

Raul Richard Stein

University of Bristol

Research area: Machine learning in particle physics detectors

Project title: Data acquisition with low-latency machine learning for neutrino physics in DUNE



About me:

I am 23 years old and come from Estonia where I completed my primary and secondary education graduating with a gold medal and an invitation to the presidential reception for top students.

I graduated from the University of York with a 1st class honours (with distinction) Physics BSc and received an award for the best overall performance. The following summer I did a summer studentship at CERN where I was first exposed to machine learning and furthered my interest in particle physics. Since then I have been more interested in data science and the wider potential of AI technology.

Following that, I explored data science in the context of fintech working in the Estonian branch of Bigbank from October 2018. I quickly realised that I would like to deepen my understanding AI & ML and decided to start the Data Intensive CDT and PhD in Particle Physics at the University of Bristol in September, 2019.

Outside of academia, I am a fitness enthusiast, and I enjoy break dancing and weightlifting. I am also interested in learning about investing and blockchain technology and related AI/ML applications in fintech, specifically predictive modelling, and portfolio optimization.

Please see my LinkedIn for more information: <https://www.linkedin.com/in/rrstein>

Data Intensive / Artificial Intelligence /Machine Learning / Advanced Computing Skills and interests:

I am mainly working in a **unix environment** on different computing clusters and using **Python with Anaconda** and relevant modules for my work. For data analysis, I regularly use **pandas, numpy, matplotlib** and for ML I have experience with **Keras + Tensorflow, XGBoost, and scikit-learn**. I have also taken an intensive **C++ programming** course as part of my studies but not yet used extensively in my work. In previous positions I have used **R for data analysis and model development** as well as MS Excel, ROOT with TMVA. I am also confident in using **(Postgre) SQL** for relational database queries.

During CDT studies I have completed several courses developing skills in **OOP, version control with Git, and code testing** as well as learning data science & visualisation concepts such as **Monte Carlo methods, Bayesian statistics & optimisation, Gaussian Processes**, etc. I have done **analysis and predictive modelling on time-series data** using an LSTM network and BDT for predicting oil prices as part of coursework. I have also limited experience with **cloud computing**, specifically using **AWS EC2** service and **docker containers** for software deployment.

I am interested in furthering my proficiency with the ML frameworks Tensorflow, Keras, and expanding to PyTorch. I also want to learn more about hardware implementations using C++ and high level synthesis for FPGAs or CUDA for GPUs. In terms of AI & ML concepts I would not exclude anything at this stage but at the moment I am most interested in the following topics: computer vision and autonomous vehicles, reinforcement learning, bayesian optimisation, and stream learning.

Science/ Research information:

I am investigating ways to accelerate machine learning models and analysis on specialised hardware for triggering applications in particle physics detectors by taking advantage of novel AI developments and the large amount of data within these environments.

I am working as part of the DUNE far detector (FD) data acquisition (DAQ) group. DUNE is a next generation neutrino experiment seeking to explain many fundamental physics questions such as matter-antimatter asymmetry, black hole formation from supernova, and proton decay. It will consist of two neutrino detectors placed in the most intense neutrino beam with one detector near the source at Fermilab, Illinois and the second, FD, 1300 km away, 1.5km underground at SURF, South Dakota. This setup and the liquid argon time projection chamber technology will allow for accurate measurements of neutrino oscillations and new subatomic phenomena.

So far, I have worked on the data transfer firmware validation in support of developing a FPGA based level 1 trigger (L1T). Currently, I am looking at using YOLO networks for object detection to identify specific neutrino events that we would like to trigger data selection on in L1T. This means we are looking at rather rough, unprocessed data with a low-latency requirement for inference (order of a few ms). If we discover a performance increase then we will investigate accelerating the network inference using an FPGA. YOLO networks are promising for this as realtime inference on common object detection (cars, animals, etc.) has been shown on CPUs and around 300 FPS has been achieved using FPGAs. We will also look into other network architectures such as Sparse Conv Nets or LSTMs for performance comparison. Furthermore, there is a possibility of investigating the potential of GPUs or novel intelligence processing units for online analysis at other stages of the trigger system.

Future goals and desires:

In the following years I wish to complete my PhD and make a useful contribution towards improving the DAQ and trigger system of particle physics detectors, specifically the DUNE FD. Personally, I hope to develop a broad understanding of data science techniques including AI & ML as well as proficiency in the practical application of these. This involves developing several skills aside from theoretical knowledge.

I believe an industrial placement will be a useful experience for me to get exposure to different types of problems, applications, and technology in the field. Post PhD I wish to pursue a career in industry and work on a project that has tangible and impactful outputs while making use of the data science skills I have developed and continuous development on these.