

An updated hybrid deep learning algorithm for identifying and locating primary vertices

SIMON AKAR¹, THOMAS J. BOETTCHER², SARAH CARL¹, HENRY F. SCHREINER³,
MICHAEL D. SOKOLOFF¹, MARIAN STAHL¹, CONSTANTIN WEISSER², MIKE WILLIAMS²

On behalf of the LHCb Real Time Analysis project,

¹ *University of Cincinnati, Cincinnati, OH, United States*

² *Massachusetts Institute of Technology, Cambridge, MA, United States*

³ *Princeton University, Princeton, NJ, United States*

ABSTRACT

We present an improved hybrid algorithm for vertexing, that combines deep learning with conventional methods. Even though the algorithm is a generic approach to vertex finding, we focus here on it's application as an alternative

Primary Vertex (PV) finding tool for the LHCb experiment.

In the transition to Run 3 in 2021, LHCb will undergo a major luminosity upgrade, going from 1.1 to 5.6 expected visible PVs per event, and it will adopt a purely software trigger. We use a custom kernel to transform the sparse 3D space

of hits and tracks into a dense 1D dataset, and then apply Deep Learning techniques to find PV locations using proxy distributions to encode the truth in training data. Last year we reported that training networks on our kernels using several Convolutional Neural Network layers yielded better than 90 % efficiency

with no more than 0.2 False Positives (FPs) per event. Modifying several elements of the algorithm, we now achieve better than 94 % efficiency with a significantly lower FP rate. Where our studies to date have been made using toy Monte Carlo (MC), we began to study **kernel density histograms** produced from complete LHCb Run 3 MC data, including full tracking in the vertex locator rather than **standalone tracking**.

PRESENTED AT

Connecting the Dots Workshop (CTD 2020)
April 20-30, 2020

1 Introduction

The LHCb experiment is currently being upgraded for the planned start of Run 3 of the LHC in 2021 to facilitate recording proton-proton collision data at $\sqrt{s} = 14$ TeV with an instantaneous luminosity of $2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ [1]. Building on the success of Run 2 [2], the experiment will continue moving towards a real-time-analysis approach and consequently remove the Level 0 hardware trigger in favor of a pure software trigger system. In this process the the entire software stack is refactored, not only to reflect the new hardware components, but also to improve performance of reconstruction and selection. One of the algorithms that is largely affected by the increased instantaneous luminosity is Primary Vertex (PV) finding. The increased luminosity translates to an expected average number of PVs from 1.1 in Run 2 to 5.6 in Run 3, which will degrade the overall performance of PV reconstruction. A new fast PV finding algorithm has been brought forward and is currently proposed as the baseline solution [3].

As an alternative we propose a hybrid algorithm, established in Ref. [4], available on gitlab [5], and schematically shown in Fig. 1, that approaches this challenge with machine learning techniques in the cluster search of PV finding - the part of the algorithm that drives the PV reconstruction efficiency. In the context of PV finding at LHCb, the hybrid algorithm starts with tracks that have been reconstructed from hits in the Vertex locator (Velo). Their location, direction, and covariance matrix information is used to reduce the problem from three to one dimension, by calculating “kernels” in bins along z , *i.e.* the beam direction. The z positions of the primary vertices are closely related to peaks in this kernel histogram, and it is then used as input to a Convolutional Neural Network (CNN) to carry out the cluster search and predict the PV locations. The output probabilities from the CNN are converted to a list of PV candidate z -locations using a simple peak finding algorithm. A toy simulation of the LHCb detector is used to train, develop and test the algorithm. In parallel, the algorithm has been deployed into the LHCb software stack where it can be used for PV finding.

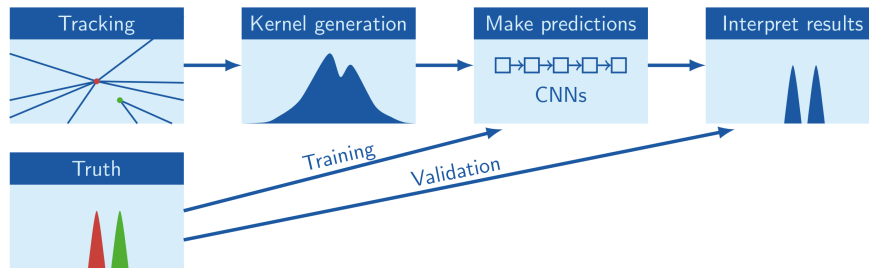


Figure 1: Schematic workflow of the hybrid deep learning algorithm for vertex finding.

2 Kernel generation

The kernel generation step converts sparse 3D data into a feature-rich 1D histogram. It consists of 4000 bins along the z -direction (beamline), each $100 \mu\text{m}$ wide, covering the active area of the Velo around the interaction point. Each z bin of the histogram is filled by the maximum kernel value in x and y , where the kernel - **growing larger the closer and the more track trajectories pass the beamline** - is defined by

$$\mathcal{K}(x, y, z) = \frac{\sum_{\text{tracks}} \mathcal{G}(\text{IP}(x, y)|z)^2}{\sum_{\text{tracks}} \mathcal{G}(\text{IP}(x, y)|z)} - \sum_{\text{tracks}} \mathcal{G}(\text{IP}(x, y)|z) . \quad (1)$$

In Eq. (1), $\mathcal{G}(\text{IP}(x, y)|z)$ is a Gaussian function, centered at $x = y = 0$ and evaluated at the impact parameter $\text{IP}(x, y)$: the distance of closest approach of a track projection to a hypothesized vertex at position x, y for a given z . The width/covariance of \mathcal{G} is given by the $\text{IP}(x, y)$ uncertainty/covariance matrix. Finding the maximum $\mathcal{K}(x, y, z)$ is a two step process, where kernel values are first computed in a coarse 8×8 grid in x, y ; then the parameters of that search are then taken as starting points for a MINUIT minimization process to find the maximum kernel value.

Equation (1) is computed from the location, direction, and covariance matrix of input tracks. This information can be provided by a standalone toy Monte Carlo proto-tracking [4] or from the proposed LHCb Run 3 production Velo tracking [6]. While the former uses a heuristic Gaussian width to estimate the $IP(x, y)$ uncertainty, the latter makes use of the measured covariance matrix of the track state closest to beam to compute the $IP(x, y)$ covariance. The difference between kernel histogram using the heuristic and measured $IP(x, y)$ uncertainty/covariance is exemplarily shown in Fig. 2. It can be seen that the kernels are more pronounced around PVs, which suggests that – once re-trained – the CNN performance should increase. Further improvement with LHCb simulation data is expected from centering \mathcal{G} at the actual beamline position at a given z position, and by using Kalman-fitted Velo tracks as input to the kernel generation.

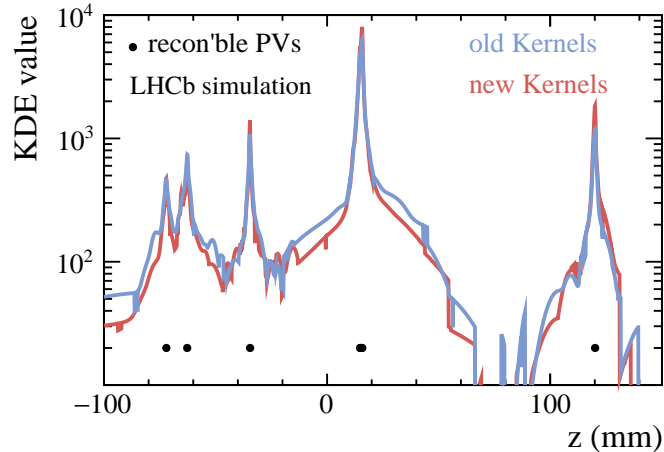


Figure 2: Comparison of kernel histograms heuristic (blue) and measured (red) $IP(x, y)$ uncertainty/covariance. A random single event with 6 reconstructible primary vertices is shown, whose z positions are marked by black dots. Here, a simulated primary vertex is called reconstructible if more than four simulated tracks within the Velo acceptance emerge from that vertex. Note that there are 2 PVs at about +15 mm, illustrating the issue of PV proximity that the cluster search faces frequently.

3 CNN based cluster search

The cluster search in our algorithm is conducted by a Convolutional Neural Network (CNN) consisting of several one dimensional convolution layers. We use `PyTorch` [7] for training and development, and `TorchScript` [8] as C++ inference engine in the LHCb software stack.

Before being able to train a CNN, we need to define what it should learn, *i.e.* give it a target. The CNN in our algorithm is designed to output a 4000 bin *target histogram*, just as the input kernel histogram **with** essentially all “noise” removed, **as sketched in Fig. 1**. A first approach to create such a target histogram is to define a Gaussian with unit area and a width of one bin ($100 \mu\text{m}$) centered around the simulated PV position. These simple target histograms have been updated to take the PV resolution as a function of track multiplicity into account, which affects width and area of the Gaussians. The resulting performance improvement is plotted in Fig. 3. Further, regions around PVs that do not pass the criterion of being reconstructible, *i.e.* have 5 or less detectable associated tracks, are “masked”. Masked regions are effectively hidden during training, so that discoveries in them are neither punished nor rewarded.

Our best performing network to date is composed of 6 convolutional layers with leaky ReLU activation functions in-between hidden layers and a softmax activation for the output layer. The widths and padding of each convolution kernel was chosen by visual inspection of the data; the number of channels were increased until benefits were no longer noticeable upon adding new channels. The training is carried out on GPUs using mini-batch gradient descent, the Adam optimizer and dropout regularization.

A custom cost function, similar to cross-entropy, has been defined and it was found that its initial symmetric form favored small false positive rates at the expense of efficiency. Therefore, a single parameter asymmetry term has been added to the cost function, serving as powerful control for selecting the false positive to efficiency tradeoff [4]. To stabilize the training process in early epochs, the last convolution layer is replaced by a Fully Connected (FC) layer. After several training epochs with this architecture, the weights of the convolutional layers are fixed, the FC layer is replaced by a convolutional layer, and the network is trained for a few more epochs. Then, all weights are floated and the CNN is trained in its final architecture.

Recently, we improved the CNN performance by adding x and y position information, found by maximizing the kernel (Eq. (1)), in a perturbative manner. Adding the information perturbatively is important because the CNN overestimates the importance of these variables compared to the z information. The perturbative addition works as follows: On the basis of the original network, another CNN with 3 convolutional layers that solely processes x and y position information is trained independently, but parallel to the original network. Such, that both CNN responses at the end of each training epoch are multiplied. The CNN with x and y information will contribute with values close to 1 in most cases, but can veto kernel peaks with high x, y gradients that have been observed in data and can lead to false positives. The performance improvement of the perturbation network, together with the addition of one convolutional layer is shown in Fig. 3.

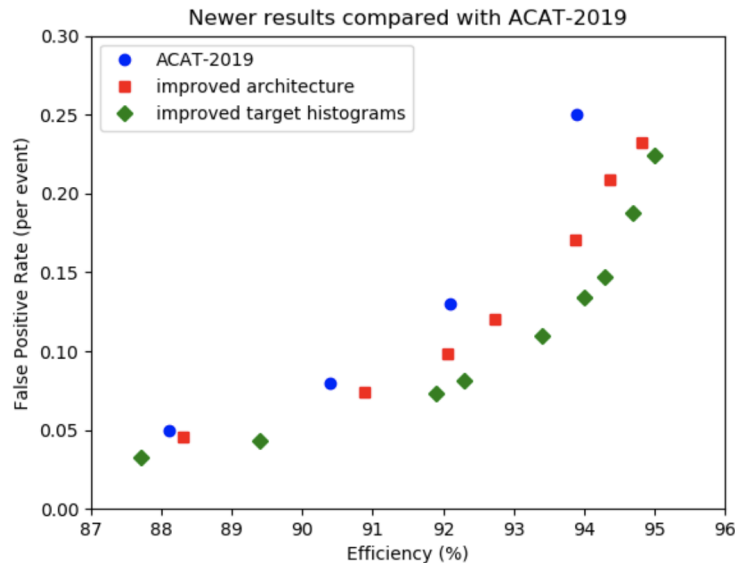


Figure 3: Performance improvements with respect to Ref. [4] achieved by modifying target histograms, adding layers to the CNN and adding x, y position information perturbatively. Details are described in the text.

4 Results

The proof of principle, that the here proposed hybrid deep learning algorithm is an efficient vertex finding tool has been established in Ref. [4]. Since then, the algorithm’s performance has been improved further by changes described in Sec. 3 and are shown in Fig. 3. We want to highlight that for a fixed PV finding efficiency of 94%, the false positive rate per event could be reduced by a factor of 2. These numbers have been obtained on toy simulation using metrics that do not fully reflect standard LHCb definitions, and are thus not directly comparable to the ones given in Ref. [3]. However, first studies using official LHCb simulation in place of the toy simulation with our metrics agree with those from pure toy simulation. We are looking forward to re-train the CNNs with LHCb simulation data and expect further improvements of the PV finding performance.

5 Conclusions and Outlook

We have presented a hybrid deep learning vertexing algorithm and its application as PV finding algorithm in the LHCb Run 3 software trigger. The algorithm has been privately deployed in the LHCb CPU software stack and its performance has been further improved using toy data.

Future milestones of the project are well defined. We plan to deploy the algorithm in Allen, the High-Level Trigger application on GPUs for LHCb [9]. Further, it is of great importance to have a one to one comparison with the currently proposed LHCb Run 3 production PV finding algorithm [3]. This means that we need to refactor the current implementation of our algorithm in the LHCb software stack to be able to run the same performance benchmark tests. In parallel, we need to re-train our CNN with official LHCb simulation data in place of toy data. We expect further performance improvements from doing so, in particular from using the measured covariance matrix of (Kalman-fitted) Velo tracks in the generation of kernel histograms. Currently, the generation of these kernel histograms is the throughput bottleneck of our algorithm due to the usage of MINUIT. We are planning to replace the kernel generation by a fast machine learning algorithm, which could be merged with the cluster search CNN into a single algorithm that predicts PV positions from the output of the Velo track reconstruction in one step. We are also investigating pruning techniques to speed up the throughput of the inference. Moreover, we plan to associate Velo tracks with found PV candidates probabilistically to reduce the false positive rate. In an adjacent step, we plan to probabilistically identify secondary vertices and associated tracks.

ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation under Cooperative Agreement OAC-1836650, OAC-1739772, and OAC-1740102. It was also supported by the University of Cincinnati Women in Science and Engineering program.

References

- [1] LHCb Collaboration, "LHCb Trigger and Online Upgrade Technical Design Report," CERN-LHCC-2014-016, LHCb-TDR-2014-016.
- [2] R. Aaij *et al.* [LHCb], "Design and performance of the LHCb trigger and full real-time reconstruction in Run 2 of the LHC," JINST **14**, no.04, P04013 (2019); [arXiv:1812.10790 [hep-ex]].
- [3] F. Reiß, "Fast parallel Primary Vertex reconstruction for the LHCb Upgrade," Talk given at this conference; LHCb-TALK-2020-044. Publication in preparation.
- [4] R. Fang, H. F. Schreiner, M. D. Sokoloff, C. Weisser and M. Williams, "A hybrid deep learning approach to vertexing," [arXiv:1906.08306 [physics.ins-det]].
- [5] `pv-finder` repository: <https://gitlab.cern.ch/LHCb-Reco-Dev/pv-finder>.
- [6] A. Hennequin, B. Couturier, V. Gligorov, S. Ponce, R. Quagliani and L. Lacassagne, "A fast and efficient SIMD track reconstruction algorithm for the LHCb Upgrade 1 VELO-PIX detector," [arXiv:1912.09901 [physics.ins-det]].
- [7] A. Paszke *et al.* [The PyTorch team], "PyTorch: An Imperative Style, High-Performance Deep Learning Library," Advances in Neural Information Processing Systems 32, 8024 (2019).
- [8] The PyTorch team, "Torch Script," <https://pytorch.org/docs/stable/jit.html>.
- [9] R. Aaij *et al.*, "Allen: A High-Level Trigger on GPUs for LHCb," Comput. Softw. Big Sci. **4**, no.1, 7 (2020); [arXiv:1912.09161 [physics.ins-det]].