
Deep Learning-Based Spatiotemporal Multi-Event Reconstruction for Delay-Line Detectors

Marco Knipfer¹ Stefan Meier² Jonas Heimerl² Peter Hommelhoff² Sergei Gleyzer¹

¹University of Alabama (UA), Department of Physics and Astronomy

²Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Physics Department

{mknipfer@crimson., sglezyer@}ua.edu

{stefan.m.meier, jonas.heimerl, peter.hommelhoff}@fau.de

Abstract

Accurate observation of two or more particles within a very narrow time window has always been a great challenge in modern physics. It opens the possibility for correlation experiments, as e.g. the important Hanbury Brown-Twiss experiment, leading to new physical insights. For low-energy electrons, one possibility is to use a micro-channel plate with subsequent delay-lines for the readout of the incident particle hits. With such a Delay-Line Detector the spatial and temporal coordinates of more than one particle can be fully reconstructed as soon as both particles have a larger separation than what is called the dead radius. For events where two electrons are closer in space and time, the determination of the individual positions of the particles requires elaborated peak finding algorithms. While classical methods work well with single particle hits, they fail to identify and reconstruct events caused by multiple particles when they arrive close in space and time. To address this challenge, a new spatiotemporal machine learning model is developed to identify and reconstruct the position and time of such multi-hit signals. The model achieves a much better resolution for near-by particle hits compared to the classical approaches, reducing the dead radius by half. This shows that machine learning models can be effective in improving the spatiotemporal performance of Delay-Line Detectors.

1 Introduction

The detection of single particles, such as atoms, ions, electrons and highly energetic photons is the cornerstone of many fields in fundamental physics research. Typically, the signal of a single particle is too faint to be measured directly and requires amplification. This is usually achieved with electron avalanches, as for example in Photon Multiplier Tubes (PMTs) [1] and Multi-Channel Plates (MCPs) [2]. MCPs offer the advantage that the incident particle can be spatially resolved with a phosphor screen or a wire-grid. The latter system is called a Delay-Line Detector (DLD) [3], a sophisticated spatially and temporally resolving detector schematically shown in Figure 1. When one or more electrons hit the MCP, they get multiplied and the resulting electron shower hits the wire grid consisting of three identical layers that are rotated by 120°. The signals from the wires (six channels) and the MCP (one channel) get amplified and an analog-to-digital converter (ADC) discretizes the signals with a bin size of 0.8 ns. By determining the peak positions in the data, one can calculate the position and time, (x, y, t) , for each incoming particle. DLDs can be used for ions, electrons and photons with energies high enough to emit an electron from the front side of the MCP. In the field of ultra-fast atomic physics, such detectors are e.g. used for the so called reaction microscopes (Cold Target Recoil Ion Momentum Spectroscopy *COLTRIMS*) [4], a technique that allowed to resolve the correlation of two simultaneously emitted electrons in non-sequential double ionization (NSDI) [5]. In another seminal experiment, a Delay-Line Detector was used to compare the bunching

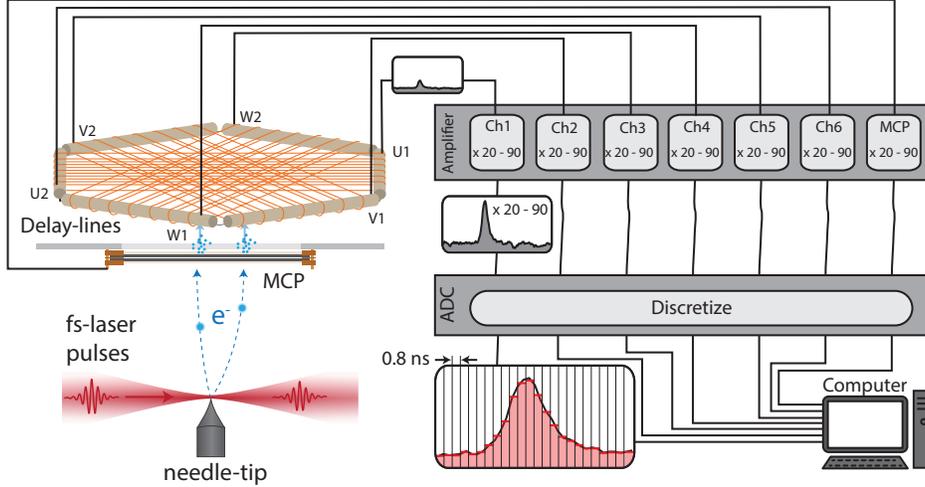


Figure 1: Schematic display of the experimental setup.

and antibunching behavior of free bosonic particles versus free fermionic particles [6]. There is redundant information for each incident particle provided by three layers of delay lines. This helps to disentangle the signals of two or more incident particles that hit the detector simultaneously. However, despite the additional information there still exist limitations regarding the multi-hit capability of these detectors. If the two particles arrive close in both location and time, their signals overlap and cannot be easily assigned to the individual particles.

Machine learning has seen an explosive use in physics, including applications of classification, regression, anomaly detection, generative modeling and others. Most of these studies exclusively focus either on the spatial or on the temporal domain. Recently spatiotemporal studies have started gaining traction in physics and other fields [7–18].

In our work, we focus on challenging spatiotemporal reconstruction scenarios where individual DLD signals are close in both space and time. We show that deep learning models can significantly enhance the multi-hit capability of Delay-Line Detectors compared to classical reconstruction methods.

2 Machine learning approach

The ML approach consists of a *Hit Multiplicity Classifier* (HMC) and *Deep Peak Finder* (DPF) models. The HMC works with normalized input and consists of a convolutional neural network (CNN) for each channel except the MCP. The outputs of the CNNs are concatenated and a final dense network with softmax activation for the last layer outputs probabilities for the four classes consisting of one to four particles. For inference, events that are classified as doubles are sent to the DPF where the peak positions for each channel are predicted separately. For the DPF we focus on double-hit events only and leave higher-multiplicity events to future work. The DPF consists of two bidirectional GRU (Gated Recurrent Unit) layers [19] followed by a dense network. Because of differences in the data, there are two models, one for the channels except MCP and one for the MCP. In the following we will call them DPF and DPF (MCP).

The signals are approximately additive and so it is possible to simulate multi-hit data for training. First, a classical peak finding algorithm is used to find the peaks in the real data. Then, the data is filtered to only keep single-hit events. Events with more than one peak in a channel are treated as multi-hit events. To simulate a double-hit event, two single-hit events are randomly shifted and added. Similarly for triple- and quadruple-hit events. Note that the peak positions in the simulated events are known to the same accuracy as for the single-hit events.

In training, the hyperparameters for the HMC and the DPF are optimized on a smaller dataset using Keras Tuner [20] and the Hyperband optimizer [21]. The HMC is trained for 300 epochs on ca. 202k simulated events split evenly into single/double/triple/quadruple using early stopping, model

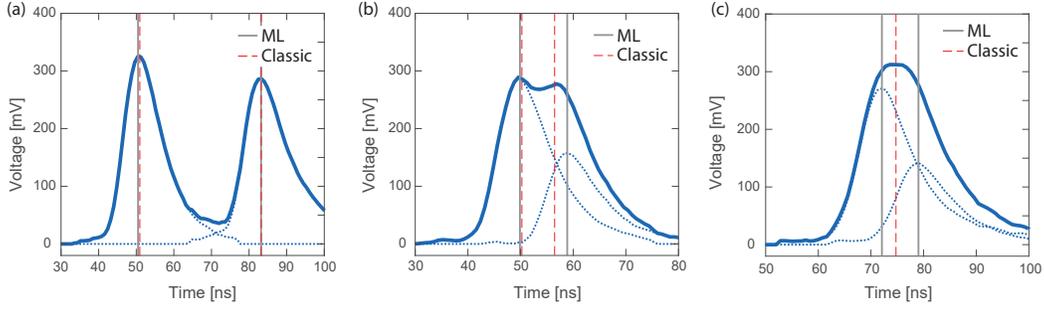


Figure 2: Examples of simulated double events. Dotted blue: two single peaks, thick blue: sum of the singles. Vertical lines: dashed: fit-based classical method, filled: Deep Peak Finder.

checkpoints and reduction of the learning rate on plateau. The DPF and DPF (MCP) are trained on 180k and 100k simulated doubles respectively for 300 epochs. All models are trained using the ADAM [22] optimizer. Training took around 3 h on Google Colab Pro Plus per model.

3 Results

Based on ca. 106k evenly distributed test events, the HMC model has a test accuracy of $acc = 0.9951$. The AUC is > 0.9998 for every class (one-vs-all). Note that the HMC works on normalized data to avoid learning the higher amplitudes due to the data being linear combinations of real data. On training data the root mean square error (RMSE) for the DPF is $rmse_{train} = 1.03$ ns and the mean average error (MAE) is $mae_{train} = 0.25$ ns. Evaluated on 1.2M simulated double-hit events, the RMSE for the peak position $rmse_{test} = 1.72$ ns and the MAE is $mae_{test} = 0.24$. For comparison, the bin size of the data is 0.8 ns. For the DPF (MCP) the results are similar: On 100k simulated MCP double peak signals: $rmse_{train} = 0.90$ ns and $rmse_{test} = 0.83$ ns, $mae_{train} = 0.27$ ns and $mae_{test} = 0.27$ ns. In Figures 2 and 3 comparisons between the classical methods and the DPF are shown. The first shows three simulated double peaks together with their underlying singles. The vertical lines show the classical fit-based prediction and the prediction from the DPF. The fit-based method is what was used to create the labels for the singles. It consists of quadratic fits around the maximal values. Figure 3 shows the RMSE as a function of peak separation for different methods. The first method is a simulation of the the hardware-based Constant-Fraction-Discriminator (CFD) which adds an inverted and shifted copy of the signal to itself. The resulting zero-crossing then indicates the peak. The second method is the above mentioned fit-based technique. The third is the Mean Pulse Curve Fit (MPCF) method. First, the mean pulse from 10k single-hit signals is calculated. Then, for doubles, the sum of two mean pulses is fitted to the data. The fit parameters are the two peak positions, amplitudes and scalings in the time direction.

In order to test the HMC, DPF and DPF (MCP) on real data, we have performed a separate experiment. We put a grid made of copper in front of the electron source. The shadow of the grid is projected onto the DLD as shown in Figure 4 (a). How clear the grid appears in the reconstructed positions is an indicator for the quality of the method. If the method is inaccurate, then the grid will look washed-out, because events will appear at positions where none should be. This will become more severe if we look at the distances between the two hits Δx and Δy : On the right side (b,c,d,e) of Figure 4 the position differences are shown as 2D histogram. Figure 4 (b), shows only differences between uncorrelated single-hit events classically evaluated and a good resolution of the grid. This can be taken as a reference for the best possible resolution. The second plot, (c), shows double-hit events evaluated with the CFD-method. There is a star-shaped dead region with a maximal extension of about 30 mm where most events cannot be reconstructed. There are also weak 6-folded artifacts. The third plot, (d), shows double-hit events evaluated with the classical peak finding algorithm. There are very strong 6-folded artifacts and the dead radius (white region) is at about 20 mm. The fourth plot, (e), shows the same double-hit events evaluated with the DPF. While the 6-folded artifacts are still visible, they are much weaker and almost gone. The dead radius is smaller at about 10 mm.

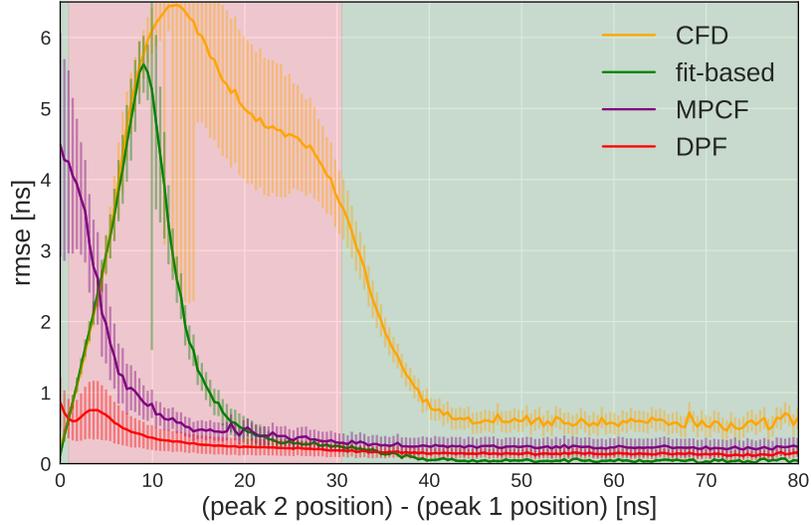


Figure 3: Root mean squared error as a function of the peak separation for the channels (not MCP). For peak positions smaller than 15–20 ns, which is around the width of a typical peak, the fit-based and Mean Pulse Curve Fit (MPCF) methods decrease drastically in performance, while the Deep Peak Finder (DPF) stays good.

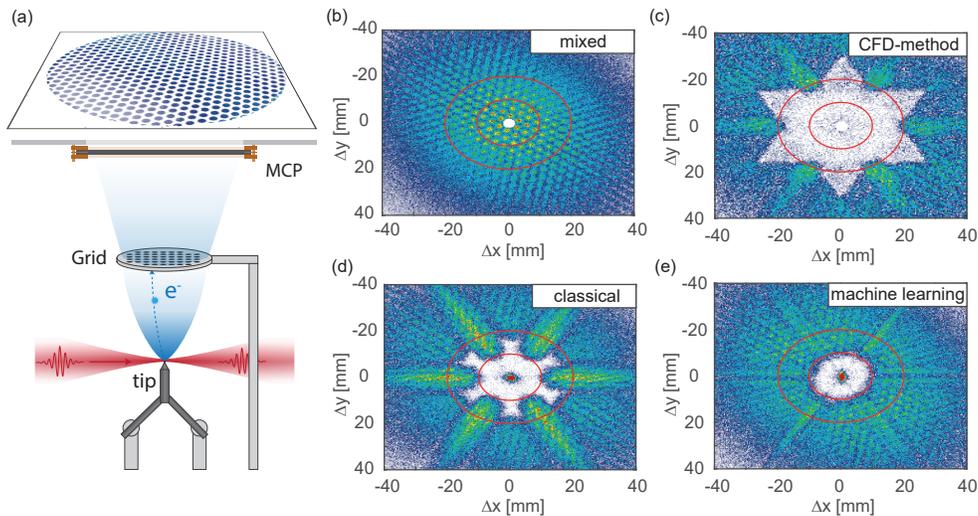


Figure 4: A grid is put in front of the MCP and is projected onto the detector (a). The plots (b-e) show histograms of the position differences between particle pairs calculated with the different methods.

4 Interpretation and conclusion

The Hit Multiplicity Classifier works well with an exceptional AUC of > 0.9998 . Figure 2 shows clearly how the Deep Peak Finder finds the correct peaks even when they are close. From Figure 2 one can see that a single peak typically has an extension of ca. 20-30 ns. Thus, for peak separations larger than this order we expect the peaks not to influence each other and no disturbance of the peak position determination happens. Indeed, for > 30 ns separations, in Figure 3 the fit-based method has almost zero RMSE because this method is what was used to create the labels. Note that in this region the DPF still performs well with an RMSE about 1/4 of the bin width which is 0.8 ns. For smaller separations however, one would expect the peaks to influence each other and disturb the peak position determination. This is confirmed in Figure 3 where for about 0.9–25 ns peak separation the

DPF clearly wins over the classical peak finder in terms of RMSE and in Figure 2 in the middle and right plot the DPF finds the underlying peaks better. Once the peaks get too close, the classical peak finder only finds one peak which is then counted as two in our evaluation. Thus, for the case where the peaks are < 0.9 ns apart, the classical wins again in 3. Note that here the classical is actually overrated, because we give it the information that there have to be two peaks. The grid plots in Figure 4 confirm that our model also performs well on real data where no ground truth exists. Improvements over the classical methods are 1) smaller dead radius of ca. 10 mm as opposed to the 20 mm of the fit-based or 30 mm of the CFD method, 2) less artifacts and 3) a better resolution of the grid for smaller separations. Note that these absolute numbers for the dead radius depend on the initial width of the peaks, thus all methods will give smaller dead radii if the peak width gets reduced due to hardware changes.

As this evaluation is applicable for any existing delay-line-based detector with the possibility of analog read-out, it is highly relevant for many ultrafast pulsed applications like non-sequential double ionization, atom probe tomography, ultrafast electron microscopy or other correlation experiments like the fermionic Hanbury Brown-Twiss experiment [23, 24]. Further plans include end-to-end prediction of the particle positions and times instead of only predicting the peak positions, implementation of the models with SOFIE [25] such that they can be used on live data, and other model architectures like transformers [26]. Currently, there are no on-going experimental studies or correlation experiments that take advantage of triplet and quadruplet events. However, that may change in the future. Therefore, we leave higher multiplicity event reconstruction to future work.

Potential broader impact of this work

This work contributes to the analysis of data generated by a specific kind of particle detectors called Delay-Line Detectors. Combined with correlation experiments, this will hopefully one day lead to better understanding of the fundamental interactions between particles and ultimately to a better understanding of the universe. To the authors' knowledge there are no direct ethical aspects to consider. Indirect ethical aspects are difficult to grasp, but to the authors' knowledge no obvious indirect ethical aspects are to consider in the near future. Future societal consequences are the same as for all fundamental physics research and almost impossible to foresee.

Acknowledgments and Disclosure of Funding

S.M., J.H. and P.H. acknowledge support by the European Research Council (Consolidator Grant NearFieldAtto and Advanced Grant AccelOnChip) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project-ID 429529648—TRR 306 QuCoLiMa (“Quantum Cooperativity of Light and Matter”) and Sonderforschungsbereich 953 (“Synthetic Carbon Allotropes”), Project-ID 182849149. J.H. acknowledges funding from the Max Planck School of Photonics. This work was supported in part by the U.S. Department of Energy (DOE) under Award No. DE-SC0012447 (S.G.). M.K. acknowledges support through the Graduate Council Fellowship of the University of Alabama.

References

- [1] B.K. Lubsandorzhev. “On the history of photomultiplier tube invention”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 567.1 (Nov. 2006), pp. 236–238. DOI: 10.1016/j.nima.2006.05.221. URL: <https://doi.org/10.1016/j.nima.2006.05.221>.
- [2] T. Gys. “Micro-channel plates and vacuum detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 787 (July 2015), pp. 254–260. DOI: 10.1016/j.nima.2014.12.044. URL: <https://doi.org/10.1016/j.nima.2014.12.044>.
- [3] O. Jagutzki et al. “A broad-application microchannel-plate detector system for advanced particle or photon detection tasks: large area imaging, precise multi-hit timing information and high detection rate”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 477.1-3 (Jan. 2002),

- pp. 244–249. DOI: 10.1016/s0168-9002(01)01839-3. URL: [https://doi.org/10.1016/s0168-9002\(01\)01839-3](https://doi.org/10.1016/s0168-9002(01)01839-3).
- [4] J. Ullrich et al. “Recoil-ion and electron momentum spectroscopy: reaction-microscopes”. In: *Reports on Progress in Physics* 66.9 (Aug. 2003), pp. 1463–1545. DOI: 10.1088/0034-4885/66/9/203. URL: <https://doi.org/10.1088/0034-4885/66/9/203>.
- [5] Th. Weber et al. “Correlated electron emission in multiphoton double ionization”. In: *Nature* 405.6787 (June 2000), pp. 658–661. DOI: 10.1038/35015033. URL: <https://doi.org/10.1038/35015033>.
- [6] T. Jelte et al. “Comparison of the Hanbury Brown–Twiss effect for bosons and fermions”. In: *Nature* 445.7126 (Jan. 2007), pp. 402–405. DOI: 10.1038/nature05513. URL: <https://doi.org/10.1038/nature05513>.
- [7] Rui Wang et al. *Bridging Physics-based and Data-driven modeling for Learning Dynamical Systems*. 2020. DOI: 10.48550/ARXIV.2011.10616. URL: <https://arxiv.org/abs/2011.10616>.
- [8] Leye Wang et al. “Cross-City Transfer Learning for Deep Spatio-Temporal Prediction”. In: *arXiv e-prints*, arXiv:1802.00386 (Feb. 2018), arXiv:1802.00386. arXiv: 1802.00386 [cs.AI].
- [9] Rui Wang et al. “Bridging Physics-based and Data-driven modeling for Learning Dynamical Systems”. In: *arXiv e-prints*, arXiv:2011.10616 (Nov. 2020), arXiv:2011.10616. arXiv: 2011.10616 [cs.LG].
- [10] Thomas P. Karnowski. “Deep Machine Learning with Spatio-Temporal Inference”. PhD thesis. University of Tennessee, 2012. URL: https://trace.tennessee.edu/utk_graddiss/1315/.
- [11] Rui Wang et al. “Towards Physics-Informed Deep Learning for Turbulent Flow Prediction”. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '20*. Virtual Event, CA, USA: Association for Computing Machinery, 2020, pp. 1457–1466. ISBN: 9781450379984. DOI: 10.1145/3394486.3403198. URL: <https://doi.org/10.1145/3394486.3403198>.
- [12] Guanya Shi et al. “Neural Lander: Stable Drone Landing Control Using Learned Dynamics”. In: *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, May 2019. DOI: 10.1109/icra.2019.8794351. URL: <https://doi.org/10.1109%5C%2Ficra.2019.8794351>.
- [13] Dongxia Wu et al. *DeepGLEAM: A hybrid mechanistic and deep learning model for COVID-19 forecasting*. 2021. DOI: 10.48550/ARXIV.2102.06684. URL: <https://arxiv.org/abs/2102.06684>.
- [14] Zihao Zhou et al. “Neural Point Process for Learning Spatiotemporal Event Dynamics”. In: *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*. Ed. by Roya Firoozi et al. Vol. 168. Proceedings of Machine Learning Research. PMLR, June 2022, pp. 777–789. URL: <https://proceedings.mlr.press/v168/zhou22a.html>.
- [15] Bin Sun and Yuehua Wu. “Estimation of the Covariance Matrix in Hierarchical Bayesian Spatio-Temporal Modeling via Dimension Expansion”. In: *Entropy* 24.4 (2022). ISSN: 1099-4300. DOI: 10.3390/e24040492. URL: <https://www.mdpi.com/1099-4300/24/4/492>.
- [16] Vasiliki D. Agou, Andrew Pavlides, and Dionissios T. Hristopoulos. “Spatial Modeling of Precipitation Based on Data-Driven Warping of Gaussian Processes”. In: *Entropy* 24.3 (2022). ISSN: 1099-4300. DOI: 10.3390/e24030321. URL: <https://www.mdpi.com/1099-4300/24/3/321>.
- [17] Ashesh Chattopadhyay, Pedram Hassanzadeh, and Saba Pasha. “Predicting clustered weather patterns: A test case for applications of convolutional neural networks to spatio-temporal climate data”. In: *Scientific Reports* 10, 1317 (Jan. 2020), p. 1317. DOI: 10.1038/s41598-020-57897-9. arXiv: 1811.04817 [physics.ao-ph].
- [18] Bao Wang et al. “Graph-Based Deep Modeling and Real Time Forecasting of Sparse Spatio-Temporal Data”. In: *arXiv e-prints*, arXiv:1804.00684 (Apr. 2018), arXiv:1804.00684. arXiv: 1804.00684 [cs.LG].
- [19] Kyunghyun Cho et al. “On the Properties of Neural Machine Translation: Encoder-Decoder Approaches”. In: *arXiv e-prints*, arXiv:1409.1259 (Sept. 2014), arXiv:1409.1259. arXiv: 1409.1259 [cs.CL].

- [20] Tom O'Malley et al. *Keras Tuner*. <https://github.com/keras-team/keras-tuner>. 2019.
- [21] Lisha Li et al. "Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization". In: *arXiv e-prints*, arXiv:1603.06560 (Mar. 2016), arXiv:1603.06560. arXiv: 1603.06560 [cs.LG].
- [22] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2014. DOI: 10.48550/ARXIV.1412.6980. URL: <https://arxiv.org/abs/1412.6980>.
- [23] Harald Kiesel, Andreas Renz, and Franz Hasselbach. "Observation of Hanbury Brown–Twiss anticorrelations for free electrons". In: *Nature* 418.6896 (July 2002), pp. 392–394. DOI: 10.1038/nature00911. URL: <http://dx.doi.org/10.1038/nature00911>.
- [24] Makoto Kuwahara et al. "Intensity Interference in a Coherent Spin-Polarized Electron Beam". In: *Physical Review Letters* 126.12 (Mar. 2021), p. 125501. DOI: 10.1103/physrevlett.126.125501. URL: <https://doi.org/10.1103/physrevlett.126.125501>.
- [25] An, Sitong and Moneta, Lorenzo. "C++ Code Generation for Fast Inference of Deep Learning Models in ROOT/TMVA". In: *EPJ Web Conf.* 251 (2021), p. 03040. DOI: 10.1051/epjconf/202125103040. URL: <https://doi.org/10.1051/epjconf/202125103040>.
- [26] Ashish Vaswani et al. "Attention Is All You Need". In: *arXiv e-prints*, arXiv:1706.03762 (June 2017), arXiv:1706.03762. arXiv: 1706.03762 [cs.CL].

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] see section 4
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] in the section on "Potential broader impact of this work"
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
 - (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] The code and the data are proprietary.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [No] We mention most of it, but not the exact hyperparameters. Not enough space.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We show the RMSE in Figure 3 together with error bars (one standard deviation) and an evaluation on real data in Figure 4.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] In section 2.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A] No crowdsourcing.
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]