Abstract Submitted
for the DNP20 Meeting of
The American Physical Society

Benchmarking New Hardware For Machine Learning In Neutrino Physics STEFANO VERGANI, University of Cambridge — Over the last ten years, the popularity of Machine Learning (ML) has grown exponentially in all scientific fields, included particle physics. The amount of data and its complexity has grown as well, and the computing power required to perform inference can nowadays hardly be managed by the existing technology. Central Processing Units (CPUs) are generally affordable and ready to use but their ability to run Artificial Intelligence (AI) is very limited. In recent years, Graphics Processing Units (GPUs) have started to be used with very good results but they expensive, require a lot of power, and they are difficult to program since they were not invented for this task. Recently, Google has produced a brand new Edge Tensor Processing Unit (TPU) made explicitly to perform inference. It is cheap, it consumes less power than a GPU, and it comes with the portable size of a USB-key. A generic Liquid Argon Time-Projection Chamber (LArTPC) has been simulated and images produced by fictitious neutrino interactions have been used to benchmark the Edge TPU. Several popular Deep Learning (DL) models have been trained with those images using TensorFlow software and the performance of the Edge TPU during inference has been tested and compared with CPUs and GPUs.

Stefano Vergani
University of Cambridge

Date submitted: 29 Jun 2020 Electronic form version 1.4