

Machine Learning for ParticleFlow and Software Compensation including Timing

Ulrich Einhaus¹ and Bohdan Dudar²

¹Karlsruhe Institute of Technology, Karlsruhe, Germany.

²Johannes Gutenberg-Universität Mainz, Mainz, Germany.

Contributing authors: ulrich.einhaus@kit.edu; bohdan.dudar@cern.ch;

Abstract

We present recent progress on studies of simulated event reconstruction in particle colliders using machine learning. The two examples given are energy reconstruction of calorimeter hits with a DGCNN and full end-to-end reconstruction of reconstructed particles from calorimeter hits and tracks with a GravNet. The results are work in progress and the impact of adding timing information is notable but small.

Talk presented at the International Workshop on Future Linear Colliders
([LCWS 2025](#)), October 20-24 2025.

This work was carried out in the framework of the ILD Concept Group.
ILD-PHYS-PROC-20XX-YY

1 Introduction

Modern approaches to event reconstruction in particle collider experiments mostly employ a variation of particle flow (PF) [1], in particular developments for future colliders. This approach aims to holistically analyse events in order to reconstruct each detector-stable particle by combining information from the tracking system with the calorimeter and muon chambers. While charged particles are usually identified in the tracker with high efficiency, neutrals depend on the calorimetric system for this. There, particle showers are reconstructed as clusters, and afterwards either connected to tracks for charged particles or identified as a neutral. In order to do this with high efficiency, the calorimeter needs a sufficient granularity to allow for a topological separation of particle showers, in particular in dense jet environments. The output of PF reconstruction are reconstructed particles, so-called particle flow objects (PFOs), and are used for jet clustering. The jet energy resolution (JER) has proven to be an effective measure for the performance of a calorimeter as well as reconstruction algorithms. In PF reconstruction, the JER depends both the intrinsic energy resolution of the calorimeter technology as well as the uncertainty from remaining confusion from mis-assigned calorimeter hits, in particular at high energies, where jets are more boosted and shower overlap is larger, as can be seen in fig. 1.




Figure 1 : Example for jet energy resolution and its different contributions, from [1].

The CALO5D project [2] aims to improve performance of high granularity calorimeters by reducing this confusion in two ways: technological enhancement of the data by adding timing information, and applying machine-learning approaches in order to utilise the event data in the most effective way. In recent years technologies have been developed to provide timing information at the level of 30 ps to detector hits. This is particularly relevant for pile-up reduction in vertexing at the extreme luminosities of the HL-LHC. Adding timing information of this precision to calorimeter hits, albeit with a more moderate resolution of 100 ps to account for the large number

of channels, can help assign hits to clusters. 100 ps is about 3 cm at the speed of light, which is in the order of the travel difference between straight neutral paths and the helical tracks of charged particles in the detector’s solenoid field until they reach the calorimeter system. At the same time, progress in machine learning (ML) techniques allows to address particle flow with neural networks (NNs). These can entail individual steps such as hit clustering or feature extraction, or aim to reconstruct particles equivalent to PFOs directly from detector hits. The advantage of NNs is that it may find and utilise correlations that are difficult to identify “by hand”, i.e. in human-informed conventional algorithms. In addition, it is near trivial to add time information to a NN, as on a technical level it is simply one additional value to provide along with all other inputs. This makes ML approaches an optimal candidate to study the impact of timing in calorimeter reconstruction.

In this paper we highlight two recent such approaches. The first employs a dynamic graph convolutional neural network (DGCNN) to improve on the energy resolution of conventionally reconstructed particles. The second uses a GravNet to perform an end-to-end reconstruction from calorimeter hits to fully reconstructed particles. Both cases need to be considered work in progress at this stage and in both we highlight the – so far rather limited – impact of adding timing information.

2 The ILD Detector Concept

The MC samples we use to train and evaluate our NN models are e^+e^- events simulated with the International Large Detector (ILD) [3] model in full geant-4 based simulation and reconstruction. ILD is one of the detector concepts proposed for future colliders, such as the ILC or FCC-ee. It has the typical onion-like subdetector layer of modern general-purpose particle collider detectors. It entails a vertexing system based on Silicon pixel sensors, a time projection chamber (TPC) with a Silicon strip sensor envelope as main tracker, and a high granularity calorimeter system consisting of a SiW ECal and a scintillator-steel HCal, all inside a 3.5 T solenoid magnet surrounded by a scintillator-based muon system. With these features ILD is optimised for ParticleFlow.

3 Software Compensation

In order to improve on the existing conventional energy reconstruction, we employ a DGCNN. Inputs to the network are all calorimeter hits and their properties: position (x, y, z) , deposited energy and timing information. To evaluate the effect of timing two network trainings were compared: one with ‘perfect timing’, i.e. the time stamp of the first energy deposition in simulation. This is compared to the same network except without time entries, i.e. ‘no timing’. The network is set up to provide a corrected deposited energy $E_{hit,i}$ for each hit i , with the target being to minimise the difference between the true total centre-of-mass energy of the event and the corresponding reconstructed value, defined as $E_{CM} = \sum_i E_{hit,i}$, with a mean squared relative error. The training was performed on samples of e^+e^- to $q\bar{q}$ (only u, d and s) events at centre-of-mass energies of 40, 200 and 500 GeV with a total of 25k events. These samples do not contain machine-related backgrounds and uncertainties. The inference was run on




Figure 2 : Reconstructed event energy distributions, using standard Pandora (PFO), Pandora with cheated clustering (Perfect PFO) and the DGCNN (NN).

these energies and in addition 91 and 350 GeV to assure the network can interpolate these and has not learned the specific centre-of-mass energies.

The resulting reconstructed event energy distributions are displayed in fig. 2 and compared to conventional ParticleFlow with Pandora. The plot shows that the network has learned to interpolate the energies of 91 and 350 GeV. The width of the distributions for the network (NN) is close to the result with Pandora (PFO). If one employs the option of ‘perfect’ reconstruction, Pandora uses truth information to cheat the hit clustering, avoiding the remaining confusion term and leaving the effective pure energy resolution of the calorimeter. The corresponding values are indicated as PFO Perfect. An ideal NN would provide a performance between what is achieved with state-of-the-art Pandora (PFO) and what would be possible if one could overcome confusion entirely (PFO Perfect).

Fig. 3 shows the relative widths of these energy distributions, averaged over all 5 centre-of-mass energies. This width, however, depends on the fraction (or quantile) of the distribution that is taken into account to calculate its standard deviation (std). The smaller this quantile, the smaller the resulting std, with 0.9 being the typical fraction used in the RMS90. The plot shows that for a full distribution the network (red) does better than Pandora (blue), at a quantile of 0.9 they perform similarly and for smaller values Pandora performs slightly better. This shows that the network can deal with outliers of the full distribution better than Pandora, but there is still lots of room for improvement, in particular seeing the difference to PFO Perfect. The comparison between the left and the right plot show the impact of adding timing to the input features of the NN. While the value for std90 is minimally better with timing, the difference is near indistinguishable. In particular this point is an active work in progress.

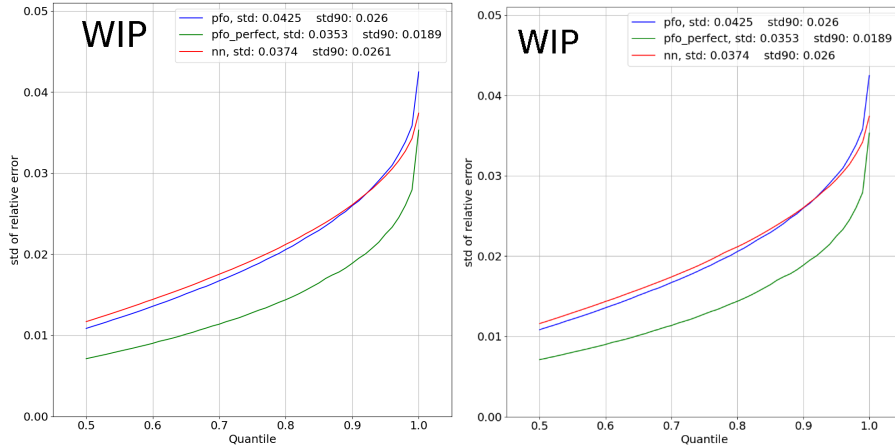


Figure 3 : Result of the DGCNN energy correction: width of the distribution of the reconstructed event energy depending of the quantile of the distribution used. Left: with no timing, right: with perfect timing.

4 End-to-end ParticleFlow

This approach uses a more complex neural network, a GravNet [4] with object condensation (OC) [5], to perform a full end-to-end reconstruction, i.e. starting with detector hits and generating reconstructed particles including their properties. The input to the network are calorimeter hits and reconstructed tracks with the properties position (calorimeter cell or track state at the calorimeter), time (activation of calorimeter cell, no time information for tracks) and deposited energy (in the calorimeter cell or track momentum). The target are particles from the MC truth record that created detector hits with their properties energy and position, which is chosen as the endpoint of the MC particle, as well as the correct association of detector hits and tracks to the corresponding reconstructed particle. The network generates an output object for each input object, i.e. for each calorimeter hit and track with the target properties energy and position. OC then introduces an additional learned property β that represents how similar the properties of an output object are. After inference, a threshold for β can be chosen (and optimised) above which the objects are then identified as predicted particles. To facilitate this, OC also adds a term to the loss function that encourages the network to generate exactly one object per true particle with a high β value and low β values for all other objects. Since these high- β objects are also used to generate a relative potential between all objects for clustering, they work as condensation points during training around which all objects belonging to the same true particle cluster. The position loss is the squared sum of the difference between predicted and true position in Cartesian coordinates, and the energy loss is the difference in predicted and true energy divided by the square root of the true energy to account for the typical behaviour of a calorimetric energy resolution.

The samples used for training and inference are e^+e^- to $q\bar{q}$ at 250 GeV, with 40k events for training and 4k events for inference. They contain ILC-like event conditions, including initial state radiation, beam energy spread and background particles




Figure 4 : Reconstructed event with end-to-end reconstruction in angular coordinates.




Figure 5 : Position reconstruction uncertainty of the end-to-end reconstruction. Left: no time information, middle: true time smeared by 100 ps, right: perfect (true) time.

from beam-beam interaction. To study the effect of timing information, samples with different levels of timing are compared: perfect timing, perfect timing with Gaussian smearing with a width of 100 ps, and no timing information. The result of the clustering can be seen in fig. 4. Here the positions of hits and particles are shown in angular coordinates in the detector. The true particles are denoted with green circles, with circle size increasing with particle energy. The predicted particles after selection via β are shown as green crosses, and they overlap fairly well with the true particles. The event has a typical structure: an initial state radiation (ISR) photon that caused the remaining e^+e^- system to return to the Z pole energy is emitted in the very backward direction, seen on the left. The remaining e^+e^- system goes into $q\bar{q}$ and forms two jets that go in the forward direction, recoiling from the ISR photon, and produce two major clusters in opposite phi directions.

The result of the position reconstruction can be seen in fig. 5. Here, the output position of the predicted particles is compared to the target position, i.e. the MC particle endpoint. Overall, the distributions are Gaussian-shaped and the position of

reconstructed particles is within about 5 cm of the true endpoint of the corresponding true particle. This is, however, more than the transversal granularity of both ECal ($5 \times 5 \text{ mm}^2$) and HCal ($3 \times 3 \text{ cm}^2$) and is also work in progress. There is a small impact of adding timing information visible: both perfect and smeared timing result in about 1 mm better position reconstruction compared to no timing, which is about three times the uncertainty of the width values themselves. It is expected that the relative impact of timing becomes more notable when the overall reconstruction is more precise.

5 Conclusion

The application of ML in event reconstruction for colliders is an active field with many ongoing developments. We showed progress towards energy and full particle reconstruction. While the energy reconstruction is close to the results of conventional particle flow, the end-to-end approach is working albeit not at a competitive performance, yet. The application of new models or adapted training or data pruning offer typical paths forward to improve on these results. The addition of timing information shows first notable effects, though still rather limited. With an improved overall performance it is expected that the effect of timing becomes more visible.

References

- [1] M. Thomson, Particle flow calorimetry and the PandoraPFA algorithm. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **611**(1), 25–40 (2009). <https://doi.org/10.1016/j.nima.2009.09.009>
- [2] CALO5D homepage. URL https://www.ipe.kit.edu/english/projects_CALO5D.php
- [3] The ILD Collaboration, International Large Detector: Interim Design Report (2020). [arXiv:2003.01116](https://arxiv.org/abs/2003.01116)
- [4] S.R. Qasim, J. Kieseler, Y. Iiyama, M. Pierini, Learning representations of irregular particle-detector geometry with distance-weighted graph networks. The European Physical Journal C **79**(7) (2019). <https://doi.org/10.1140/epjc/s10052-019-7113-9>
- [5] J. Kieseler, Object condensation: one-stage grid-free multi-object reconstruction in physics detectors, graph, and image data. The European Physical Journal C **80**(9) (2020). <https://doi.org/10.1140/epjc/s10052-020-08461-2>