieve state-of-the-art computing performance that scales linearly with the number of space-points, showing great promise for the next generation of HEP experiments. The inference pipeline has been optimized on GPU systems, on the assumption that the next generation of HEP experiments will have widespread access to accelerators either locally in heterogeneous systems [27, 53] or remotely [54, 55].

Within the simplifying assumptions of the TrackML dataset, we have shown how the Exa.TrkX pipeline could meet the tracking performance requirements of current collider experiments. Preliminary studies suggest that this performance should be robust against systematic effects like detector noise, misalignment, and pile-up.

Much remains to be done to validate these promising results. To this end, the Exa.TrkX project is collaborating

Fastest performance is obtained with FAISS for $k > 35$, Pytorch3D for $k \leq 35$.

with physicists from ATLAS [56], CMS [57], DUNE [58], ICARUS [59], and MuonE [60].

The goal is to adapt the Exa.TrkX pipeline to each experiment's needs and simulated datasets, measure its performance and robustness against systematic effects according to the experiment metrics. For example, it is crucial for HL-LHC experiments to study the performance of tracking algorithms in dense environments, like high- p_T jets. Given the interest in long-lived particle observation at the HL-LHC, it will also be important to study the performance of the Exa.TrkX pipeline for tracks coming from a displaced vertex.¹⁰

On the computational side, there are several optimization opportunities to explore systematically, including mixed precision training, multi-GPU training and inference with graph data parallelisation (that is, one event spread across multiple GPUs) [61]; locality sensitive hashing to speed-up KNN/graph construction stage [62], model quantization, operator fusion and other improvements with TensorRT [63], clustering of final node embeddings rather than hard connected components method with GravNet-style architectures [64].

The distributed training results presented in this work are promising but still preliminary. To fully exploit the capabilities of upcoming HPC systems and to further reduce training time while potentially pushing further on model size, it will be beneficial to perform further studies on large scale training of GNNs for track reconstruction. Given the size of the input graphs, this problem may be amenable to training techniques which parallelise the processing of input graphs across multiple GPUs in training.

Finally, it will be interesting to measure the computing performance of (parts of) the Exa.TrkX pipeline on domain-specific accelerators like Google TPU [65] and Graph-Core IPU [66], comparing power consumption, latency and throughput with “traditional” GPUs.

7 Software availability

A growing number of groups are studying the application of graph networks to HEP reconstruction (see [67] for a recent review). Some of these works [24,27–31,33] have strong connections with the Exa.TrkX project. To promote collaboration and reproducibility, the Exa.TrkX software is available from the HEP Software Foundation's Trigger and Reconstruction GitHub.¹¹ A pipeline of re-usable modules is implemented within the Pytorch Lightning system, which allows

¹⁰ It may be worth noticing that in LArTPC applications [28] all tracks come from a displaced vertex.

¹¹ <https://hsf-reco-and-software-triggers.github.io/Tracking-ML-Exa.TrkX>.

for uncluttered and simple model definitions. As each stage of the pipeline is dependent, logging utilities are integrated that allow a specific combination of stages and hyperparameters to be trackable and reproducible. Extensive documentation is provided to help track reconstruction groups start exploring geometric learning. The roadmap for this repository includes adding performance metrics to the codebase; a taxonomy of model features; and short tutorials in each of the available applications.

Acknowledgements This research was supported in part by: – the U.S. Department of Energy's Office of Science, Office of High Energy Physics, under Contracts No. DE-AC02-05CH11231 (CompHEP Exa.TrkX) and No. DE-AC02-07CH11359 (FNAL LDRD 2019.017); – the Exascale Computing Project (17-SC-20-SC), a joint project of DOE's Office of Science and the National Nuclear Security Administration; the National Science Foundation under Cooperative Agreement OAC-1836650. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. We are grateful to Google Co. for providing early access to Nvidia A100 instances in the context of the US ATLAS/Google Cloud Platform collaboration. Finally, we thank Marcin Wolter (IFJ PAN), Ben Nachman, Alex Sim and Kesheng Wu (LBNL) for the useful discussions.

Data Availability Statement This manuscript has associated data in a data repository. [Authors' comment: <https://competitions.codalab.org/competitions/20112>.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.
Funded by SCOAP³.

References

1. I.B. Alonso, O. Brüning, P. Fessia, M. Lamont, L. Rossi, L. Tavian, M. Zerlauth, High luminosity large hadron collider HL-LHC technical design report. CERN Yellow Rep. **10** (2020). <https://doi.org/10.23731/CYRM-2020-0010>. <https://e-publishing.cern.ch/index.php/CYRM/issue/view/127>
2. S. Amrouche et al., The tracking machine learning challenge: accuracy phase. [arXiv:1904.06778](https://arxiv.org/abs/1904.06778) [hep-ex]
3. A. Strandlie, R. Frühwirth, Track and vertex reconstruction: from classical to adaptive methods. Rev. Mod. Phys. **82**, 1419–1458 (2010). <https://doi.org/10.1103/RevModPhys.82.1419>
4. ATLAS Collaboration, Performance of the ATLAS track reconstruction algorithms in dense environments in LHC run 2.

- Eur. Phys. J. C **77**(10), 673 (2017). <https://doi.org/10.1140/epjc/s10052-017-5225-7>. arXiv:1704.07983
5. CMS Collaboration, S. Chatrchyan et al., Description and performance of track and primary-vertex reconstruction with the CMS tracker. *JINST* **9**(10), P10009 (2014). <https://doi.org/10.1088/1748-0221/9/10/P10009>. arXiv:1405.6569 [physics.ins-det]
 6. R.. O. Duda, P. E. Hart, Use of the Hough transformation to detect lines and curves in pictures. *Commun. ACM* **15**(1), 11–15 (1972). <https://doi.org/10.1145/361237.361242>
 7. J. Gradin, M. Mårtensson, R. Brenner, Comparison of two hardware-based hit filtering methods for trackers in high-pileup environments. *JINST* **13**(04), P04019 (2018). <https://doi.org/10.1088/1748-0221/13/04/P04019>. arXiv:1709.01034 [physics.ins-det]
 8. D. Funke, T. Hauth, V. Innocente, G. Quast, P. Sanders, D. Schieferdecker, Parallel track reconstruction in CMS using the cellular automaton approach. *J. Phys. Conf. Ser.* **513**, 052010 (2014). <https://doi.org/10.1088/1742-6596/513/5/052010>
 9. D. Rohr, S. Gorbunov, M.O. Schmidt, R. Shahoyan, GPU-based online track reconstruction for the ALICE TPC in run 3 with continuous read-out. *EPJ Web Conf.* **214**, 01050 (2019). <https://doi.org/10.1051/epjconf/201921401050>. arXiv:1905.05515 [physics.ins-det]
 10. ATLAS Collaboration, Computing and Software Public Results (2017). <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ComputingandSoftwarePublicResults>
 11. CMS Collaboration, CMS Tracking POG Performance Plots For 2017 with PhaseI pixel detector. (2017). <https://twiki.cern.ch/twiki/bin/view/CMSPublic/TrackingPOGPerformance2017MC>
 12. ATLAS Collaboration, Fast Track Reconstruction for HL-LHC. Tech. Rep. ATL-PHYS-PUB-2019-041, CERN, Geneva (2019). <https://cds.cern.ch/record/2693670>
 13. HEP.TrkX, HEP advanced tracking algorithms with cross-cutting applications (2016). <https://heptrkx.github.io/>
 14. S. Farrell et al., Novel deep learning methods for track reconstruction, in *4th International Workshop Connecting The Dots 2018* (CTD2018) Seattle, Washington, USA, March 20–22, 2018 (2018). arXiv:1810.06111 [hep-ex]
 15. Exa.TrkX, HEP advanced tracking algorithms at the exascale (2019). <https://exatrnx.github.io/>
 16. M.M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, P. Vandergheynst, Geometric deep learning: going beyond Euclidean data. *IEEE Signal Process Mag* **34**(4), 8–42 (2017). <https://doi.org/10.1109/MSP.2017.2693418>
 17. N. Choma et al., Track seeding and labelling with embedded-space graph neural networks. **6** (2020). arXiv:2007.00149 [physics.ins-det]
 18. S. Farrell et al., The HEP.TrkX Project: deep neural networks for HL-LHC online and offline tracking, in *Proceedings, Connecting The Dots/Intelligent Tracker* (CTD/WIT 2017): Orsay, France, March 6–9, 2017, vol. 150. (2017), p. 00003. <https://doi.org/10.1051/epjconf/201715000003>
 19. CMS Collaboration, V. Bertacchi, DeepCore: convolutional neural network for high p_T jet tracking. arXiv:1910.08058 [physics.ins-det]
 20. S. Hochreiter, J. Schmidhuber, Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997). <https://doi.org/10.1162/neco.1997.9.8.1735>
 21. A. Tsaris, D. Anderson, J. Bendavid, P. Calafiura, G. Cerati, J. Esseiva, S. Farrell, L. Gray, K. Kapoor, J. Kowalkowski, M. Mudigonda, P.P. Spentzouris, M. Spiropoulou, J.-R. Vlimant, S. Zheng, D. Zurawski, The HEP.TrkX project: Deep learning for particle tracking. *J. Phys. Conf. Ser.* **1085**, 042023 (2018). <https://doi.org/10.1088/1742-6596/1085/4/042023>
 22. S. Amrouche et al., The tracking machine learning challenge: accuracy phase. arXiv:1904.06778 [hep-ex]
 23. K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using rnn encoder-decoder for statistical machine translation (2014)
 24. C. Biscarat, S. Caillou, C. Rougier, J. Stark, J. Zahreddine, Towards a realistic track reconstruction algorithm based on graph neural networks for the hl-lhc. arXiv:2103.00916 [physics.ins-det]
 25. X. Ju et al., Graph neural networks for particle reconstruction in high energy physics detectors, in *33rd Annual Conference on Neural Information Processing Systems* vol. 3 (2020). arXiv:2003.11603 [physics.ins-det]
 26. T.D. Le, H. Imai, Y. Negishi, K. Kawachiya, Tfllms: Large model support in tensorflow by graph rewriting. arXiv:1807.02037 [cs.LG]
 27. J. Pata, J. Duarte, J.-R. Vlimant, M. Pierini, M. Spiropulu, MLPF: efficient machine-learned particle-flow reconstruction using graph neural networks. arXiv:2101.08578 [physics.data-an]
 28. J. Hewes, A. Aurisano, G. Cerati, J. Kowalkowski, C. Lee, W. keng Liao, A. Day, A. Agrawal, M. Spiropulu, J.-R. Vlimant, L. Gray, T. Klijsma, P. Calafiura, S. Conlon, S. Farrell, X. Ju, D. Murnane, Graph neural network for object reconstruction in liquid argon time projection chambers (2021)
 29. A. Heintz et al., Accelerated charged particle tracking with graph neural networks on FPGAs, in *34th Conference on Neural Information Processing Systems*, vol. 11 (2020). arXiv:2012.01563 [physics.ins-det]
 30. P.J. Fox, S. Huang, J. Isaacson, X. Ju, B. Nachman, Beyond 4d tracking: using cluster shapes for track seeding. arXiv:2012.04533 [physics.ins-det]
 31. S. Amrouche, M. Kiehn, T. Golling, A. Salzburger, Hashing and metric learning for charged particle tracking. arXiv:2101.06428 [hep-ex]
 32. W. Guan, G. Perdue, A. Pesah, M. Schulz, K. Terashi, S. Vallecorsa, J.-R. Vlimant, Quantum machine learning in high energy physics. arXiv:2005.08582 [quant-ph]
 33. C. Tüysüz, K. Novotny, C. Rieger, F. Carminati, B. Demirköz, D. Dobos, F. Fracas, K. Potamianos, S. Vallecorsa, J.-R. Vlimant, Performance of particle tracking using a quantum graph neural network. arXiv:2012.01379 [quant-ph]
 34. P.W. Battaglia, R. Pascanu, M. Lai, D.J. Rezende, K. Kavukcuoglu, Interaction networks for learning about objects, relations and physics. *CoRR*, abs/1612.00222 (2016). arXiv:1612.00222
 35. ATLAS Collaboration, Technical Design Report for the ATLAS Inner Tracker Pixel Detector. Tech. Rep. CERN-LHCC-2017-021. ATLAS-TDR-030, CERN, Geneva (2017). <https://cds.cern.ch/record/2285585>
 36. A. Collaboration, Technical Design Report for the ATLAS Inner Tracker Pixel Detector, Tech. Rep. ATLAS-TDR-030, CERN, Geneva (2017)
 37. T. Ben-Nun, T. Hoefler, Demystifying parallel and distributed deep learning: an in-depth concurrency analysis. arXiv:1802.09941 [cs.LG]
 38. A. Sergeev, M.D. Balso, Horovod: fast and easy distributed deep learning in TensorFlow. arXiv:1802.05799 [cs.LG]
 39. M. Abadi et al., TensorFlow: large-scale machine learning on heterogeneous distributed systems. arXiv:1603.04467 [cs.DC]
 40. ATLAS Collaboration, Expected tracking performance of the ATLAS inner tracker at the HL-LHC, Tech. Rep. ATL-PHYS-PUB-2019-014, CERN, Geneva (2019). <https://cds.cern.ch/record/2669540>
 41. R. Okuta, Y. Unno, D. Nishino, S. Hido, C. Cupy, A numpy-compatible library for nvidia gpu calculations, in *31st Conference on Neural Information Processing Systems (NIPS 2017)* (2017). http://learningsys.org/nips17/assets/papers/paper_16.pdf

42. M. Fey, J.E. Lenssen, Fast graph representation learning with PyTorch Geometric, in *ICLR Workshop on Representation Learning on Graphs and Manifolds* (2019)
43. J. Johnson, M. Douze, H. Jégou, Billion-scale similarity search with GPUs. [arXiv:1702.08734](https://arxiv.org/abs/1702.08734)
44. CuGraph, (2020) <https://github.com/rapidsai/cugraph>. Accessed 01 Mar 2021
45. M. Ester, H.-P. Kriegel, J. Sander, X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise, in *Kdd* (AAAI Press, 1996), pp. 226–231
46. D. Chicco, *Siamese Neural Networks: An Overview* (Springer, New York, 2021), pp. 73–94. https://doi.org/10.1007/978-1-0716-0826-5_3
47. B. Harwood, B.G.V. Kumar, G. Carneiro, I. Reid, T. Drummond, Smart mining for deep metric learning, in *ICCV 2017: International Conference on Computer Vision*, vol. 10. (2017), p. 2840–2848. <https://doi.org/10.1109/ICCV.2017.307>
48. N. Ravi, J. Reizenstein, D. Novotny, T. Gordon, W.-Y. Lo, J. Johnson, G. Gkioxari, Accelerating 3D Deep Learning with PyTorch3D. [arXiv:2007.08501](https://arxiv.org/abs/2007.08501)
49. P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, Y. Bengio, Graph attention networks (2017)
50. X. Glorot, A. Bordes, Y. Bengio, Deep sparse rectifier neural networks, in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, ed. by G. Gordon, D. Dunson, M. Dudík, vol. 15 of Proceedings of Machine Learning Research. JMLR Workshop and Conference Proceedings, Fort Lauderdale, FL, USA, 11–13. (2011), p. 315–323. <http://proceedings.mlr.press/v15/glorot11a.html>
51. K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition (2015)
52. L. Perez, J. Wang, The effectiveness of data augmentation in image classification using deep learning. *CoRR* **abs/1712.04621** (2017). [arXiv:1712.04621](https://arxiv.org/abs/1712.04621)
53. F. Fahim, B. Hawks, C. Herwig, J. Hirschauer, S. Jindariani, N. Tran, L.P. Carloni, G.D. Guglielmo, P. Harris, J. Krupa, D. Rankin, M.B. Valentin, J. Hester, Y. Luo, J. Mamish, S. Orgrenci-Memik, T. Aarrestad, H. Javed, V. Loncar, M. Pierini, A.A. Pol, S. Summers, J. Duarte, S. Hauck, S.-C. Hsu, J. Ngadiuba, M. Liu, D. Hoang, E. Kreinar, Z. Wu, hls4ml: an open-source code-sign workflow to empower scientific low-power machine learning devices (2021)
54. J. Krupa, K. Lin, M. Acosta Flechas, J. Dinsmore, J. Duarte, P. Harris, S. Hauck, B. Holzman, S.-C. Hsu, T. Klijnsmas et al., Gpu coprocessors as a service for deep learning inference in high energy physics. *Mach. Learn. Sci. Technol.* **2**(3), 035005 (2021). <https://doi.org/10.1088/2632-2153/abec21>
55. V. Kuznetsov, L. Giommi, D. Bonacorsi, Mlaas4hep: machine learning as a service for hep (2020)
56. ATLAS Collaboration, ATLAS Collaboration, The ATLAS Experiment at the CERN Large Hadron Collider, *JINST*, **3**, S08003 (2008). <https://doi.org/10.1088/1748-0221/3/08/S08003>
57. CMS Collaboration, S. Chatrchyan et al., The CMS experiment at the CERN LHC. *JINST*, **3**, S08004 (2008). <https://doi.org/10.1088/1748-0221/3/08/S08004>
58. Deep underground neutrino experiment. <http://www.dunescience.org/>
59. ICARUS Collaboration, L. Bagby et al., Overhaul and Installation of the ICARUS-T600 Liquid Argon TPC Electronics for the FNAL Short Baseline Neutrino Program. *JINST* **16**(01), P01037 (2021). <https://doi.org/10.1088/1748-0221/16/01/P01037>. [arXiv:2010.02042](https://arxiv.org/abs/2010.02042) [physics.ins-det]
60. G. Abbiendi et al., Measuring the leading hadronic contribution to the muon $g-2$ via μe scattering. *Eur. Phys. J. C* **77**(3), 139 (2017). <https://doi.org/10.1140/epjc/s10052-017-4633-z>. [arXiv:1609.08987](https://arxiv.org/abs/1609.08987) [hep-ex]
61. S. Scardapane, I. Spinelli, P.D. Lorenzo, Distributed training of graph convolutional networks. *IEEE Trans. Signal Inf. Process. Netw.* **7**, 87–100 (2021). <https://doi.org/10.1109/tsipn.2020.3046237>
62. P. Indyk, R. Motwani, Approximate nearest neighbors: towards removing the curse of dimensionality, in *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98 (Association for Computing Machinery, New York, 1998), p. 604–613. <https://doi.org/10.1145/276698.276876>
63. NVIDIA TensorRT, (2020) <https://docs.nvidia.com/deeplearning/tensorrt/index.html>. Accessed 2021-03-01
64. S.R. Qasim, J. Kieseler, Y. Iiyama, M. Pierini, Learning representations of irregular particle-detector geometry with distance-weighted graph networks. *Eur. Phys. J. C* **79**(7), 608 (2019). <https://doi.org/10.1140/epjc/s10052-019-7113-9>. [arXiv:1902.07987](https://arxiv.org/abs/1902.07987) [physics.data-an]
65. N.P. Jouppi et al., In-datacenter performance analysis of a tensor processing unit. *SIGARCH Comput. Archit. News* **45**(2), 1–12 (2017). <https://doi.org/10.1145/3140659.3080246>. [arXiv:1704.04760](https://arxiv.org/abs/1704.04760) [cs.AR]
66. Z. Jia, B. Tillman, M. Maggioni, D.P. Scarpazza, Dissecting the graphcore ipu architecture via microbenchmarking. [arXiv:1912.03413](https://arxiv.org/abs/1912.03413) [cs.DC]
67. J. Duarte, J.-R. Vlimant, Graph neural networks for particle tracking and reconstruction. [arXiv:2012.01249](https://arxiv.org/abs/2012.01249) [hep-ph]