UNIVERSITY OF CAMBRIDGE

Computer Vision in Neutrino Physics

Dr Leigh Whitehead

University of Cambridge

9th August 2023

51st SLAC Summer Institute (SSI 2023)

About Me

- Senior Research Associate at the University of Cambridge
  - PhD from the University of Warwick working on T2K
  - Post-doc on MINOS / MINOS+ at UCL
  - Research Fellow at CERN working on DUNE / ProtoDUNE

- Convene the ProtoDUNE sim / reco / physics analysis group
- Most of my career has been spent working on event reconstruction software
  - Including traditional techniques, neural networks, BDTs, etc
- Focused on Deep Learning since 2017
  - CNNs, GANs, GNNs… etc

Outline

- CNNs for event classification
  - The first CNNs in neutrino physics: NOvA and MicroBooNE
  - DUNE neutrino event classification*
- CNNs for pixel classification
  - Semantic segmentation in MicroBooNE, and ProtoDUNE-SP small-patch CNN*
- Black-box and bias concerns
  - Occlusion tests in DUNE, and the MINERvA DANN approach
- Transfer Learning*

* NB: These are examples from work that I have done

Introduction

- The title of the talk contains the phrase *computer vision*
  - In this talk I will be solely focussing on convolutional neural networks (CNNs)
- CNNs have been leading algorithms for image-based challenges for many years
- I can't cover everything here!
  - I give a few other references in the final slide of the talk
  - Any omission certainly isn't an indication of the quality of work

Neutrino physics detectors

- Many neutrino physics detectors lend themselves to image recognition
  - Lots of 2D readout systems naturally produce image-like data
  - Scintillator tracking detectors usually have 2 x 2D images per event
  - LArTPCs typically have 3 x 2D images per event
- Interactions typically happen uniformly within the detector volume
- Complexity of interactions varies significantly over a fairly small range of energy
  - See multiple track-like and shower-like components in each event

CNNs for Event Classification

The First CNNs: NOvA

- NOvA is a long-baseline neutrino oscillation experiment
  - Far detector is made from bars filled with liquid scintillator giving 2 x 2D readout
  - Detects neutrinos from Fermilab's NuMI beam

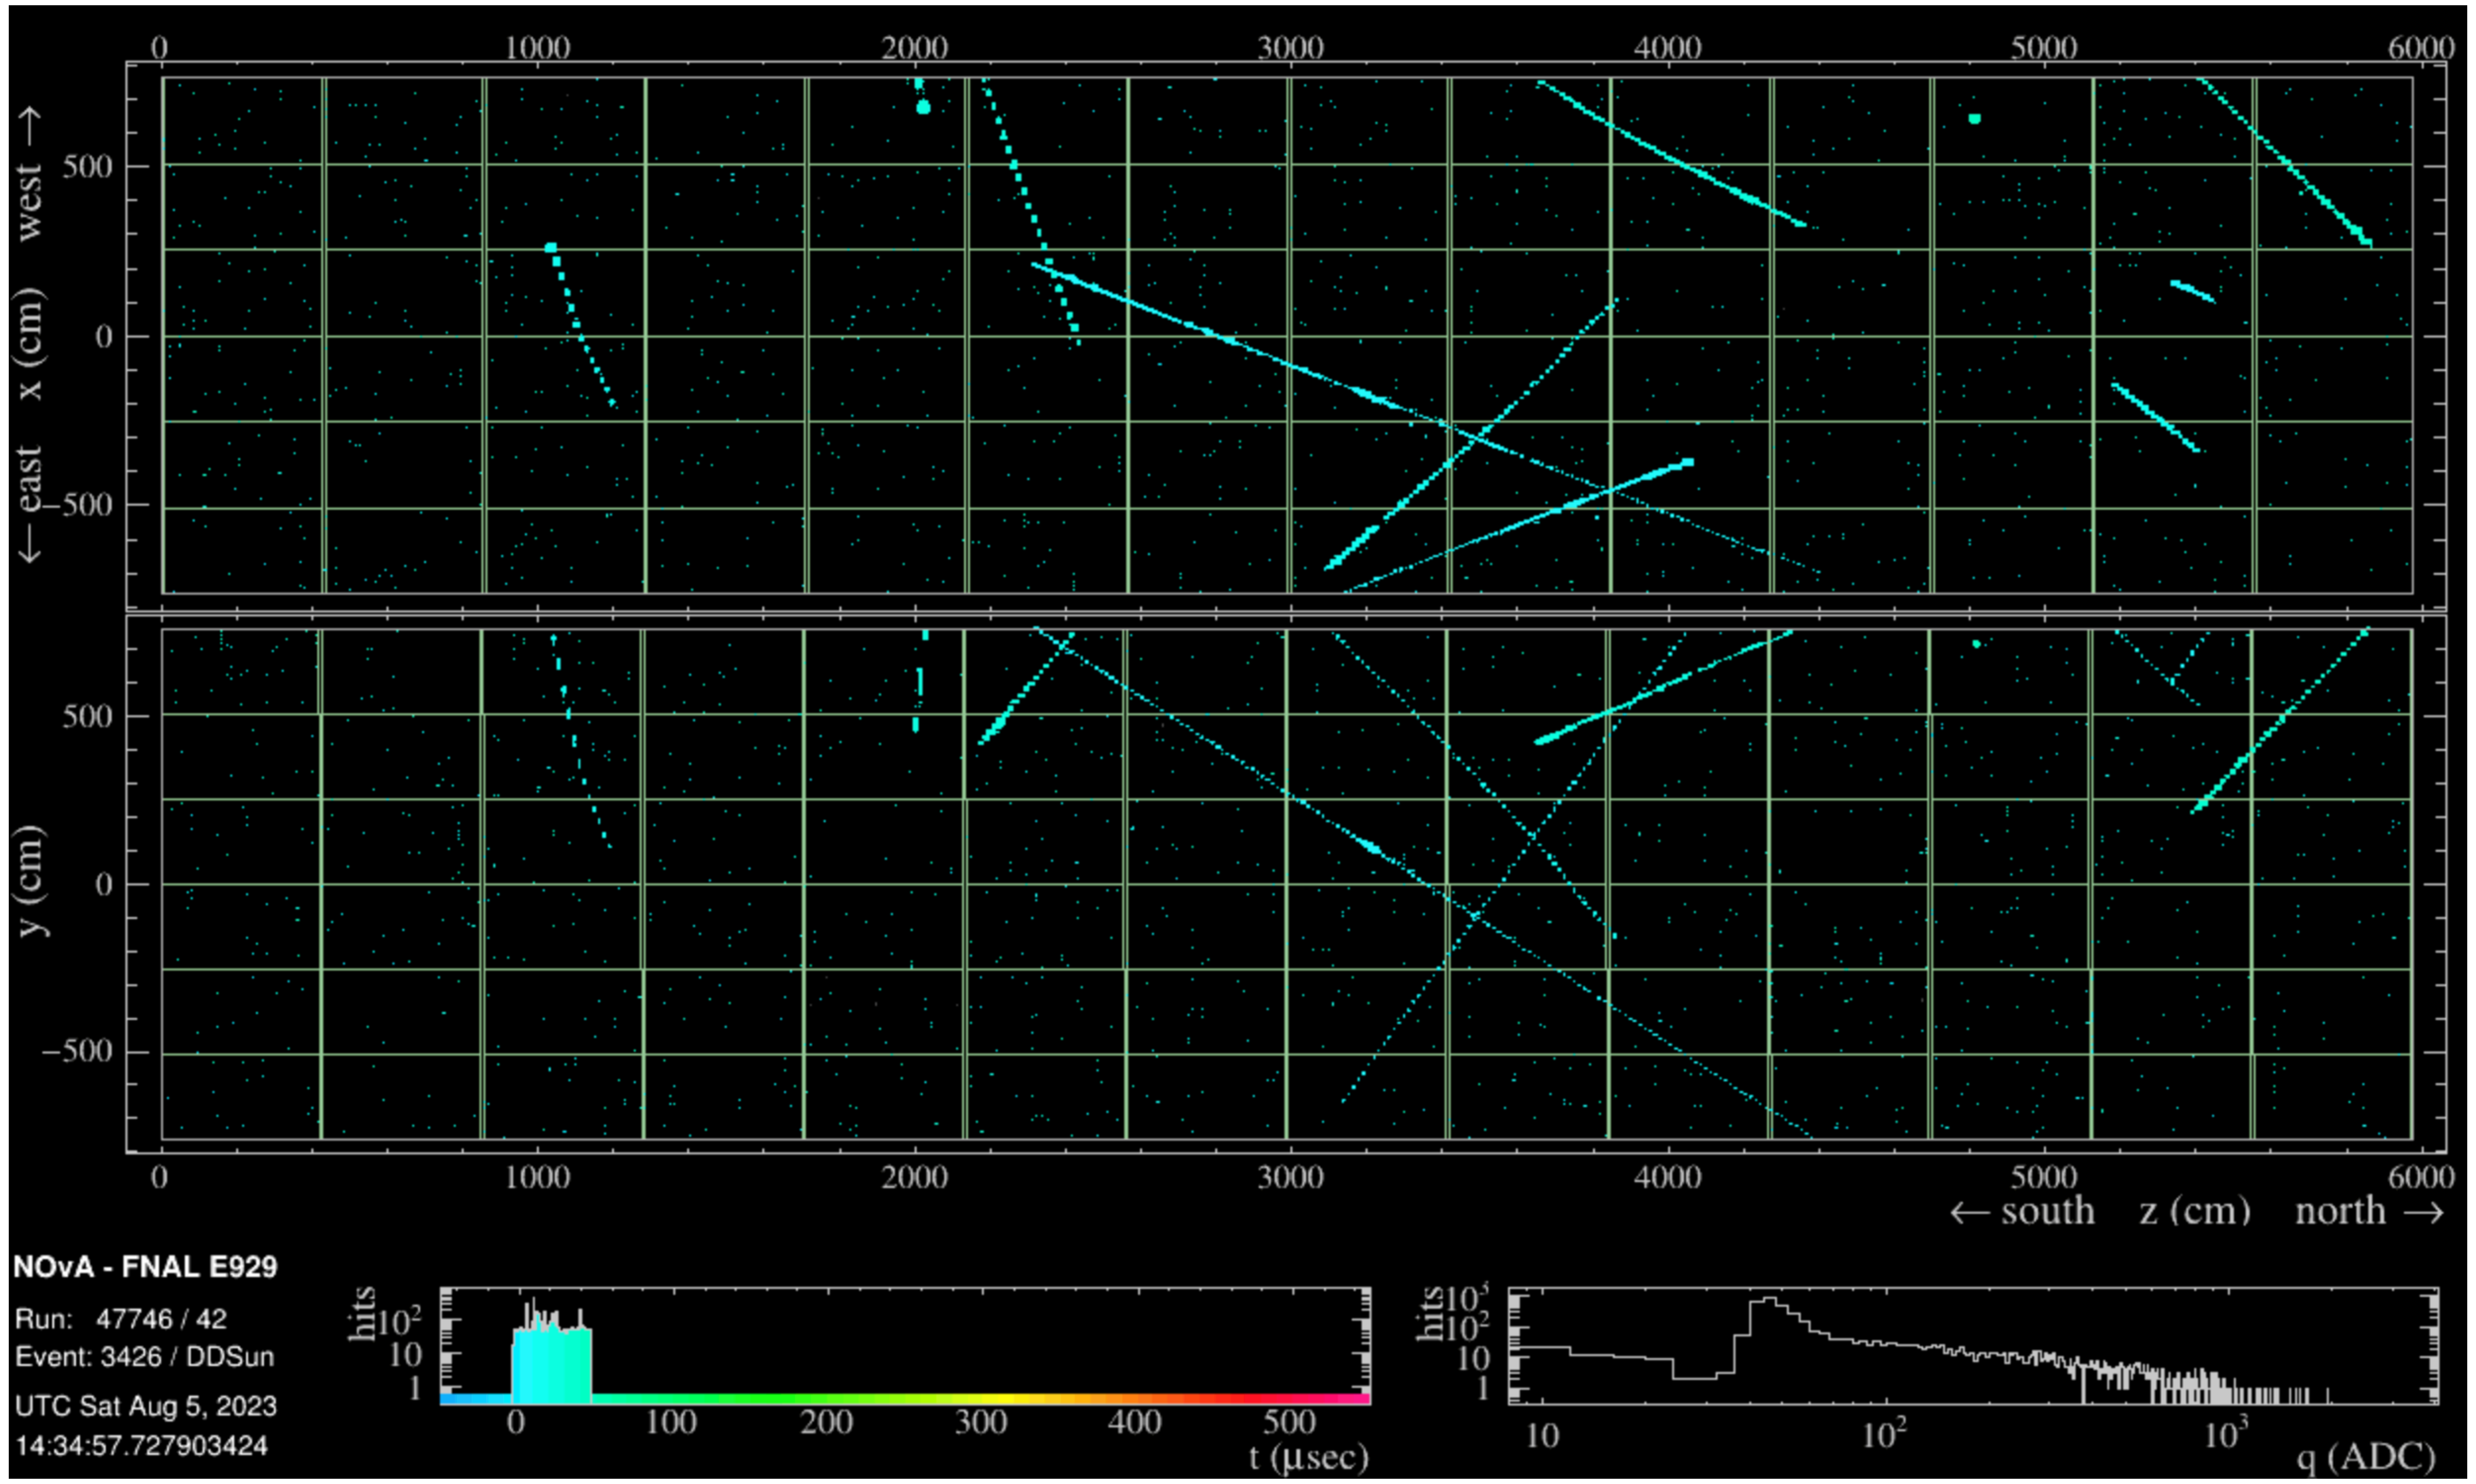

The First CNNs: NOvA

- The NOvA experiment was the first to use a CNN
  - Paper published in 2016[1]
  - A CNN was used for event classification
  - Based on the GoogLeNet (Inception) architecture
  - 40% increase in efficiency with no loss of purity for their main CC ν_e analysis

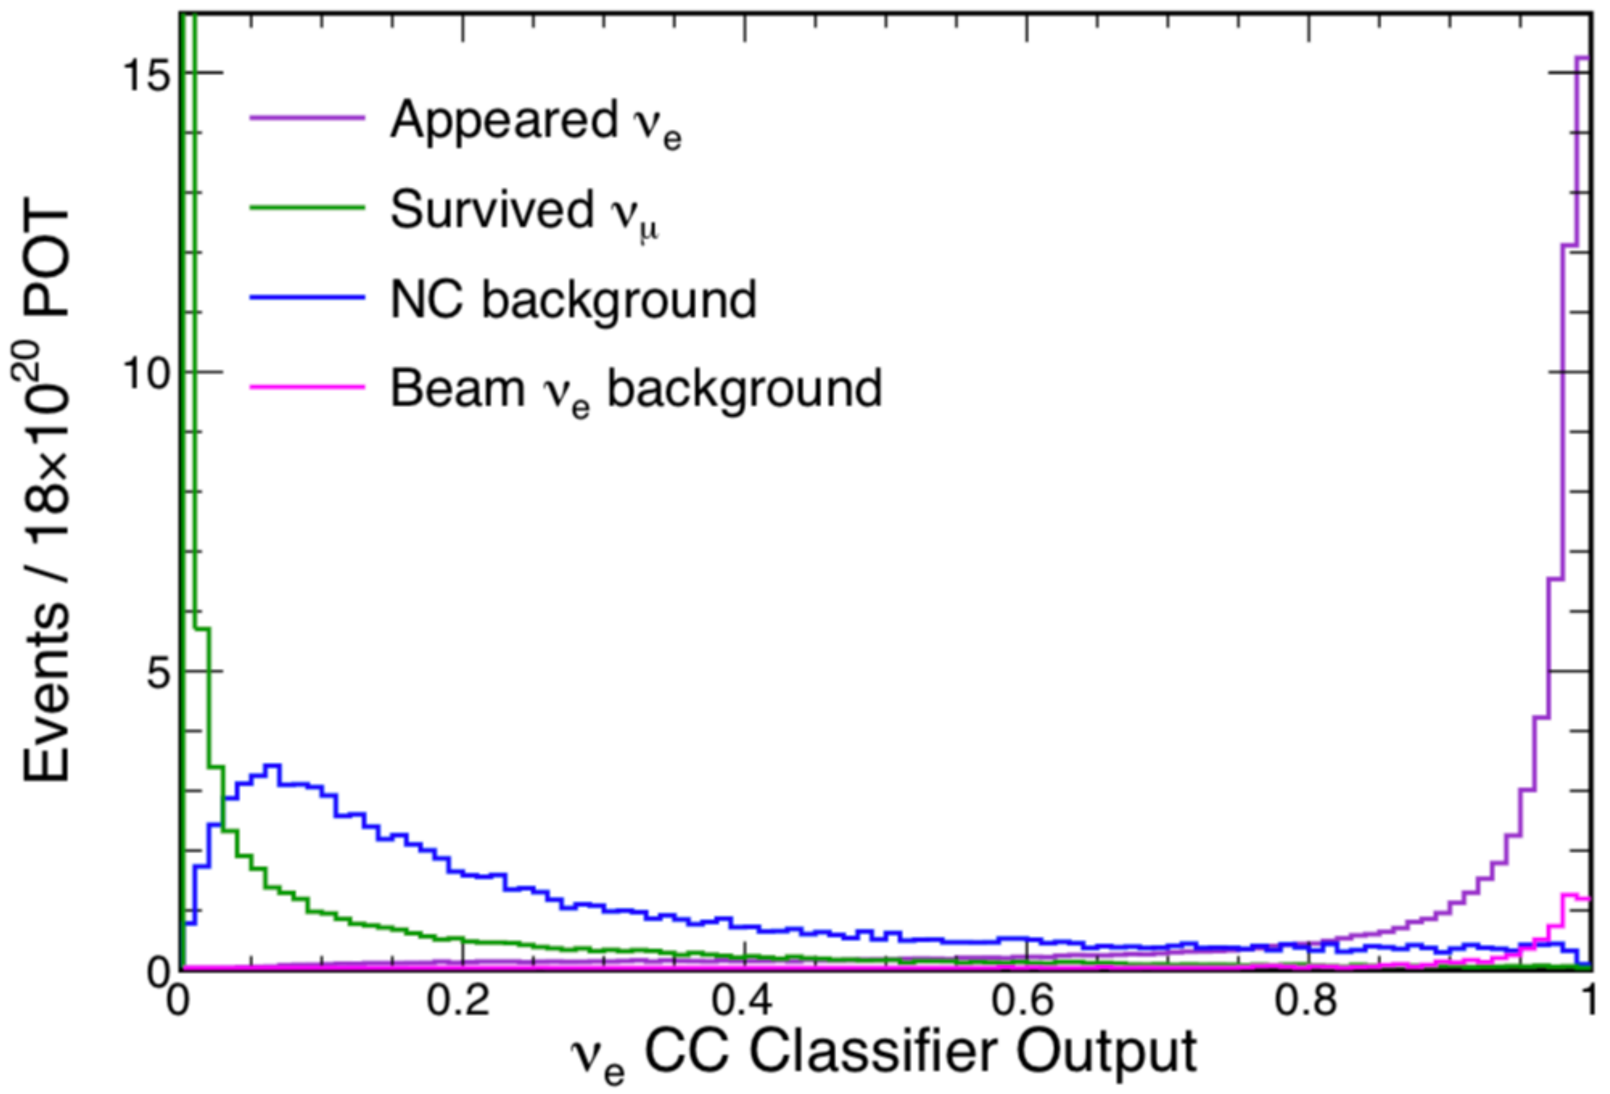

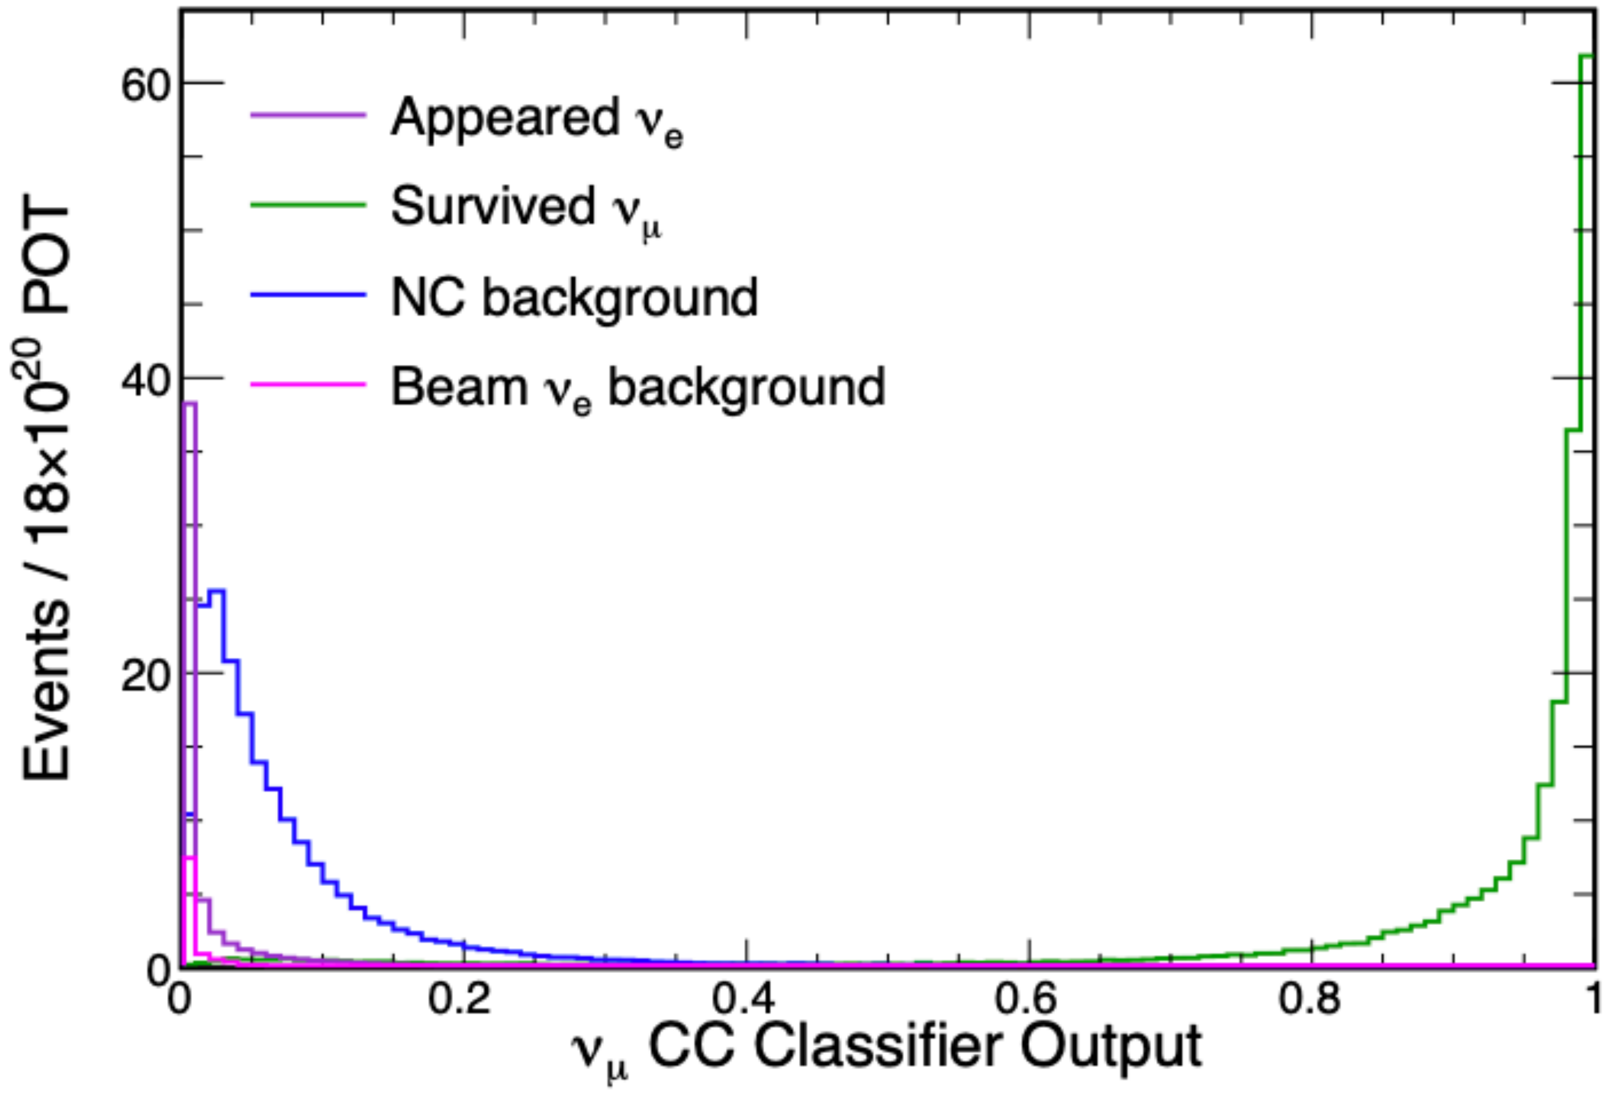

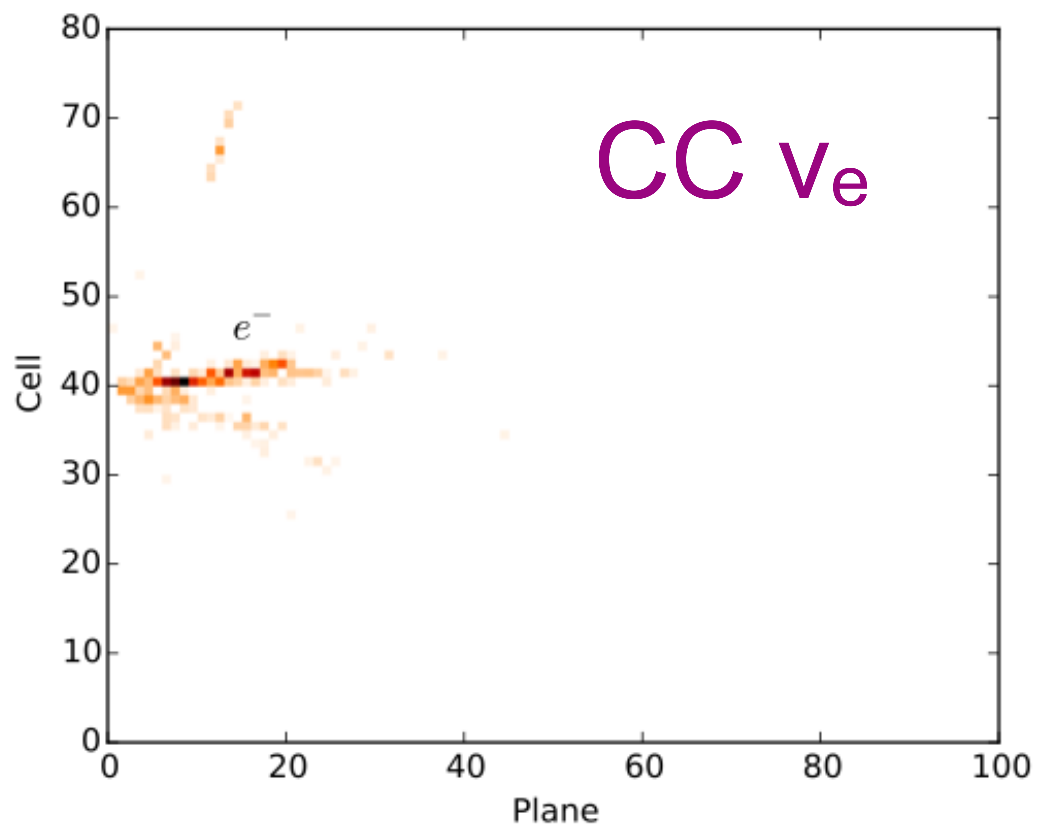

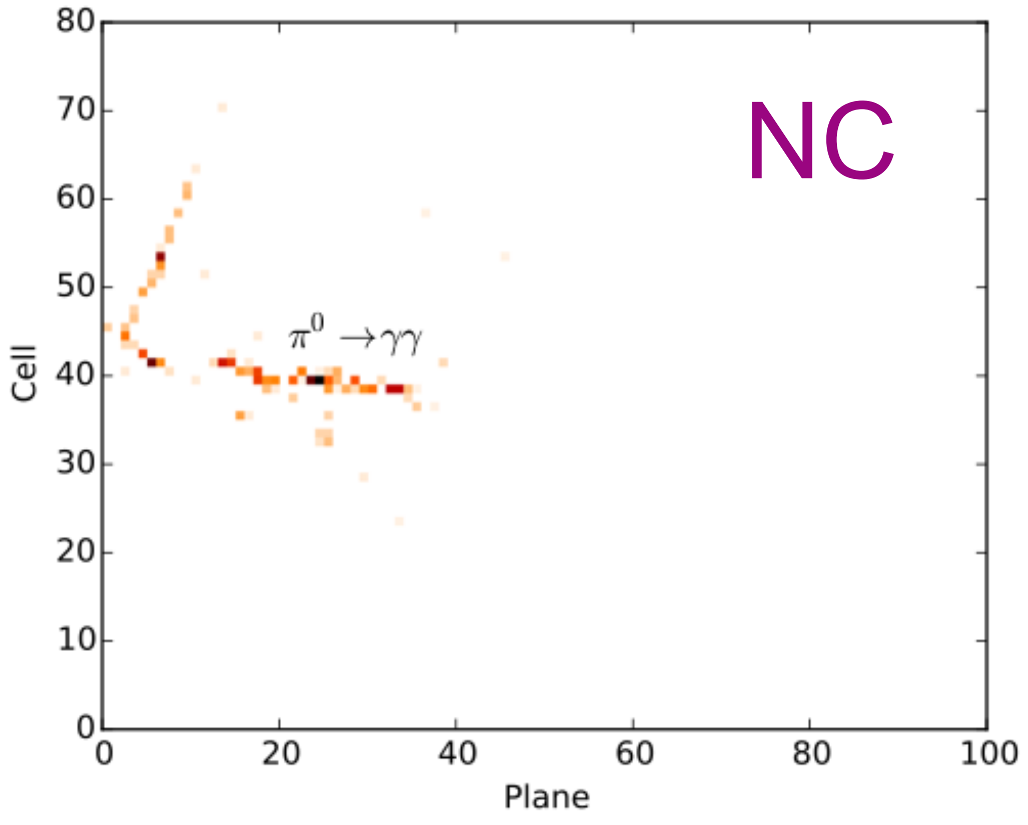

[1] A. Aurisano, et al., *A convolutional neural network neutrino event classifier,* Journal of Instrumentation 11 (2016) 09, P09001

The First CNNs: MicroBooNE

- MicroBooNE was a short-baseline neutrino detector

- It was built to investigate the LSND and MiniBooNE low energy excess of electron-like events

- Liquid argon TPC with 3 x 2D readout

- Collected neutrinos from Fermilab's Booster and NuMI beams

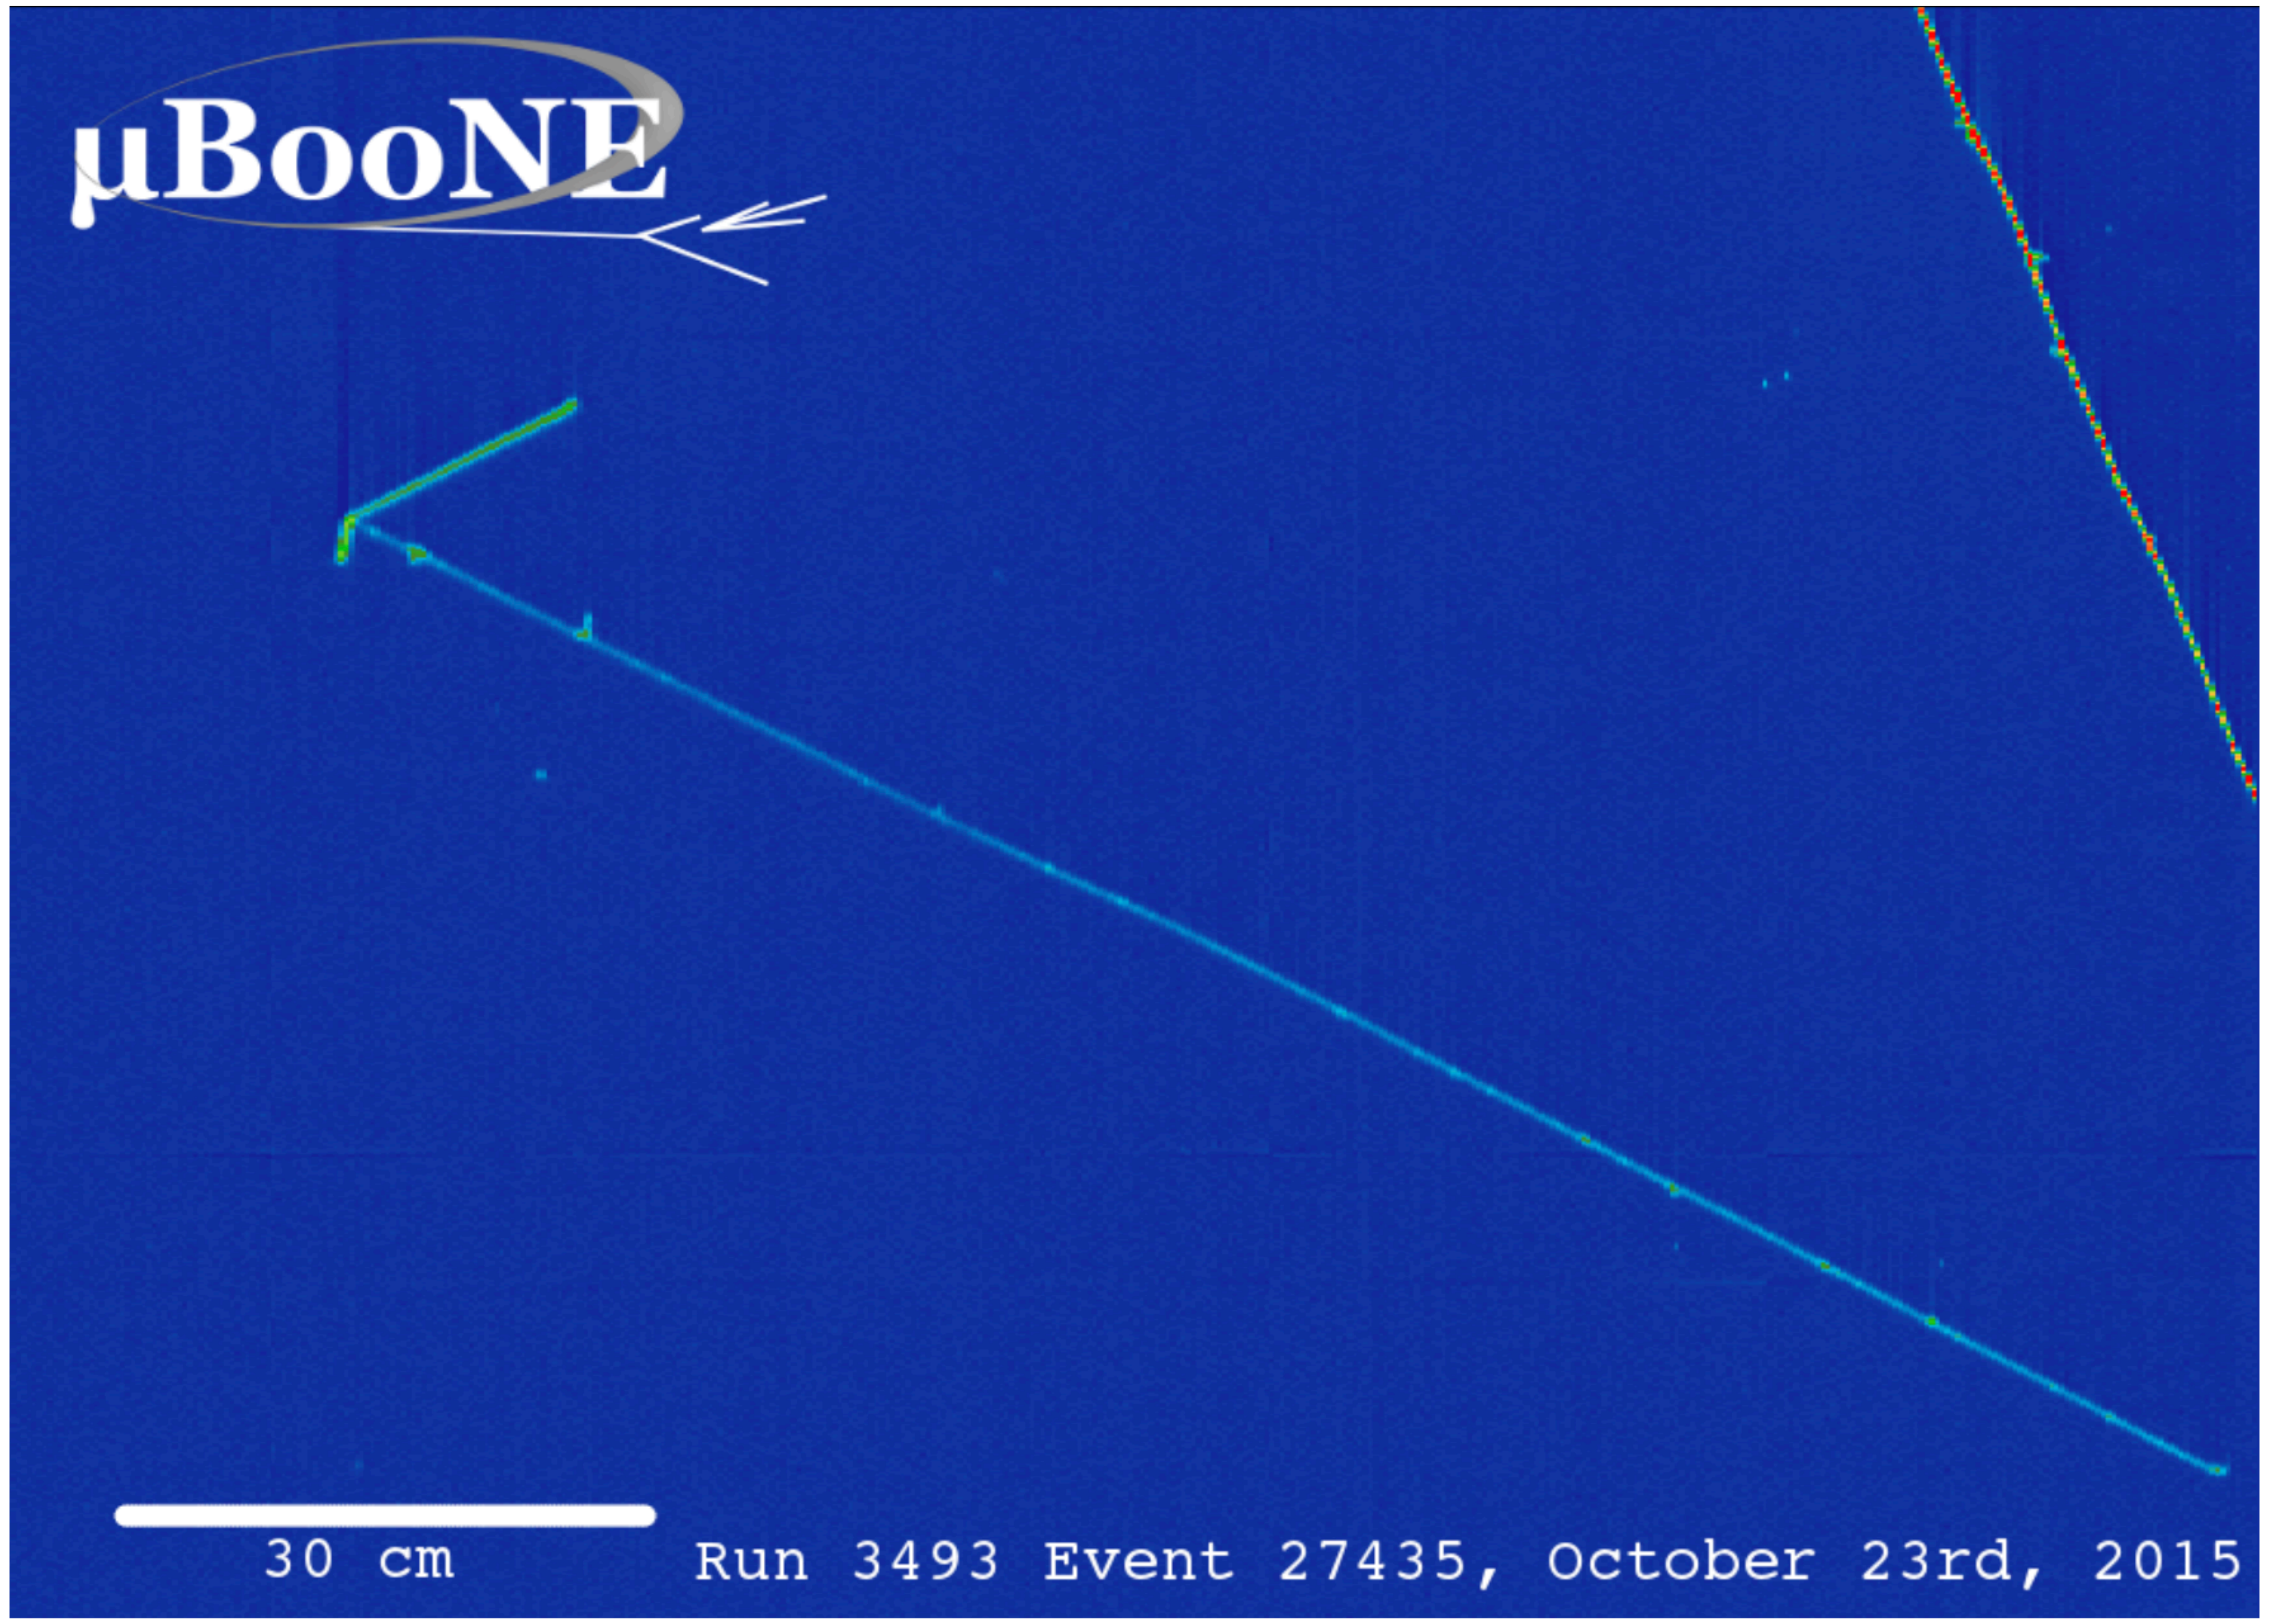

The First CNNs: MicroBooNE

- MicroBooNE: first LArTPC experiment to use a CNN
  - Study published in 2017[1]
- Paper includes a number of use-cases, I'll focus on one
  - Used for particle and neutrino detection and classification
  - Based on Faster R-CNN architecture

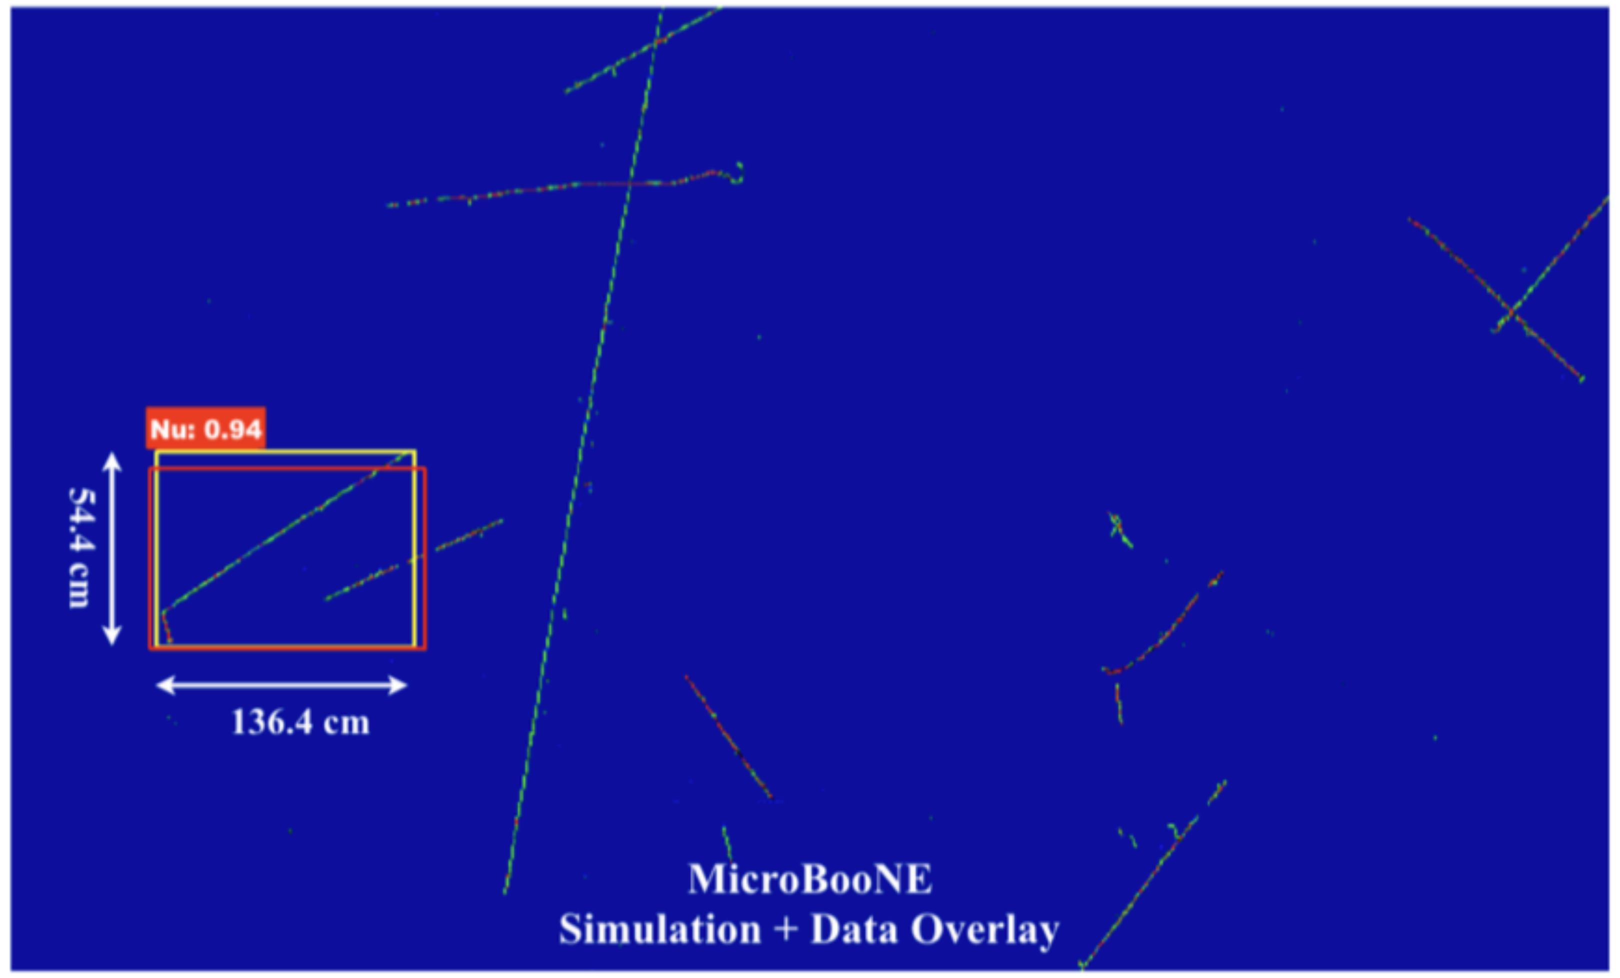

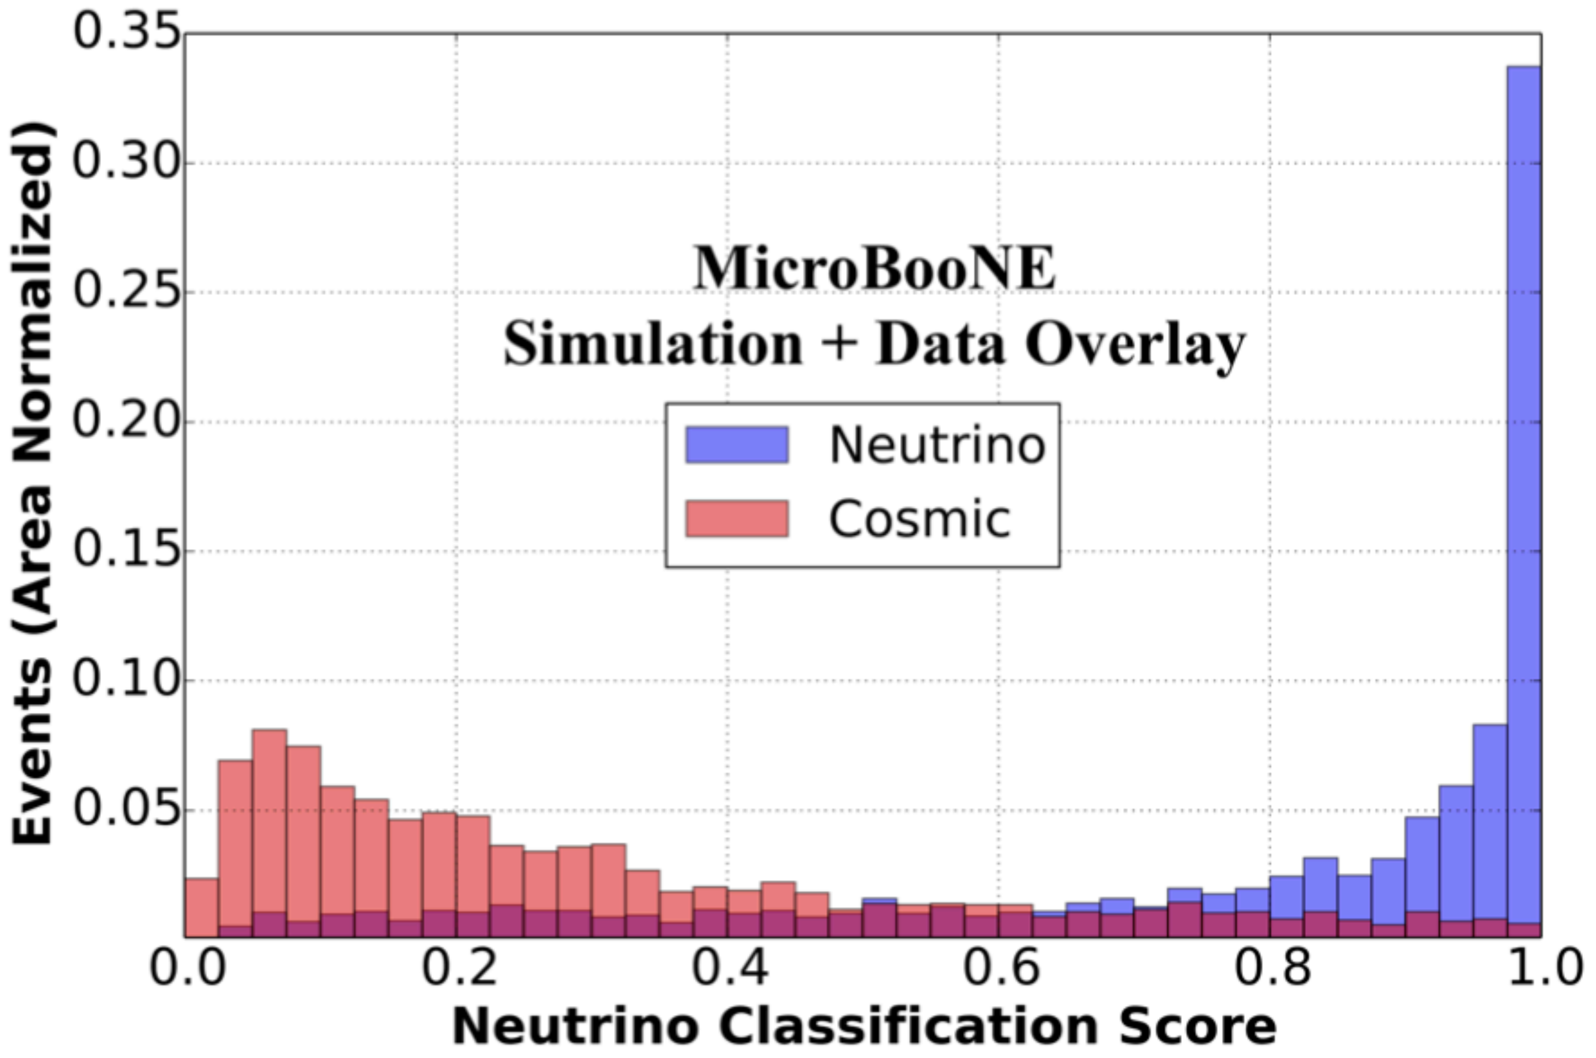

[1] MicroBooNE Collaboration, *Convolutional Neural Networks Applied to Neutrino Events in a Liquid Argon Time Projection Chamber*, JINST 12 (2017) 03, P03011

DUNE Neutrino Event Classifier

- As a concrete example, I will discuss the neutrino event classifier from the Deep Underground Neutrino Experiment (DUNE)

- DUNE is a future long-baseline neutrino oscillation experiment

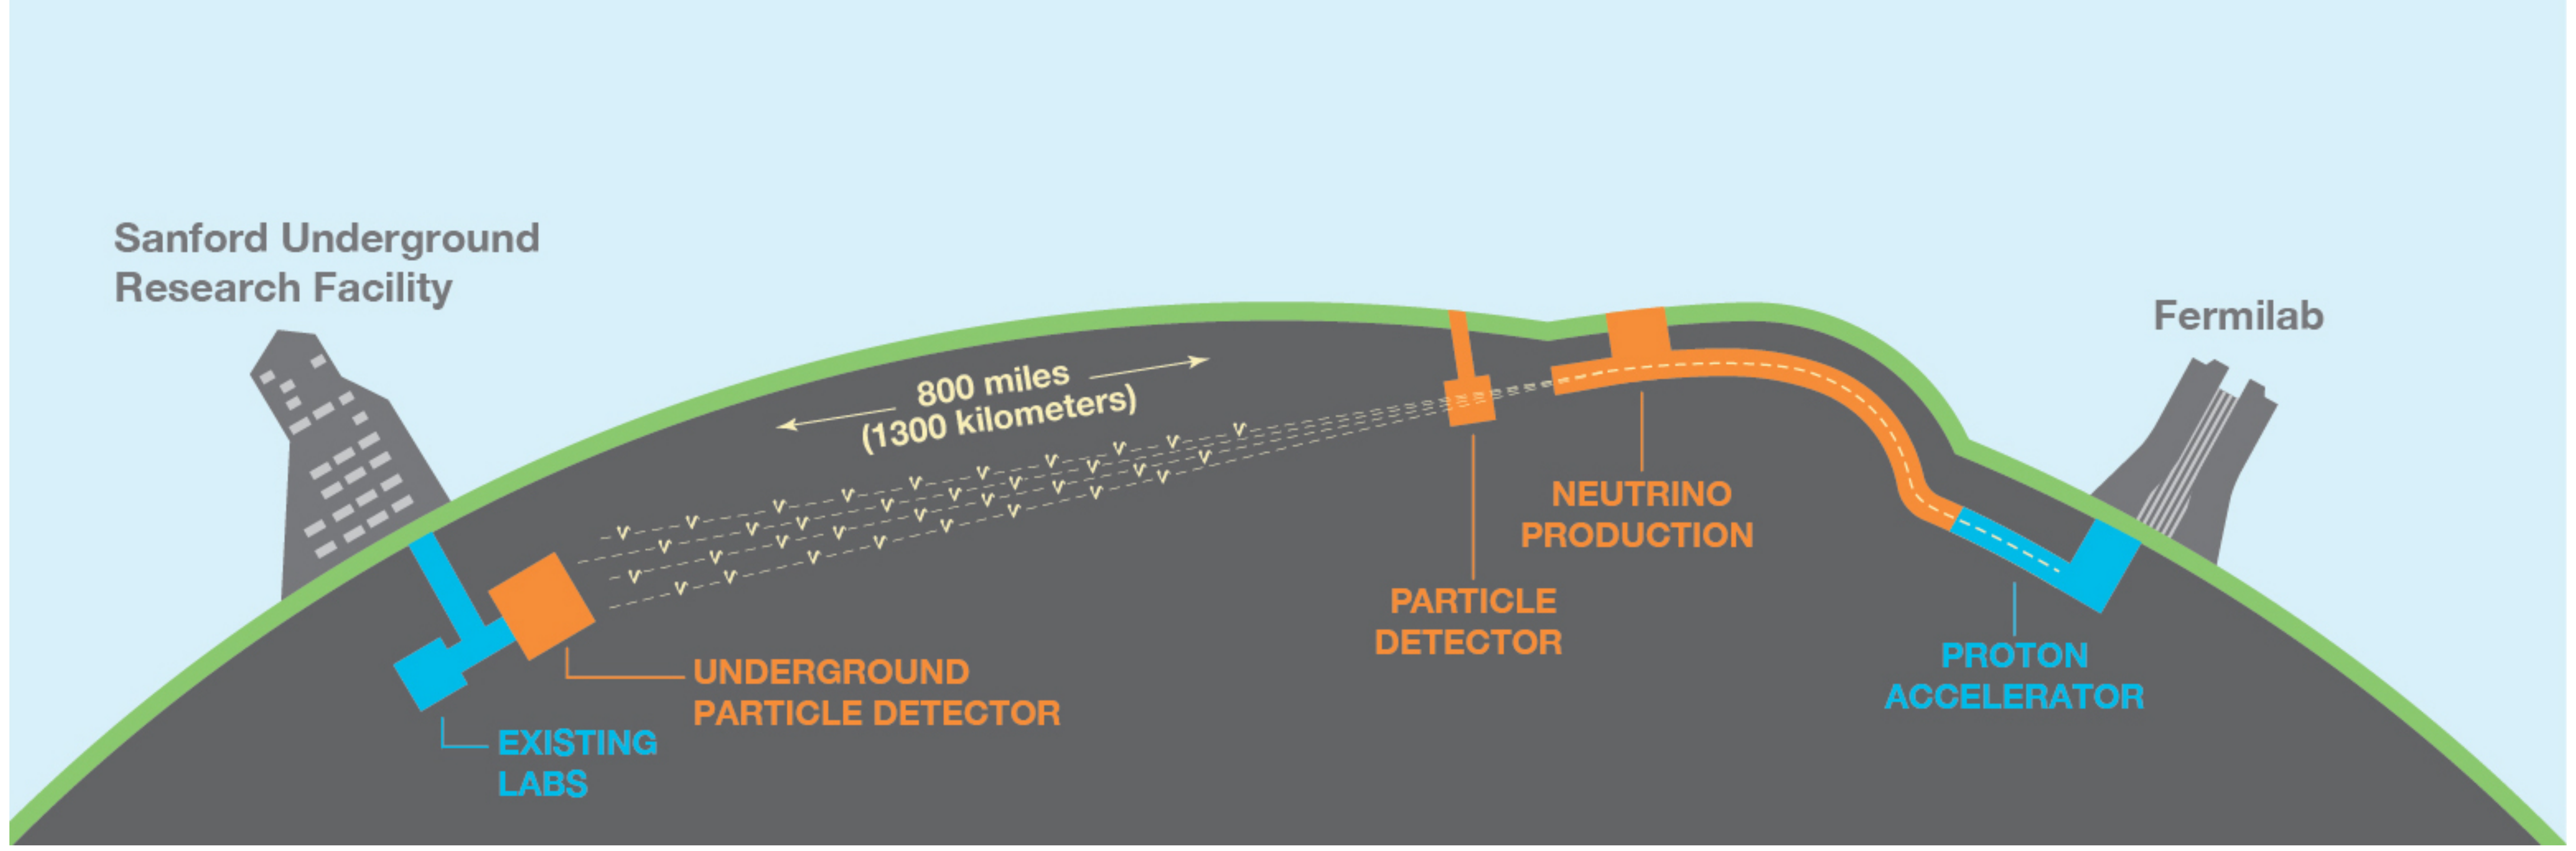

- Primary physics goal is to measure CP-violation in the neutrino sector
  - In order to do that, we need to distinguish different types of neutrinos

Introduction to DUNE

- The DUNE Far Detector will be made from four modules
  - Each module will be a liquid argon time projection chamber (LArTPC)
  - Three 2D projections of each interaction sharing one common coordinate

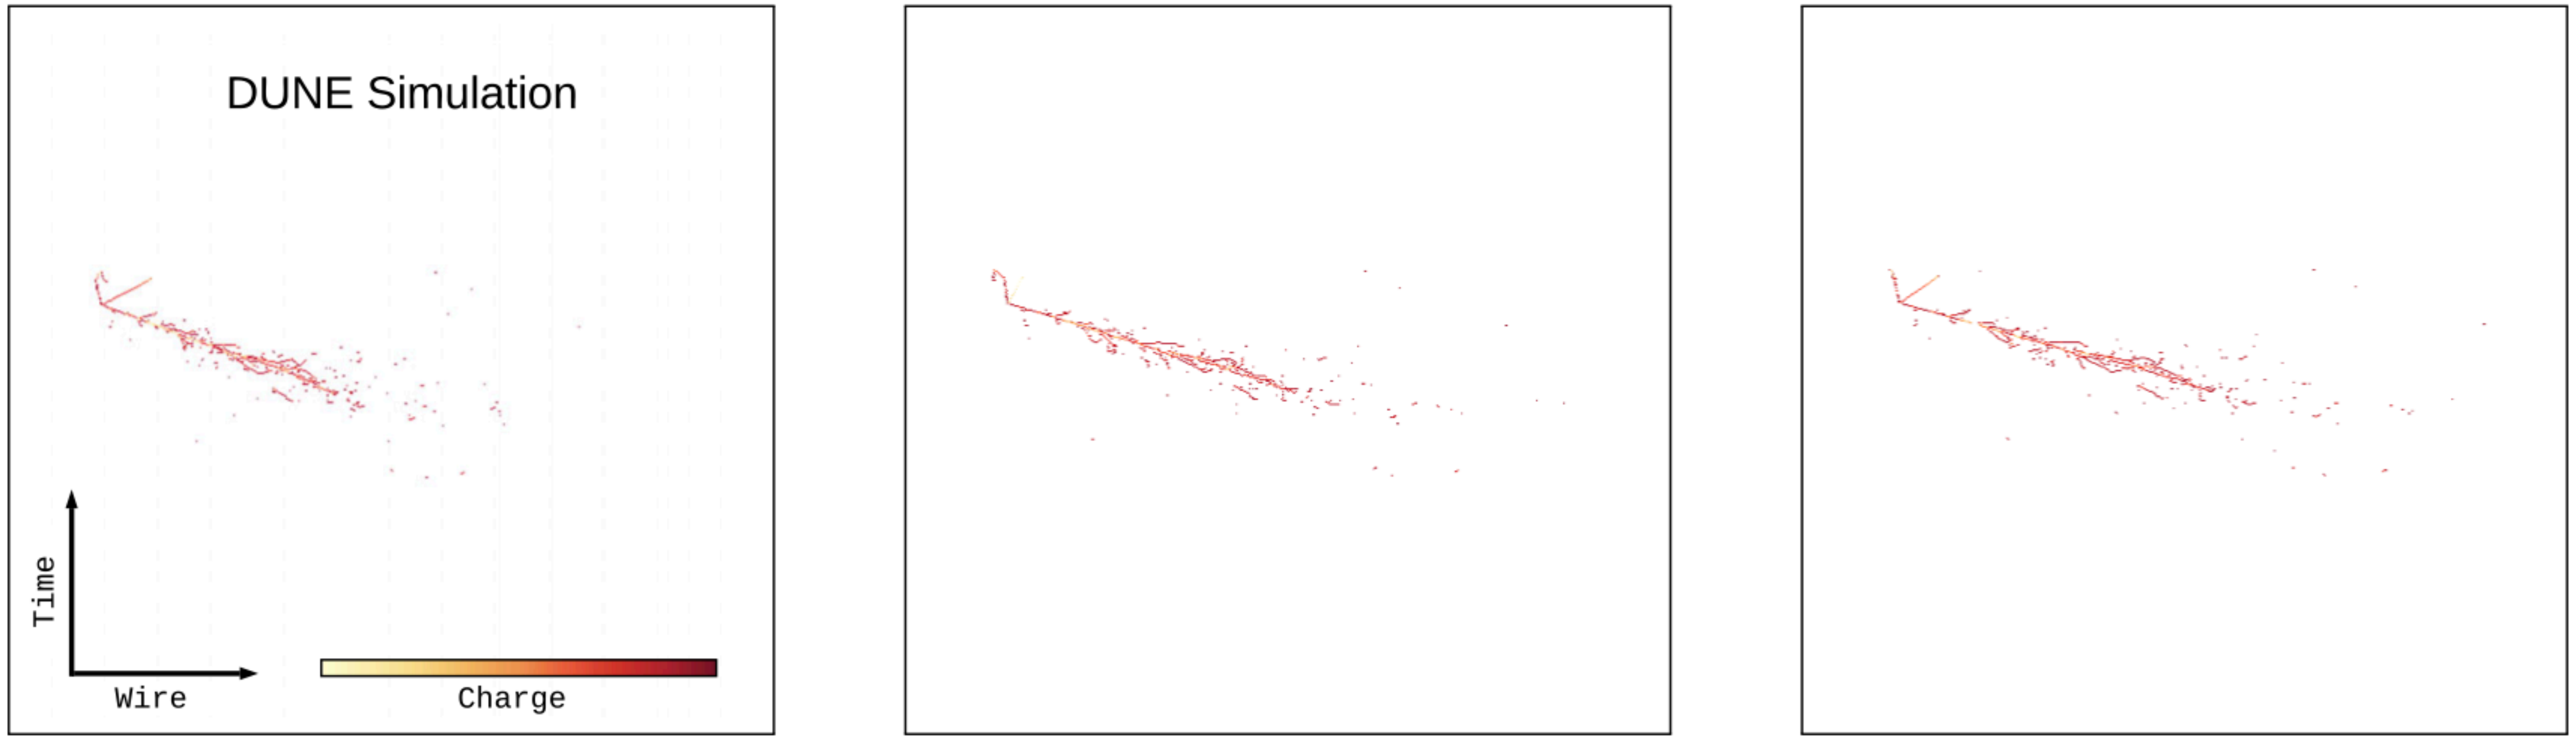

- The events are naturally representable as images
  - A CNN is an obvious choice of algorithm to extract information from the data

Figures reproduced from: DUNE Collaboration, *Neutrino interaction classification with a convolutional neural network in the DUNE far detector,* Phys.Rev.D 102 (2020) 9, 092003.

DUNE

- Need to identify CC ν_μ and CC ν_e reject background events
- Thus, the DUNE CVN[1,2] (it is a CNN) aims to classify beam neutrino events as:
  - CC ν_μ, CC ν_e, CC ν_τ, and NC
  - CC ν_τ are rare and hard to classify, so I won't discuss them further

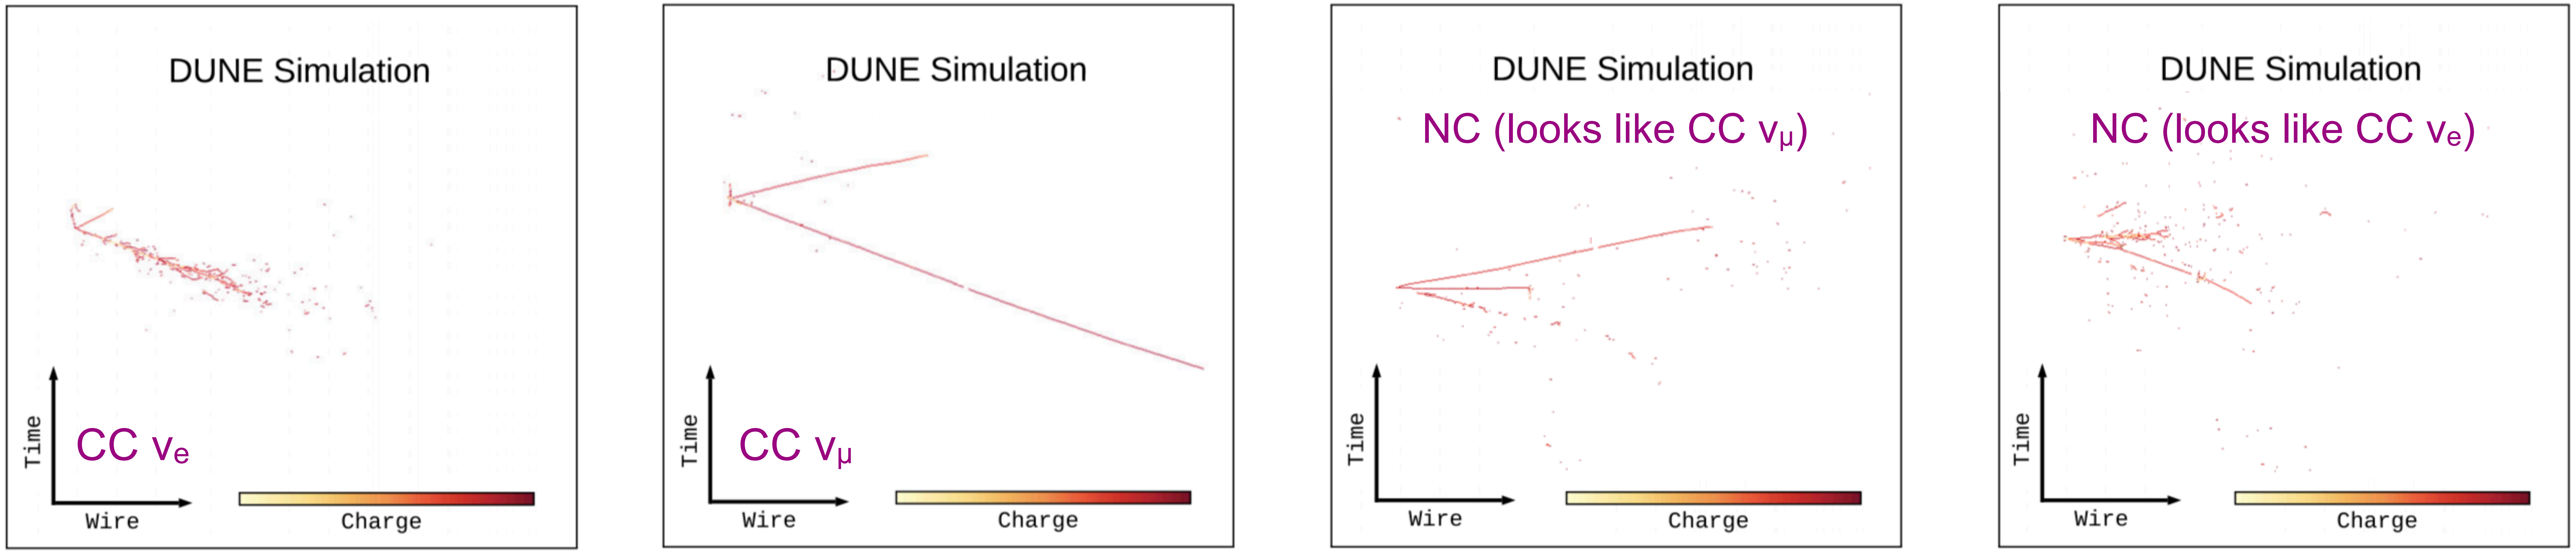

[1] DUNE Collaboration, *Neutrino interaction classification with a convolutional neural network in the DUNE far detector,* Phys. Rev. D **102** 9, 092003 (2020)
[2] S. Alonso Monsalve, *Novel usage of deep learning and high-performance computing in long-baseline neutrino oscillation experiments*, PhD Thesis, Universidad Carlos III Madrid (2021)

DUNE CVN

- Architecture based on SE-ResNet-34[1,2]
- Inputs processed separately for the first few blocks and then merged
- Main output is the flavour classifier
  - The top one shown in the figure
- Other particle counting outputs will be further studied in the future
- Trained on over 3 million events

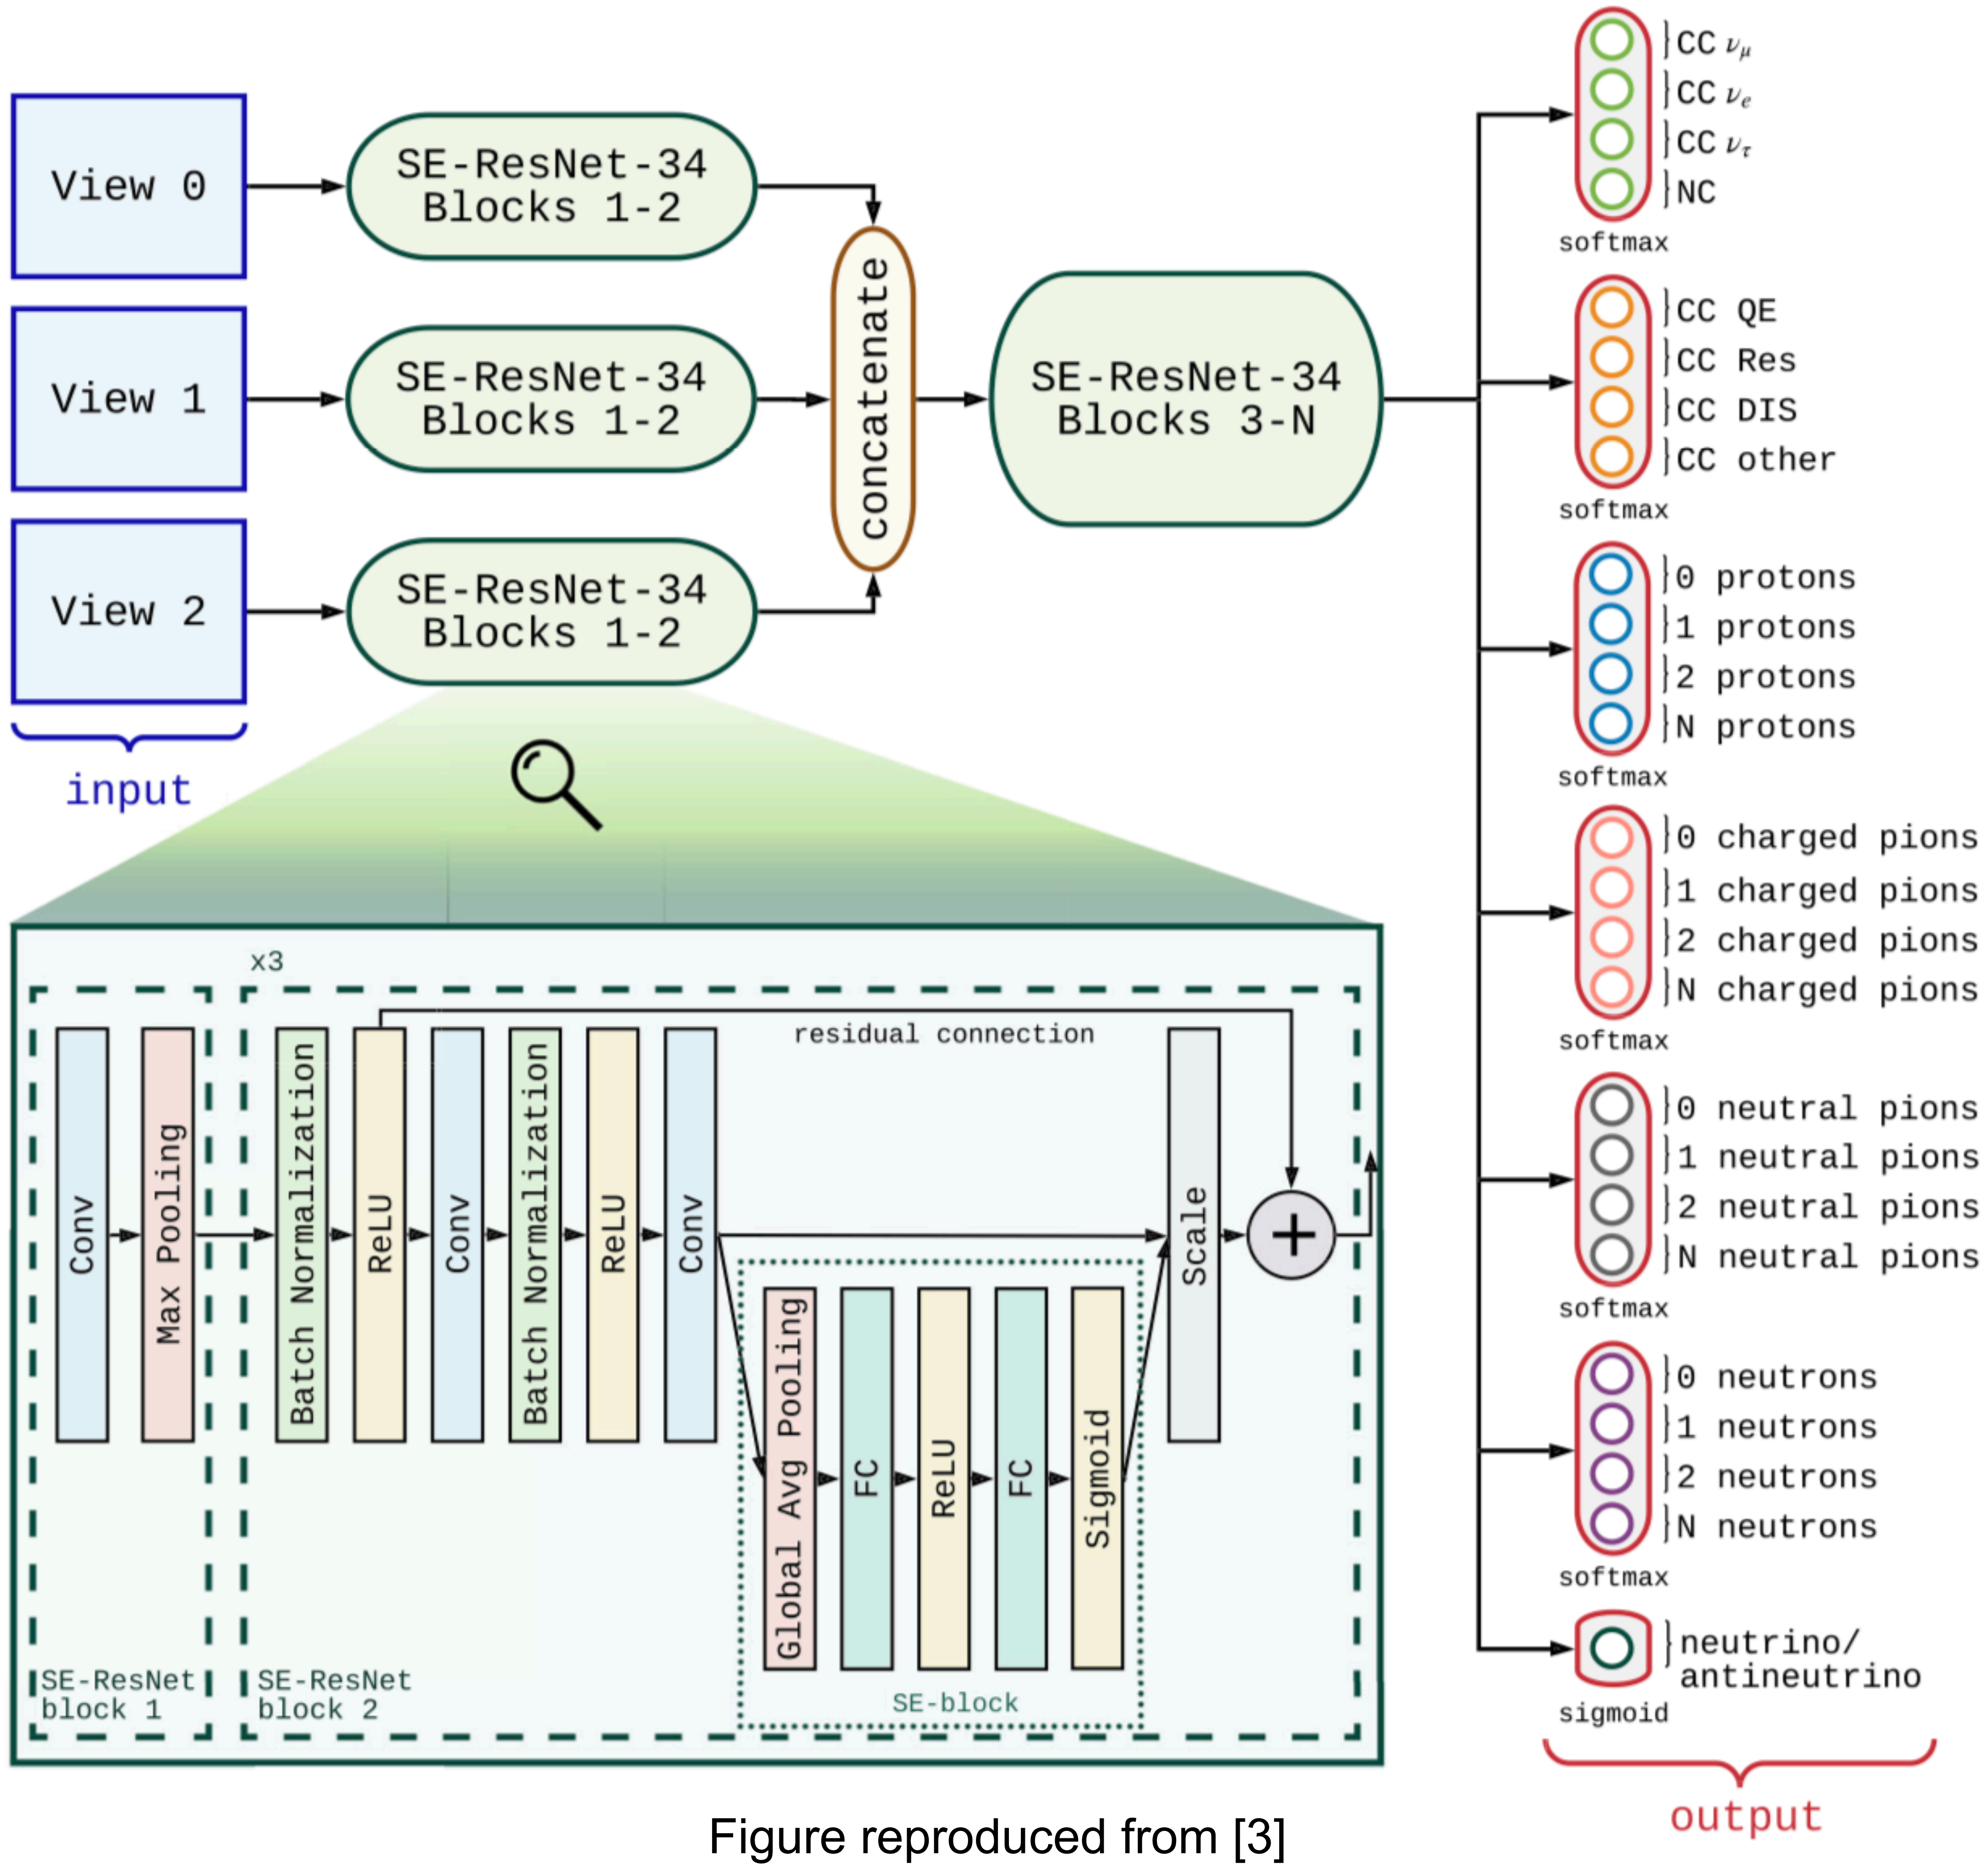

Figure reproduced from [3]

[1] K. He, X. Zhang, S. Ren, and J. Sun, *Deep Residual Learning for Image Recognition*, 1512.03385; K. He, X. Zhang, S. Ren, and J. Sun, *Identity Mappings in Deep Residual Networks*, 1603.05027
[2] J. Hu, L. Shen, and G. Sun, *Squeeze-and-Excitation Networks*, 1709.01507
[3] DUNE Collaboration, *Neutrino interaction classification with a convolutional neural network in the DUNE far detector*, Phys. Rev. D **102** 9, 092003 (2020)

DUNE CC ν_e selection

- See very good signal background separation

Neutrino mode

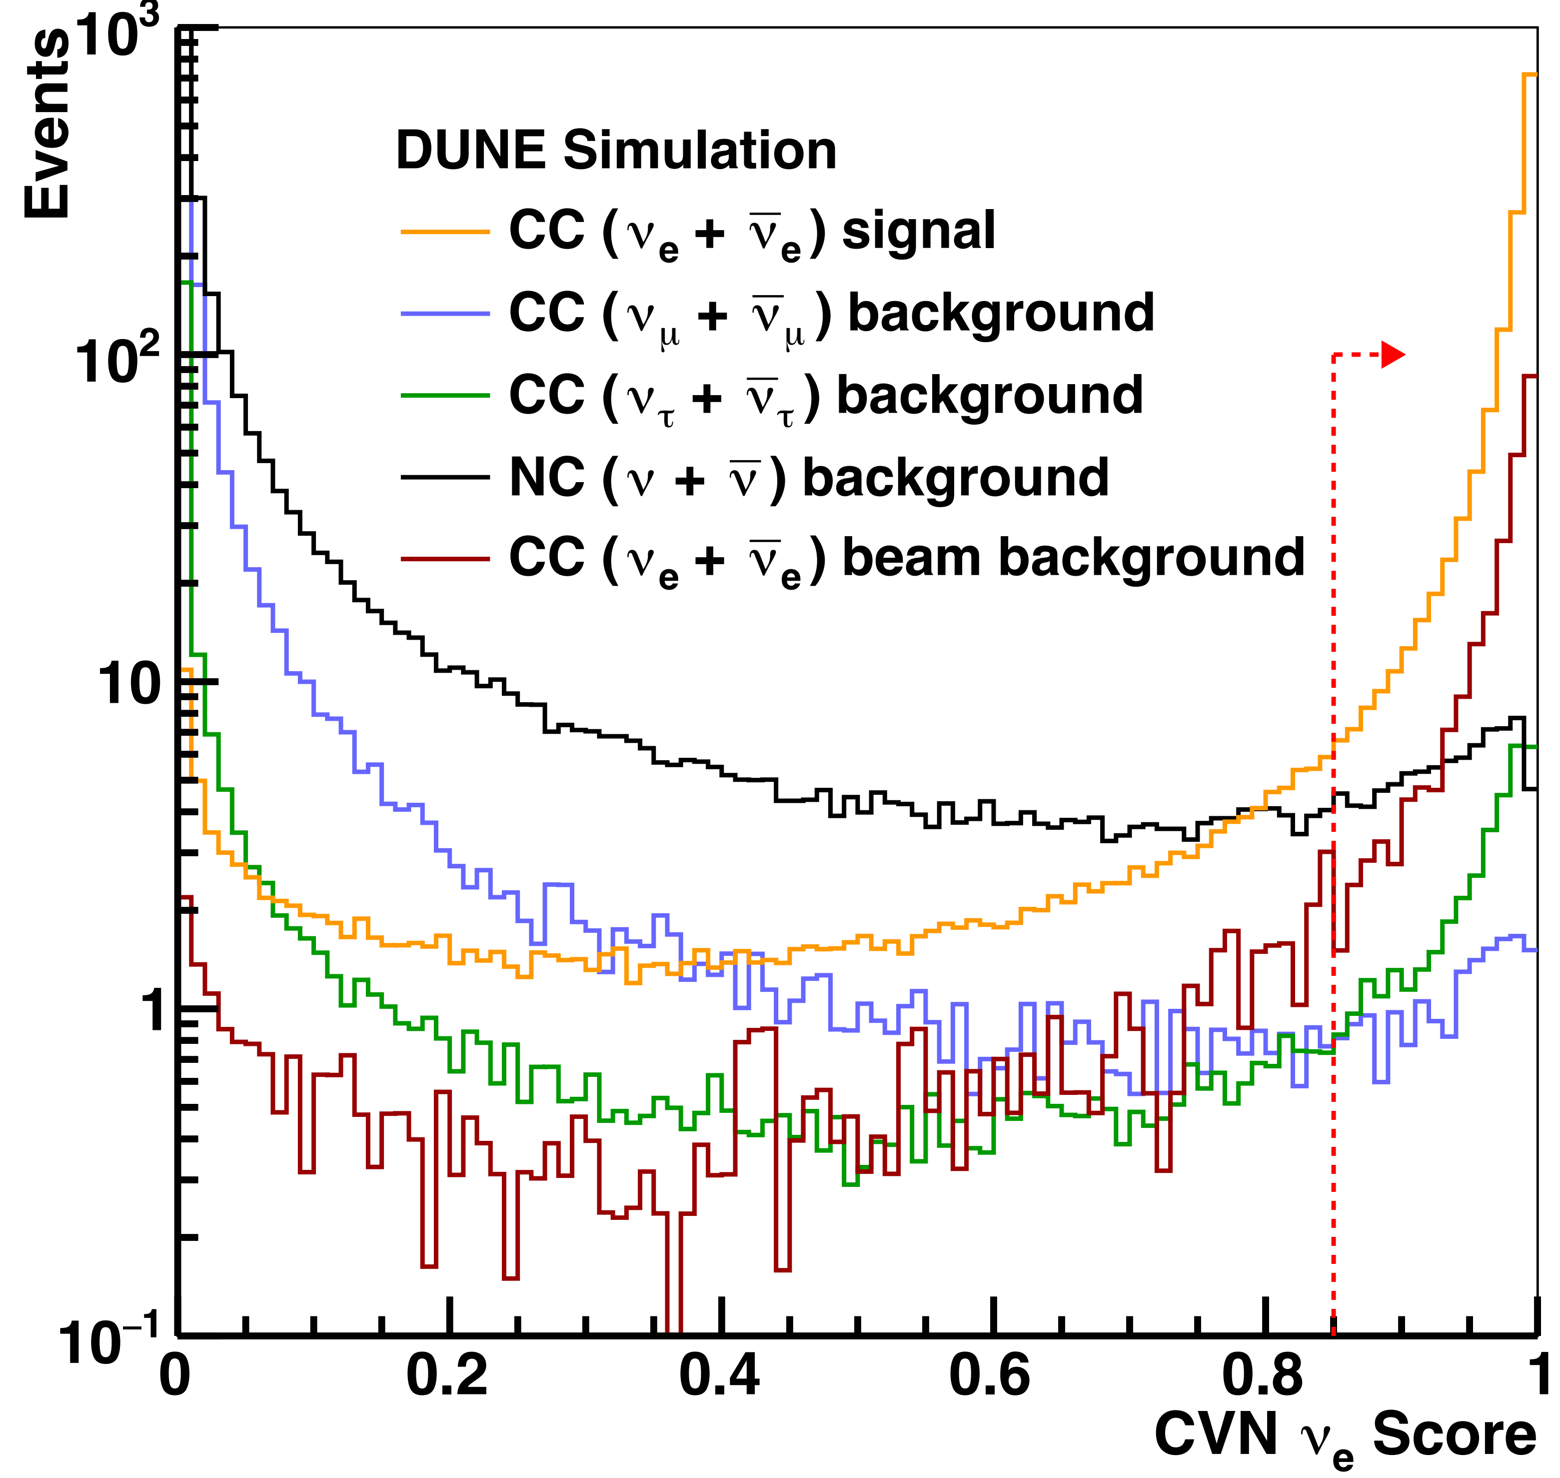

Antineutrino mode

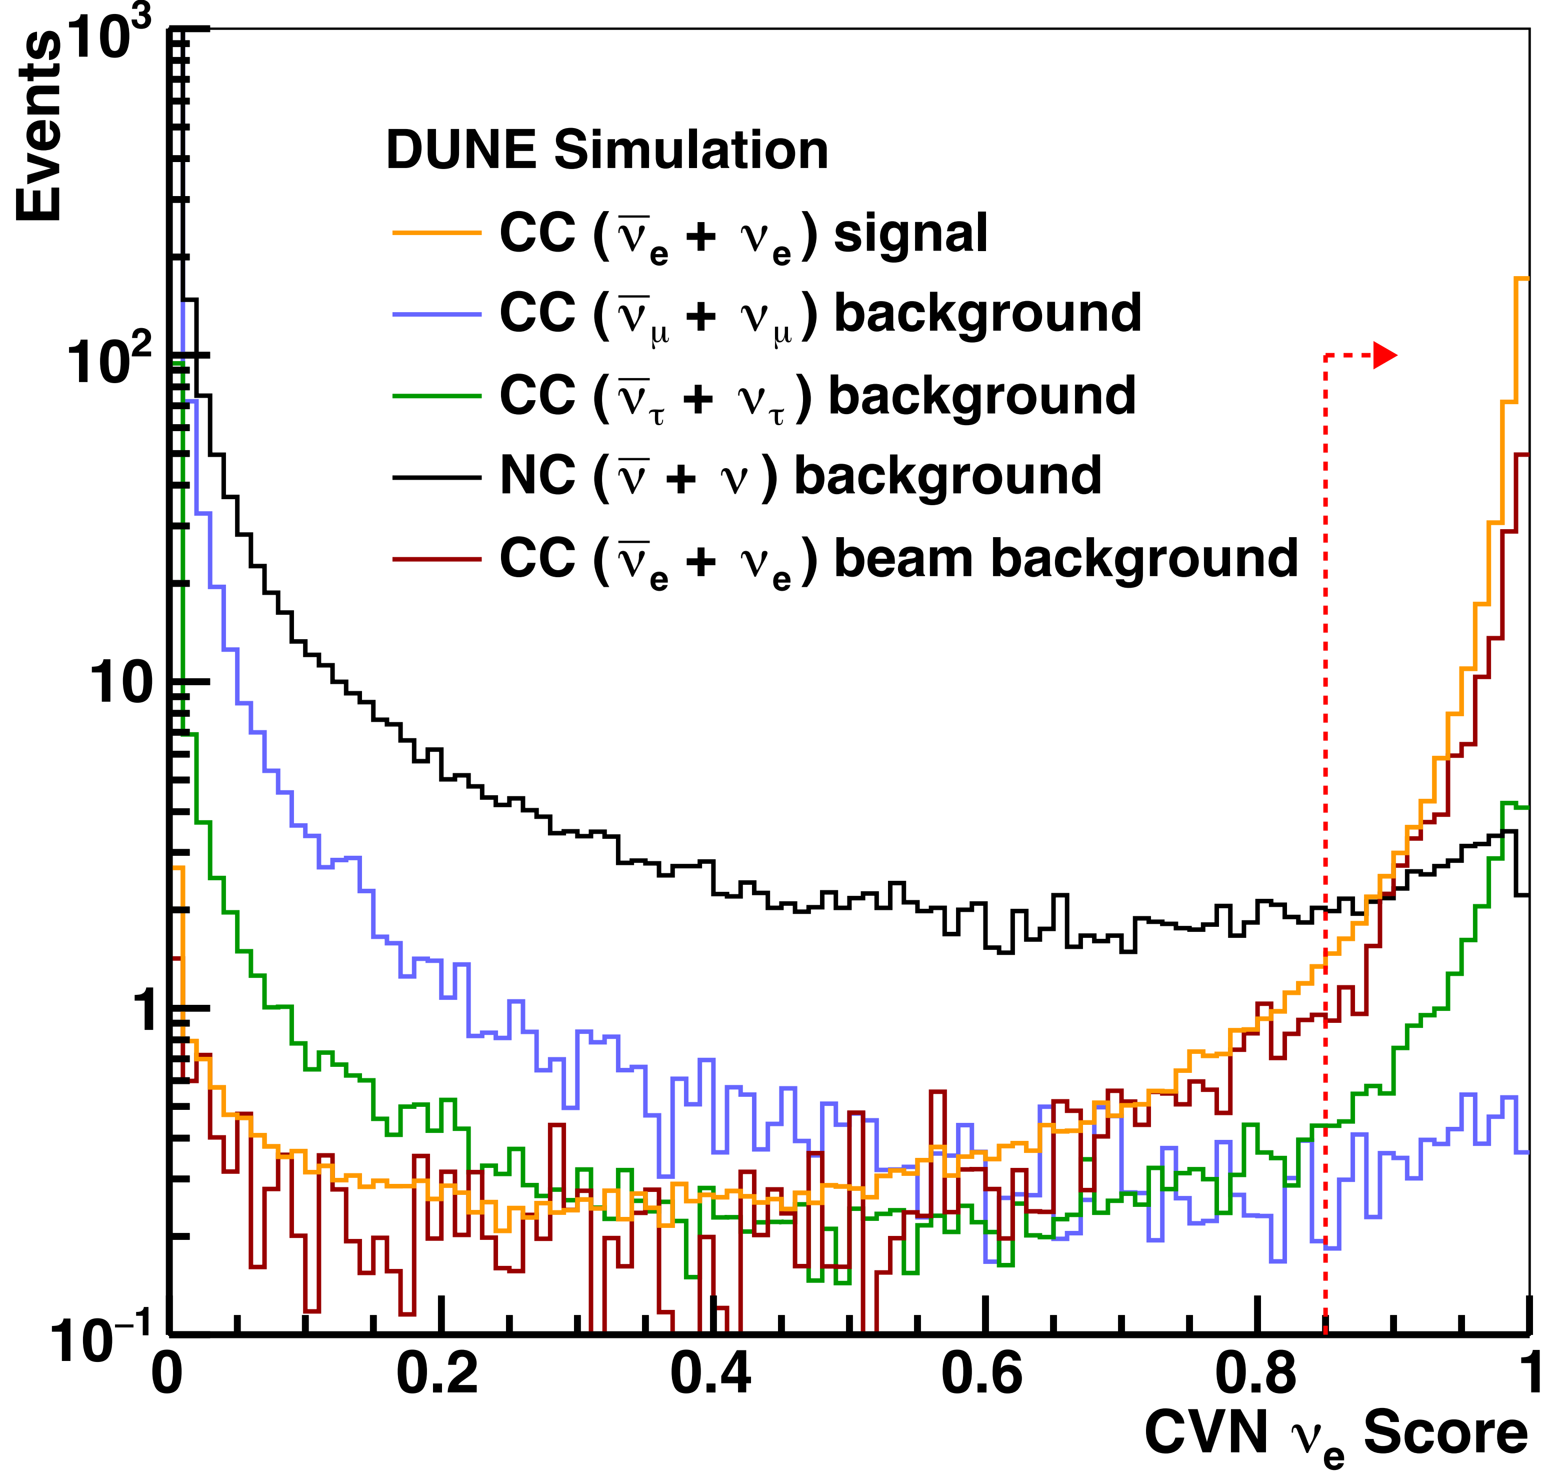

Arrows show events selected for the CC ν_e appearance sample

Figures reproduced from: DUNE Collaboration, *Neutrino interaction classification with a convolutional neural network in the DUNE far detector*, Phys. Rev. D **102** 9, 092003 (2020)

DUNE CC ν_μ selection

- See very good signal background separation

Neutrino mode

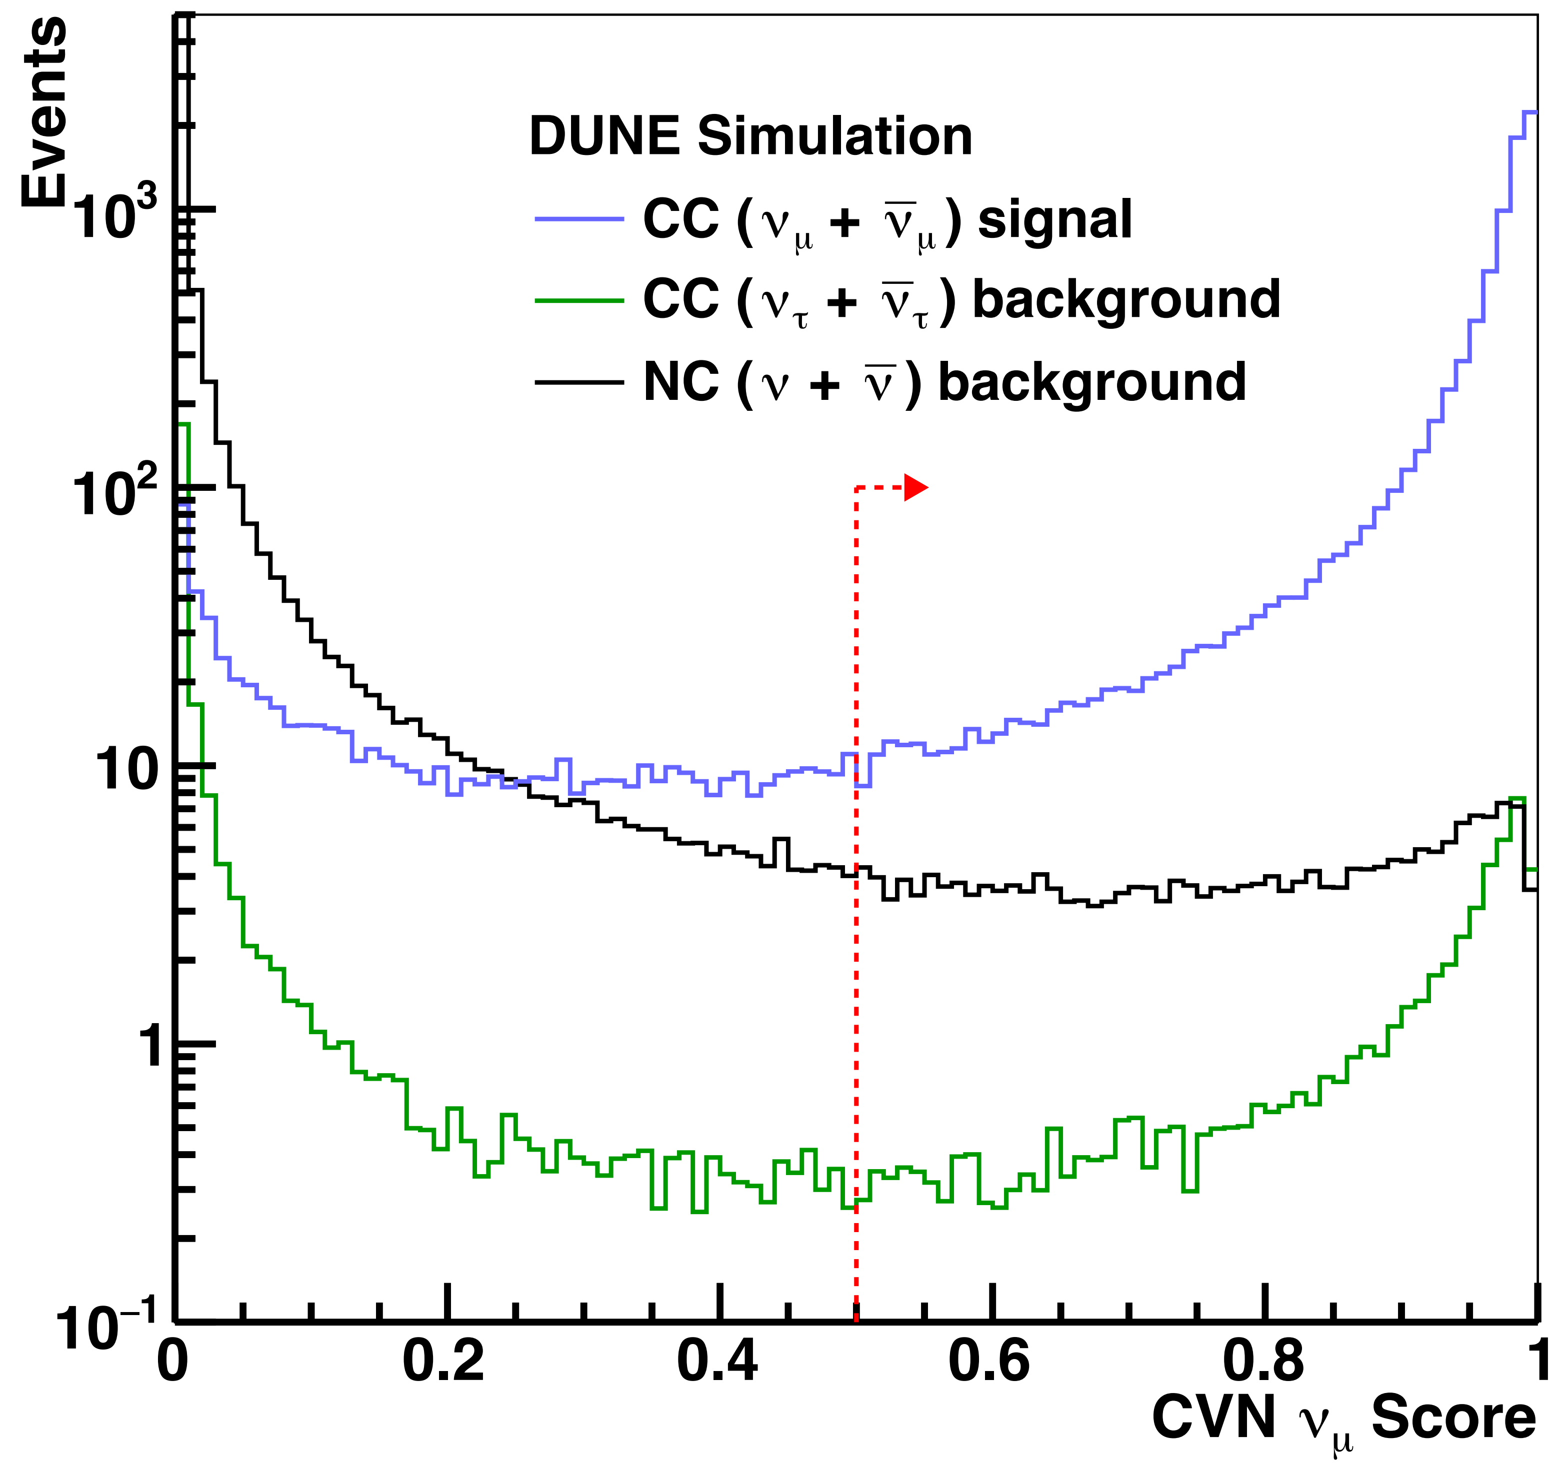

Antineutrino mode

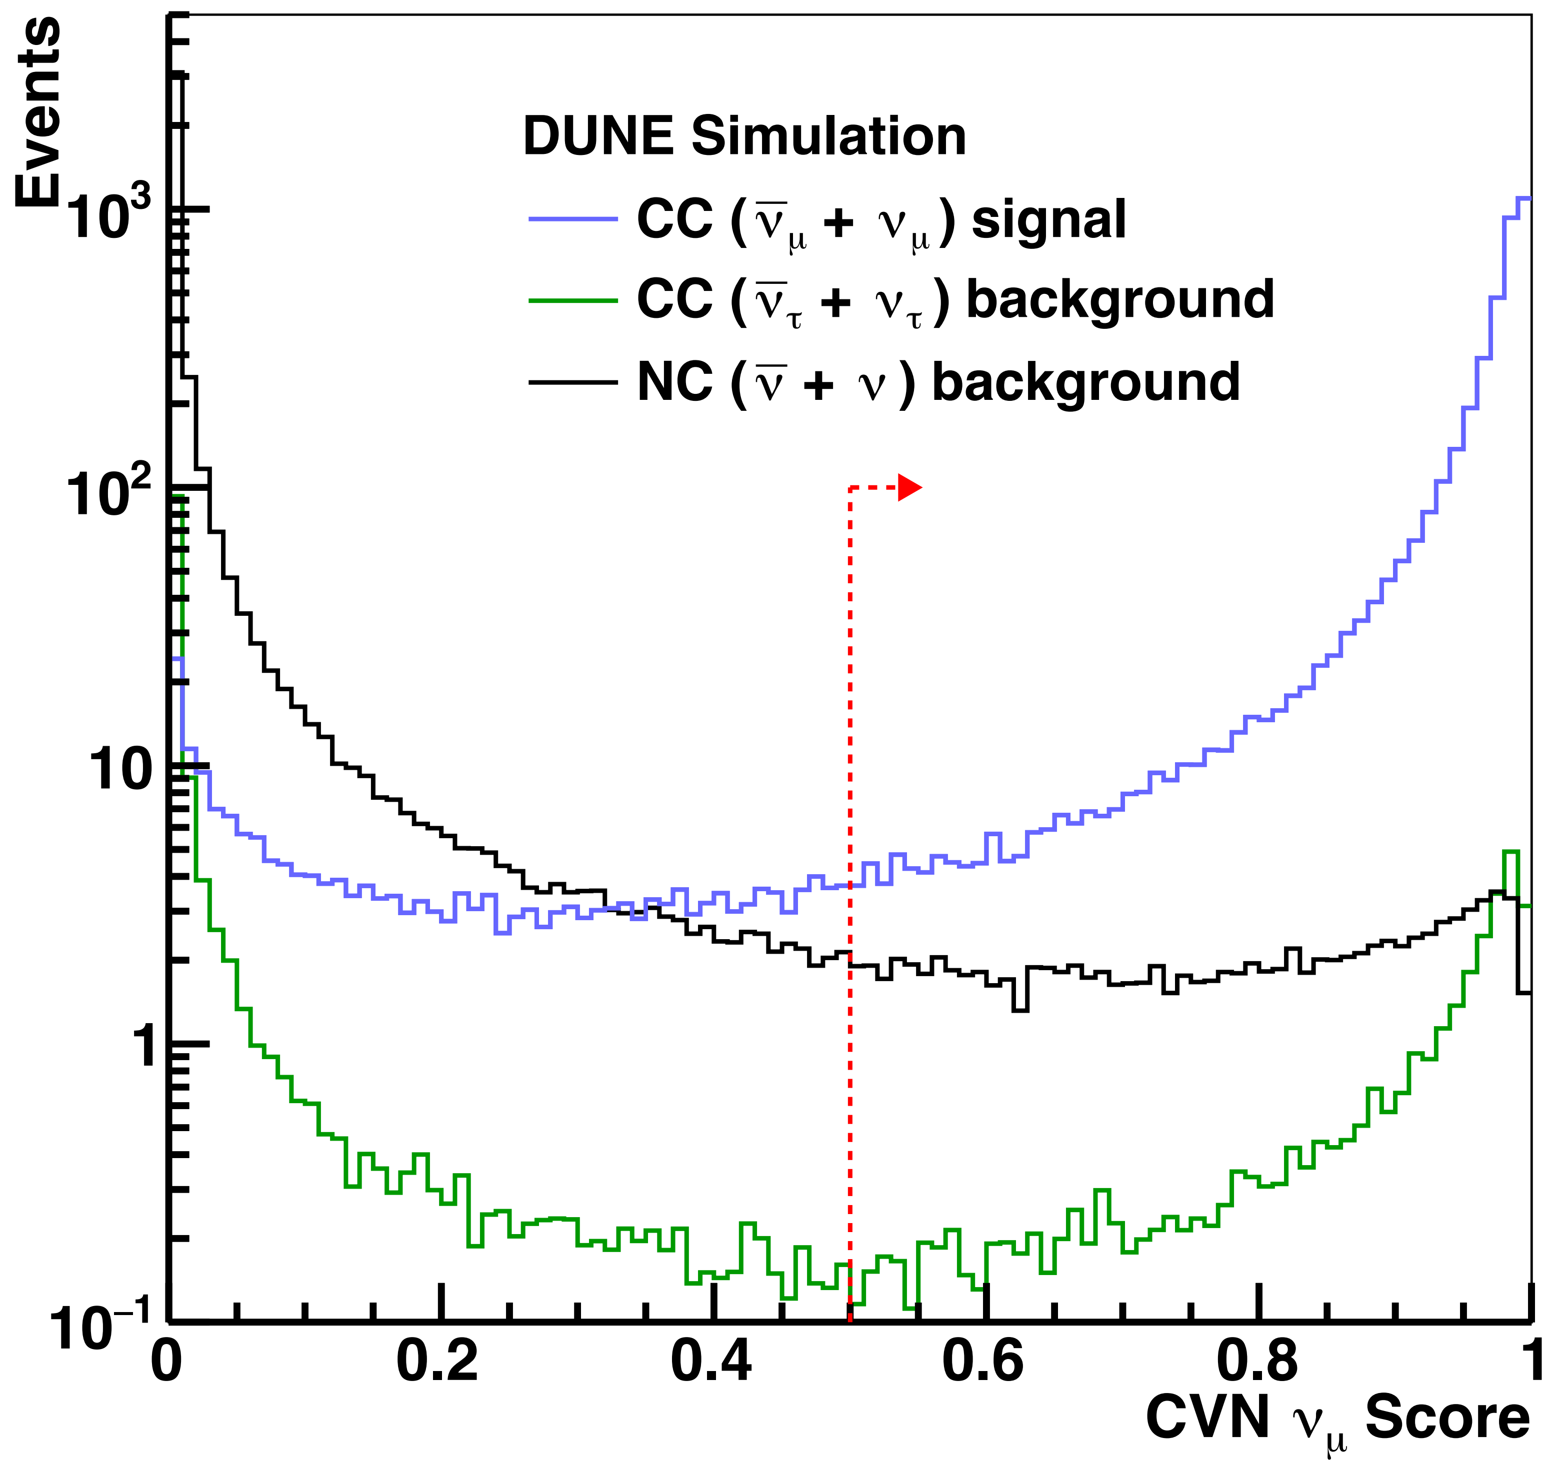

Arrows show events selected for the CC ν_μ disappearance sample

Figures reproduced from: DUNE Collaboration, *Neutrino interaction classification with a convolutional neural network in the DUNE far detector*, Phys. Rev. D **102** 9, 092003 (2020)

DUNE selection efficiencies

- We obtain highly efficiency analyses from the CVN event selection

Neutrino mode

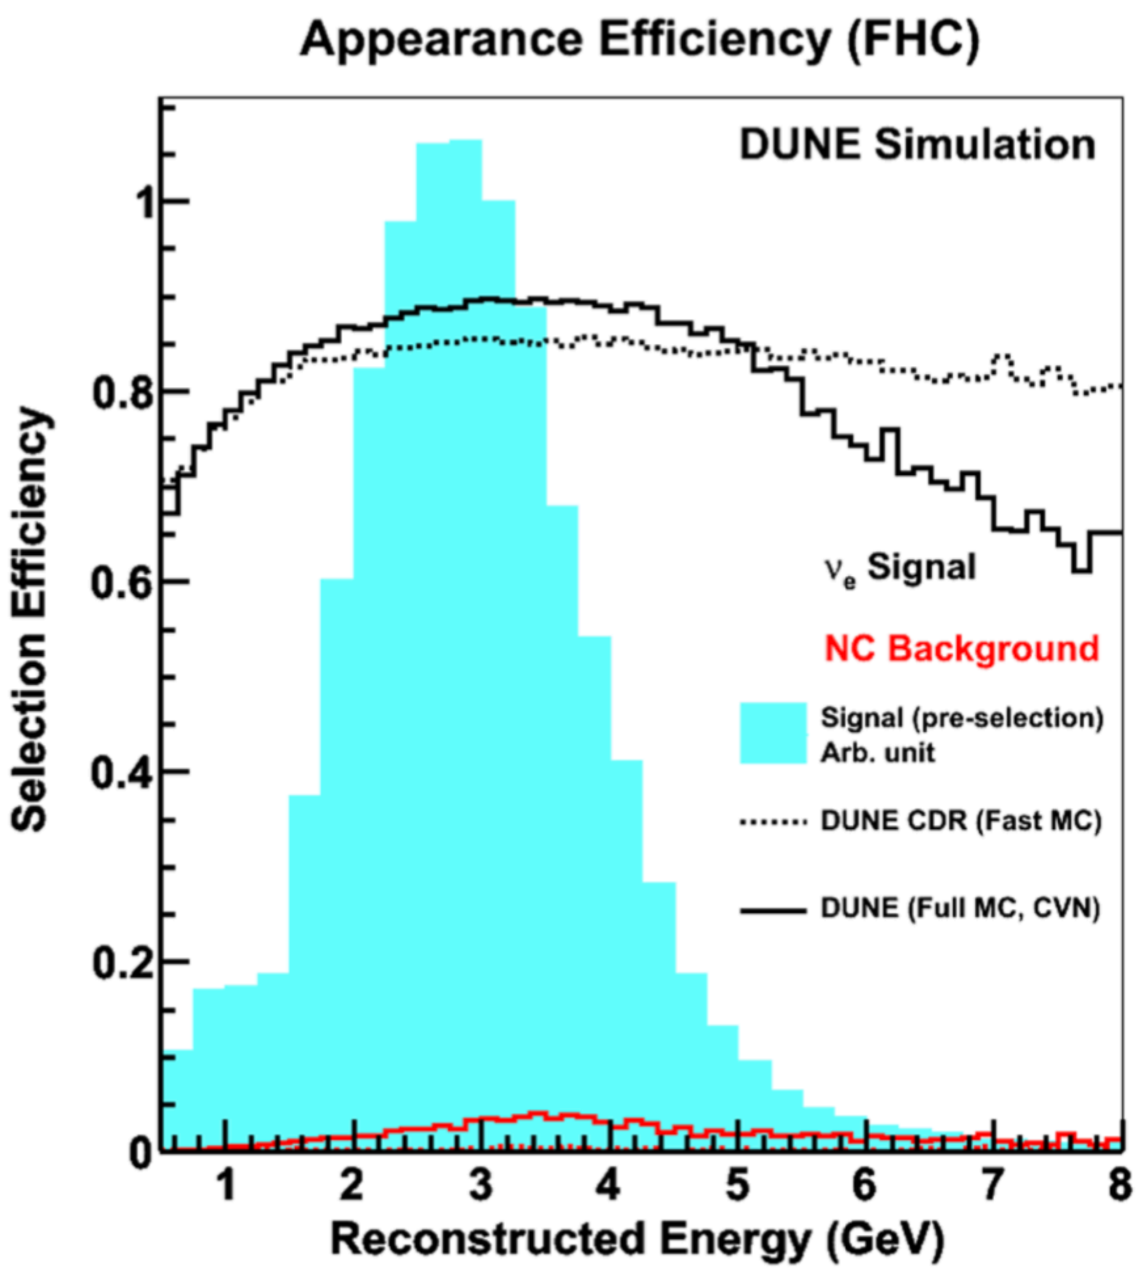

Antineutrino mode

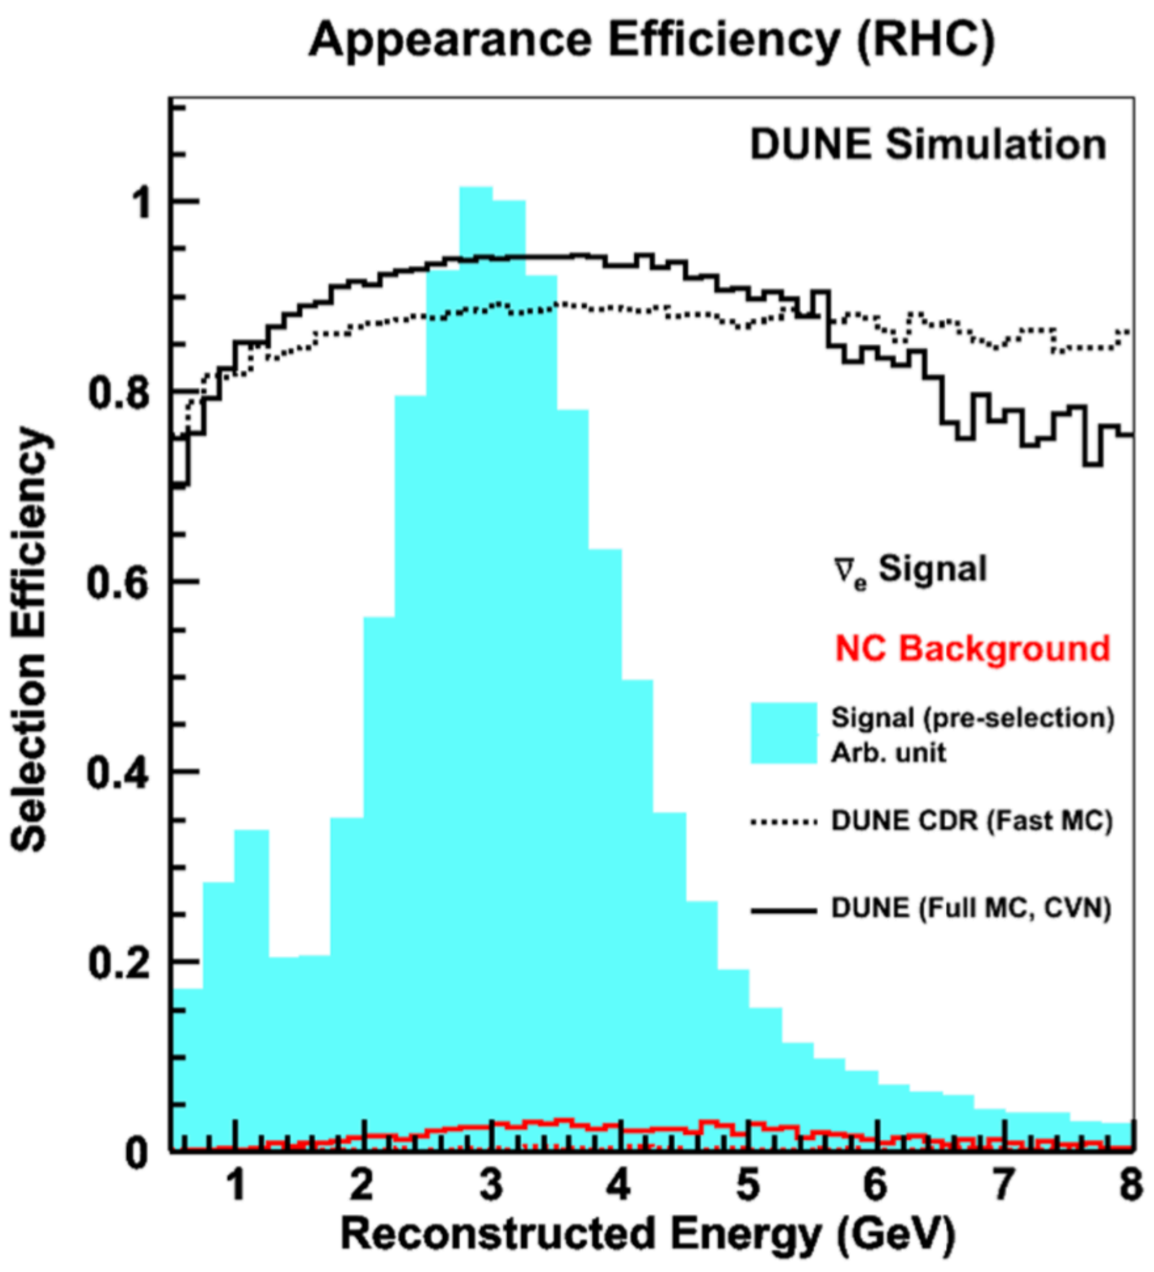

Neutrino mode

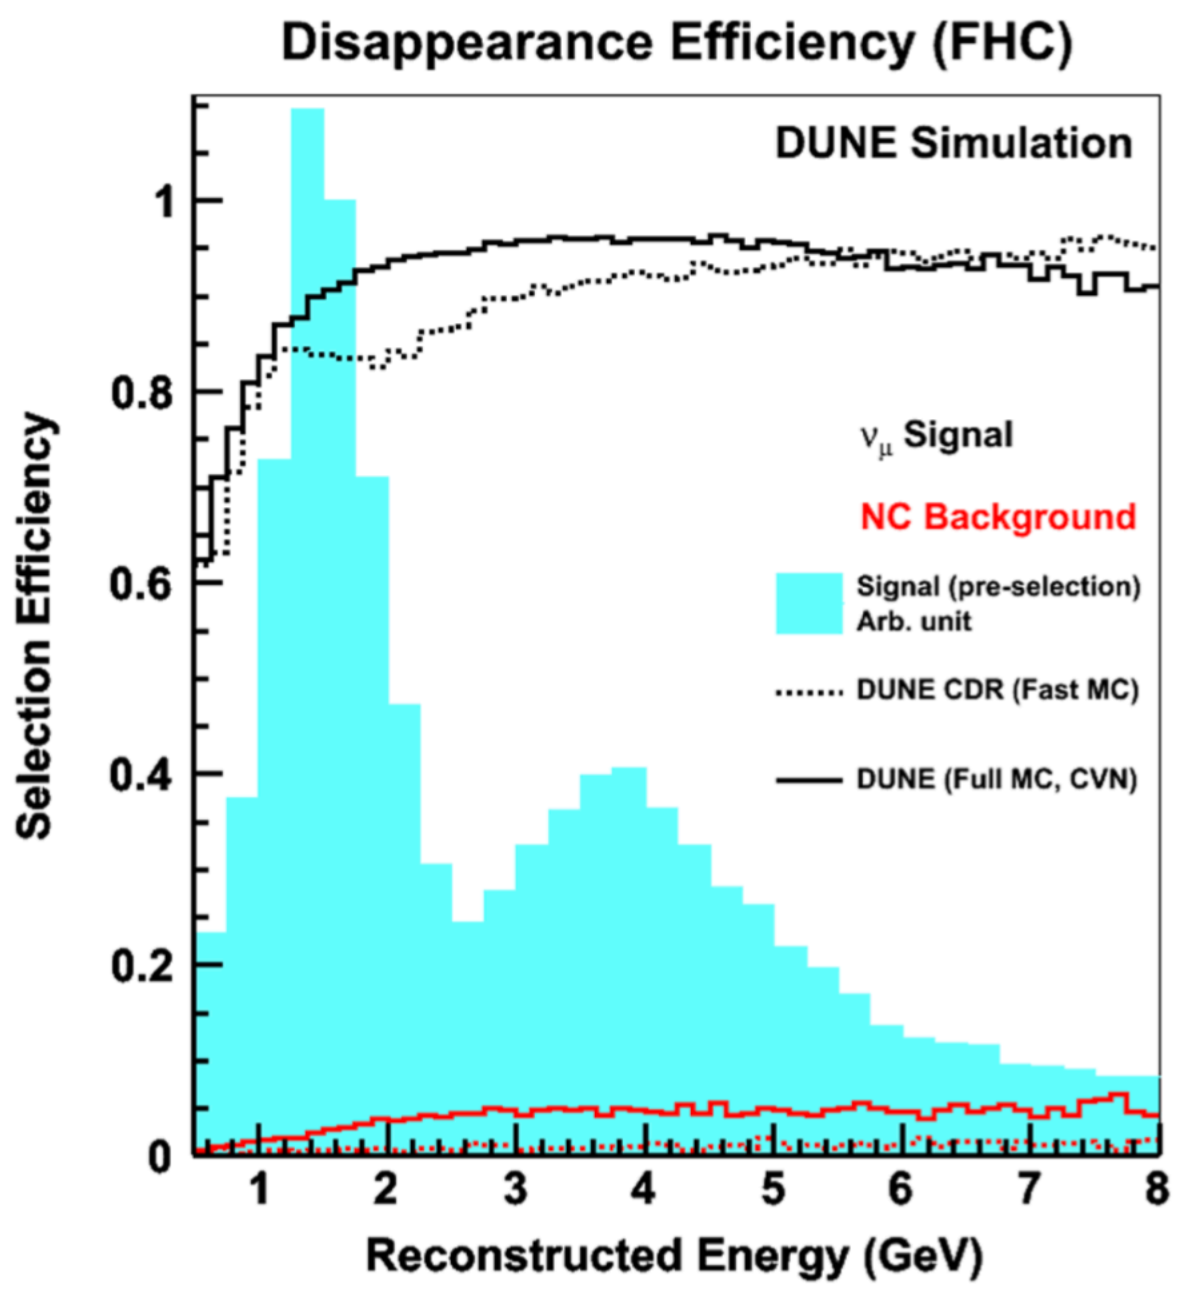

Antineutrino mode

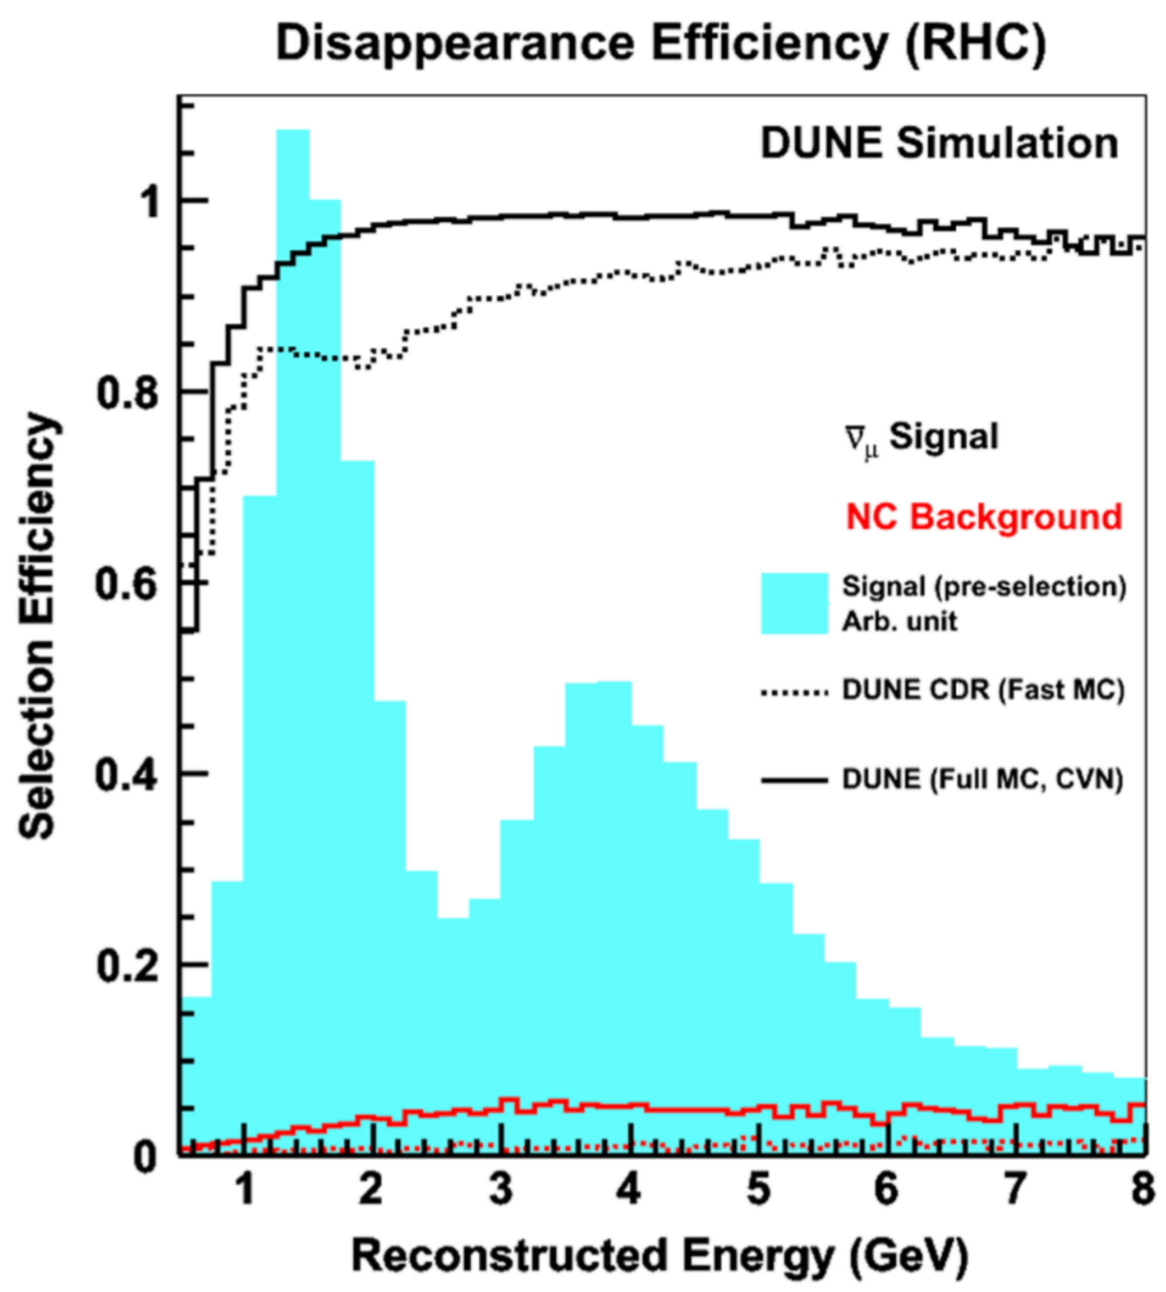

Efficiency for selecting CC ν_e interactions

Efficiency for selecting CC ν_μ interactions

Figures reproduced from: DUNE Collaboration, *Neutrino interaction classification with a convolutional neural network in the DUNE far detector,* Phys. Rev. D **102** 9, 092003 (2020)

DUNE CVN - Particle counting

- We tested some of the particle counting outputs
  - Proof of principle of using the CVN for exclusive final state selections
- Multiply together different scores:
  - CC ν_μ, 1p, $0\pi^\pm$, $0\pi^0$
  - NC, 0p, $0\pi^\pm$, $1\pi^0$

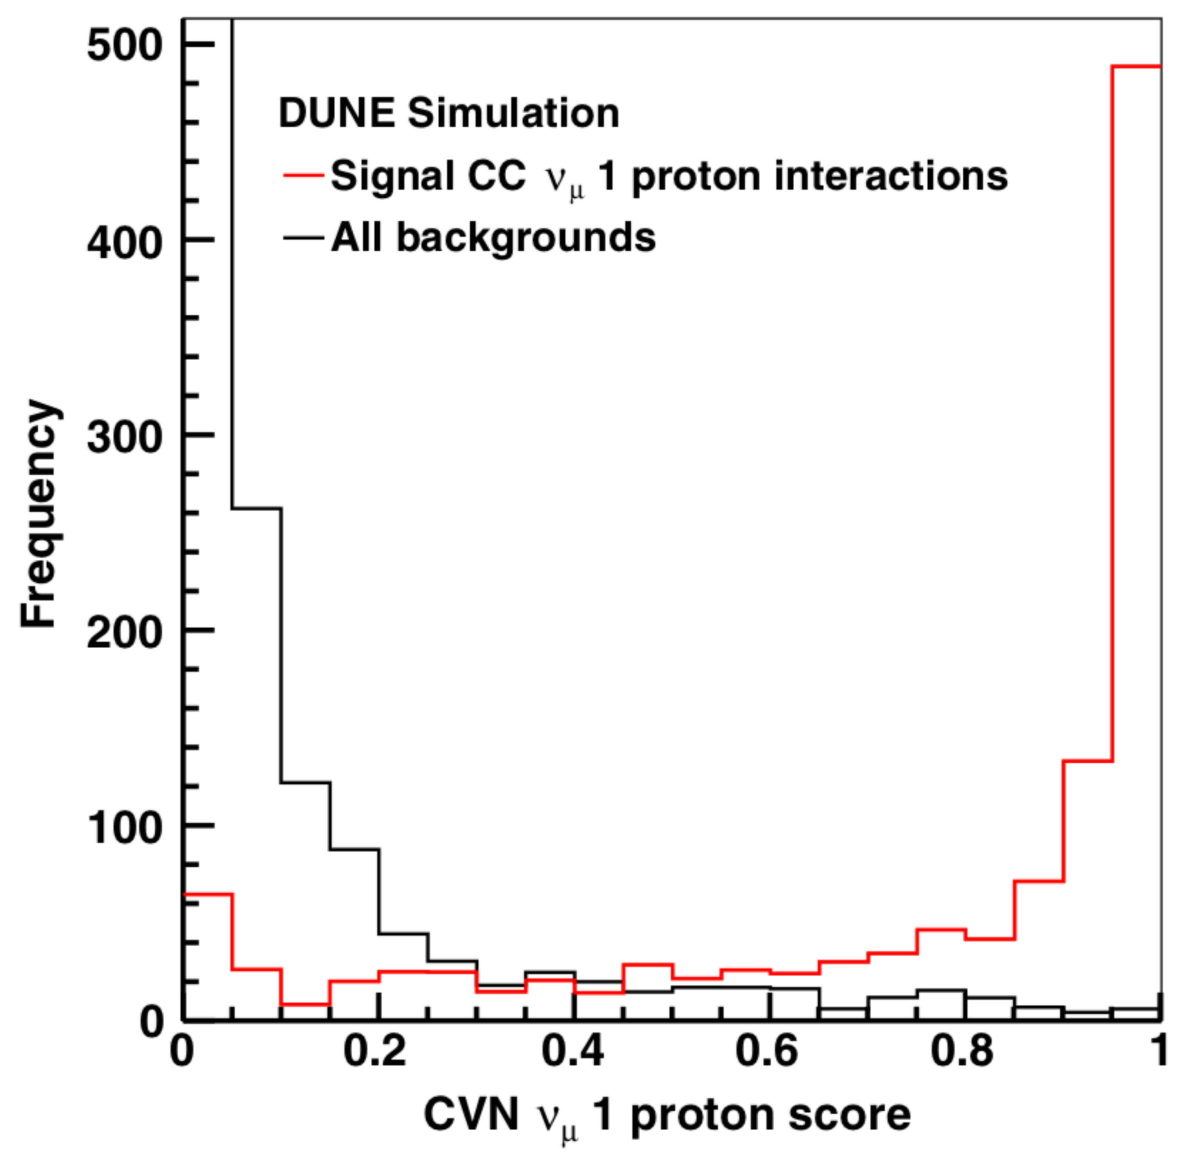

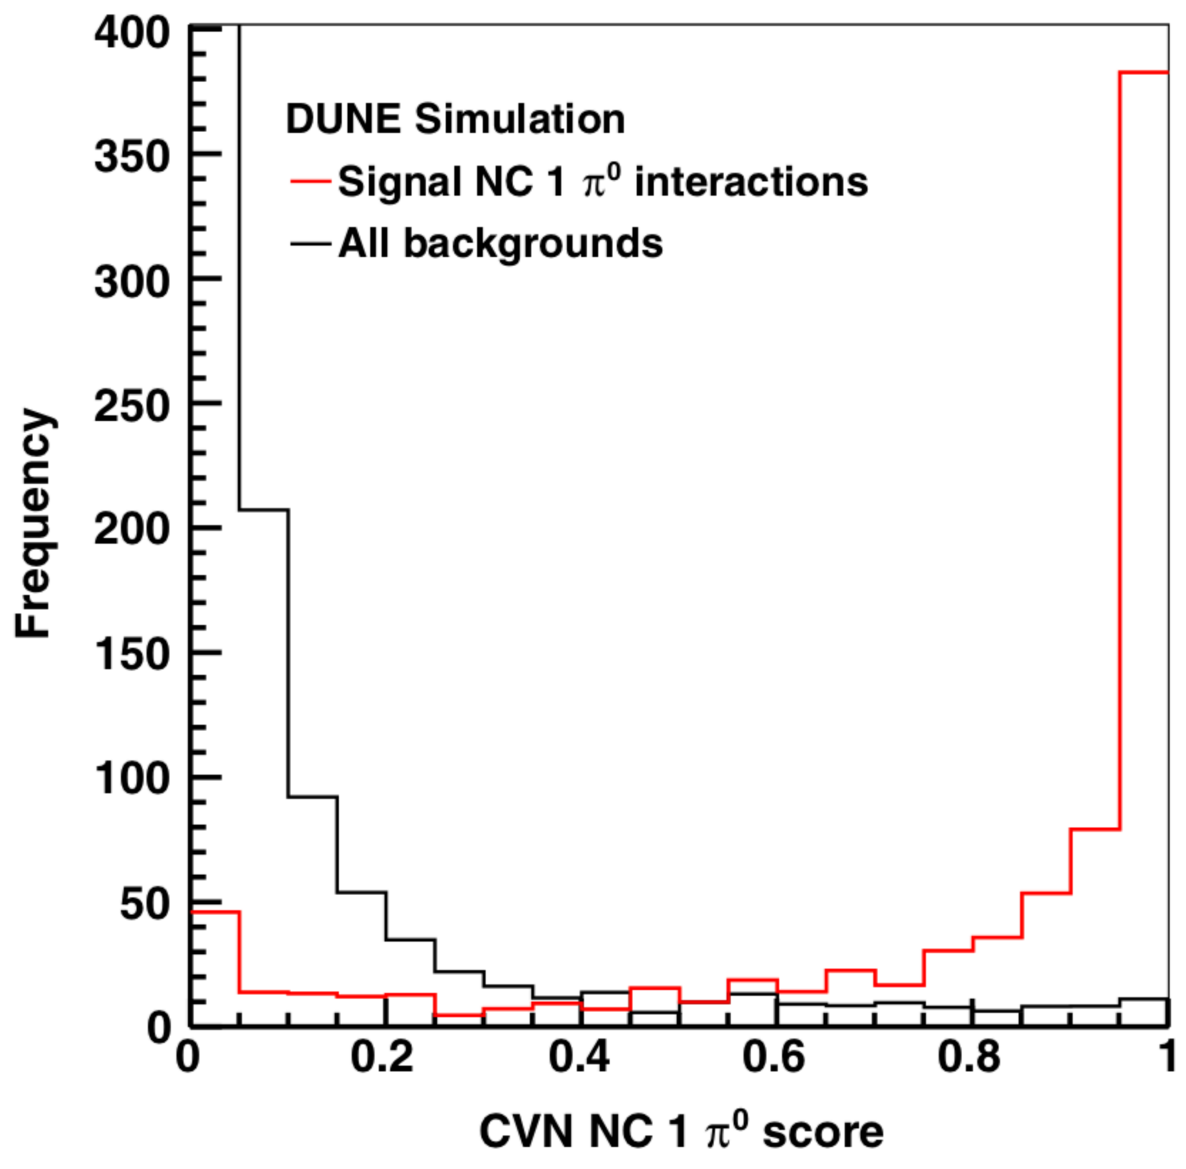

- Clearly these would need to be strongly validated before use on data
  - Much more likely to be biased by the choice of event generator

Figures reproduced from: DUNE Collaboration, *Neutrino interaction classification with a convolutional neural network in the DUNE far detector*, Phys. Rev. D **102** 9, 092003 (2020)

CNNs for Pixel Classification

Pixel classification

- Semantic Segmentation is the standard method used for pixel classification
  - It can also be used for instance segmentation
  - Example from MicroBooNE[1] (uses a sparse CNN), U-Net architecture
- Small-patch classification
  - ProtoDUNE small-patch network[2]

[1] P. Abratenko *et* al., Semantic segmentation with a sparse convolutional neural network for event reconstruction in MicroBooNE, Phys. Rev. D **103**, 052012 (2021)
[2] A. Abed Abud, *et al.*, Reconstruction of interactions in the ProtoDUNE-SP detector with Pandora. *Eur. Phys. J. C* **83**, 618 (2023)

MicroBooNE semantic segmentation

- Aim: to classify hits (and hence particles) as one of five classes:
  - Minimum ionising particle, heavily ionising particle, shower, delta-ray or Michel electron

- Architecture: Sparse U-Res-Net

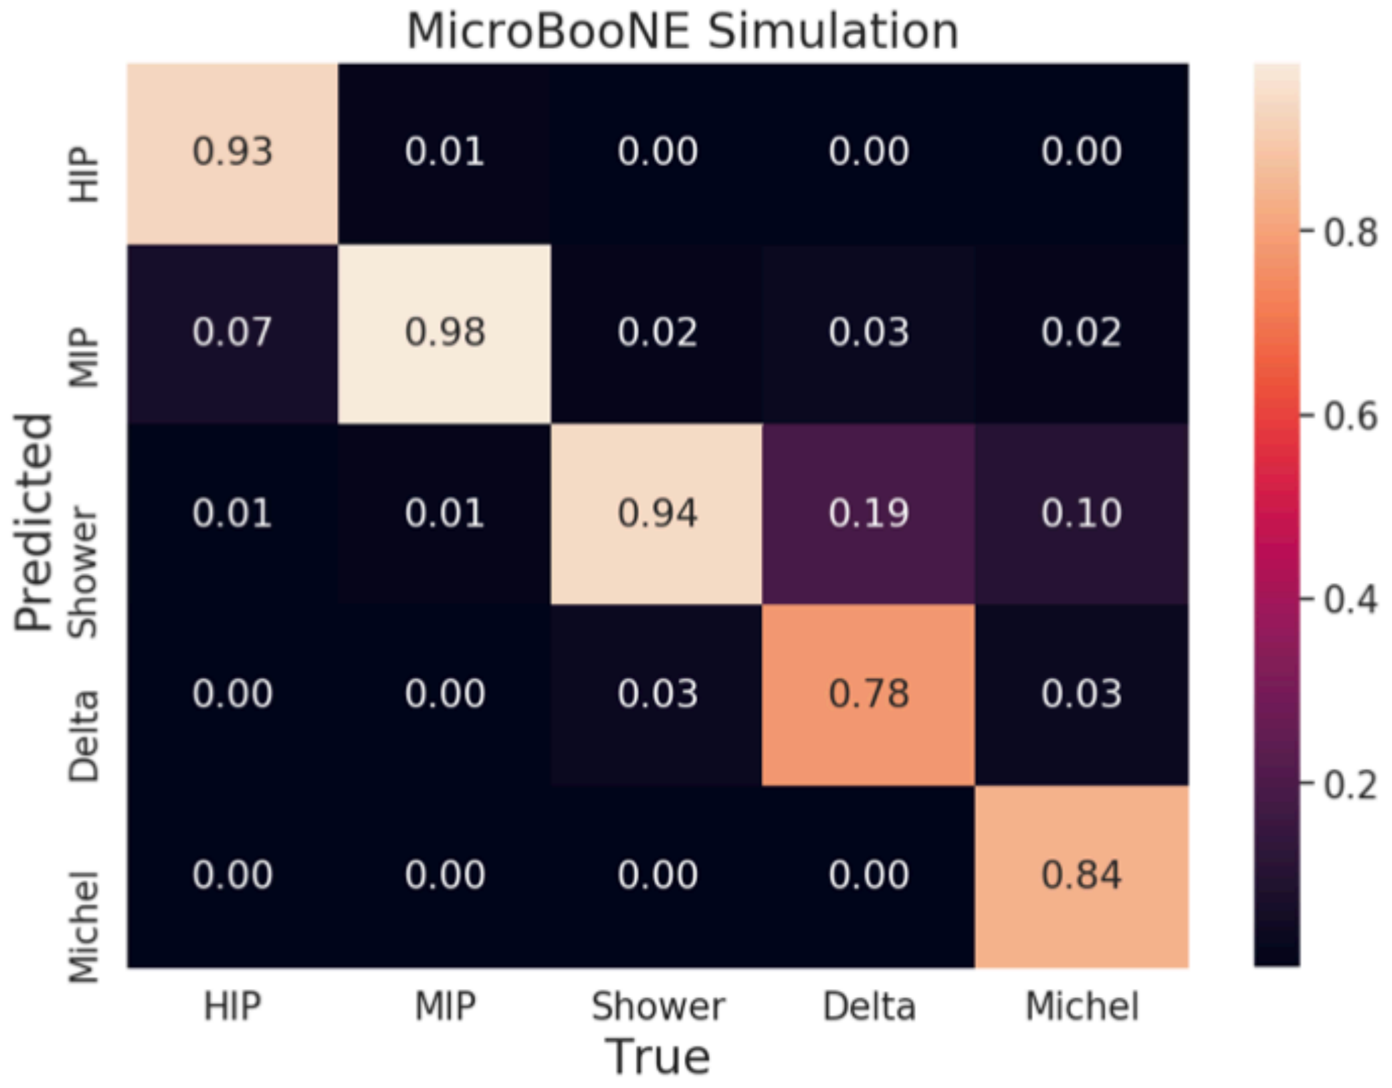

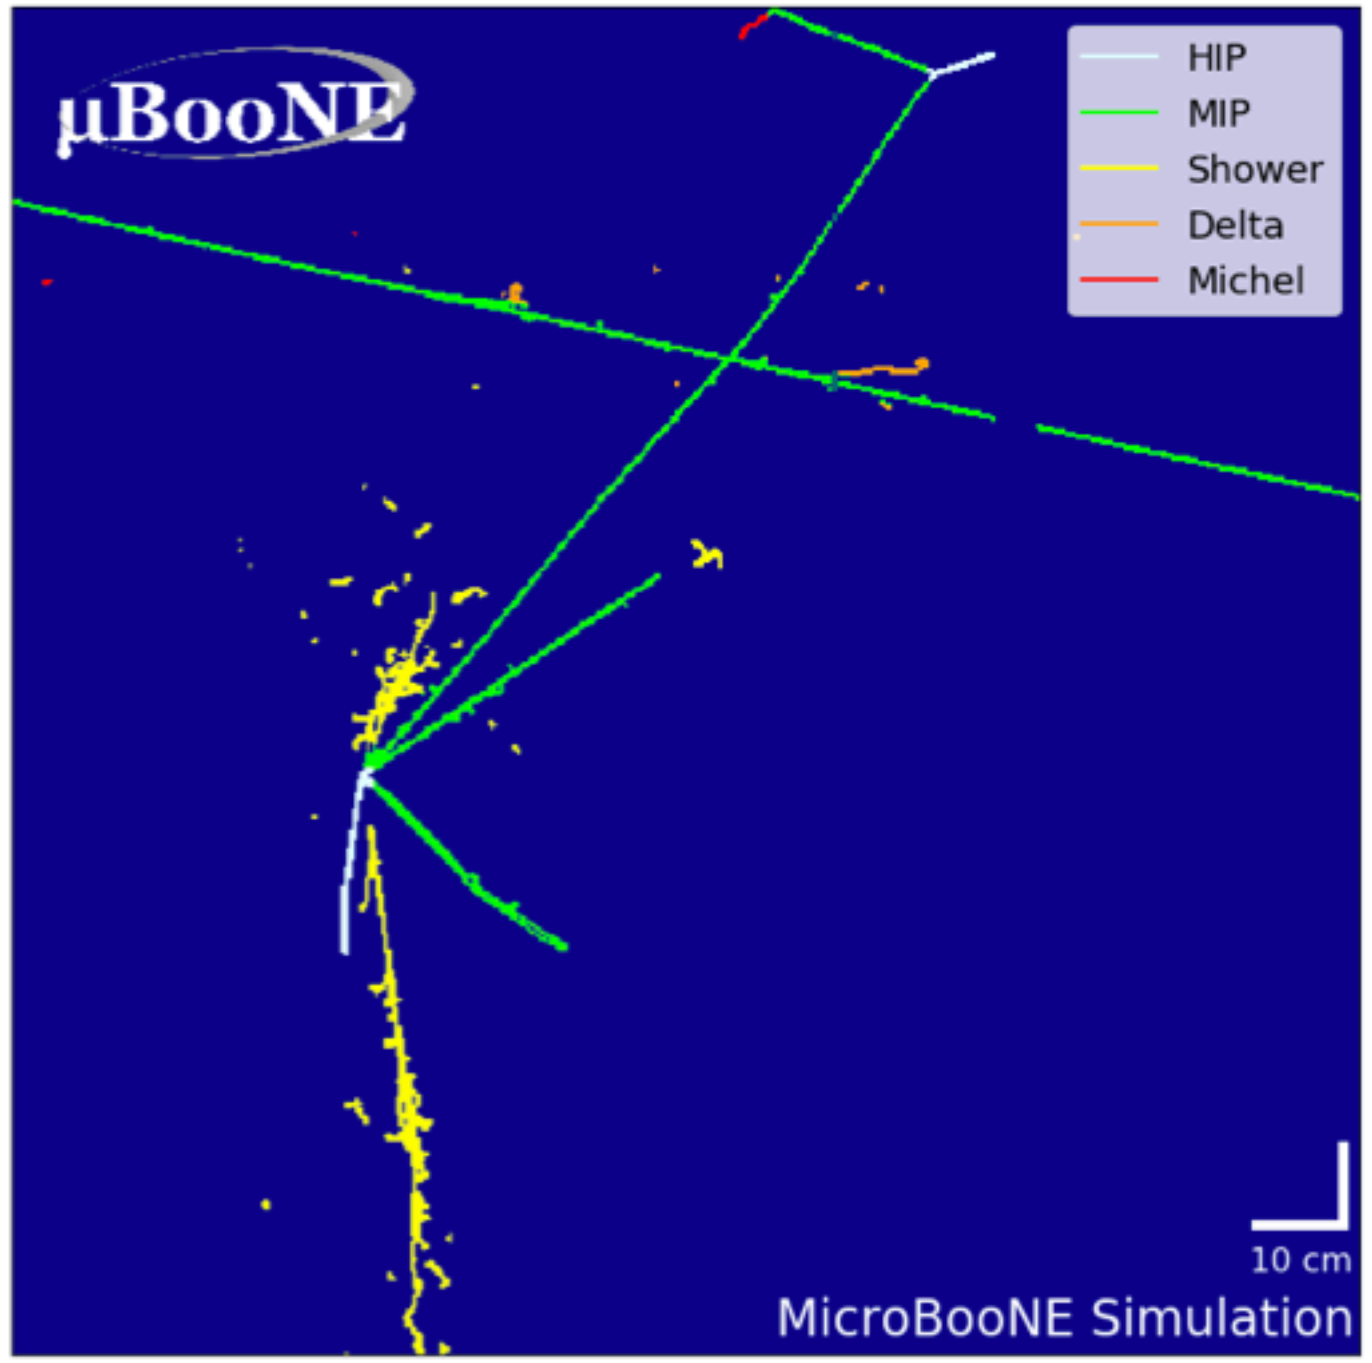

Figures reproduced from P. Abratenko *et* al., Semantic segmentation with a sparse convolutional neural network for event reconstruction in MicroBooNE, Phys. Rev. D **103**, 052012 (2021)

MicroBooNE semantic segmentation

- Benchmarked on simulation

- Some nice examples from data too

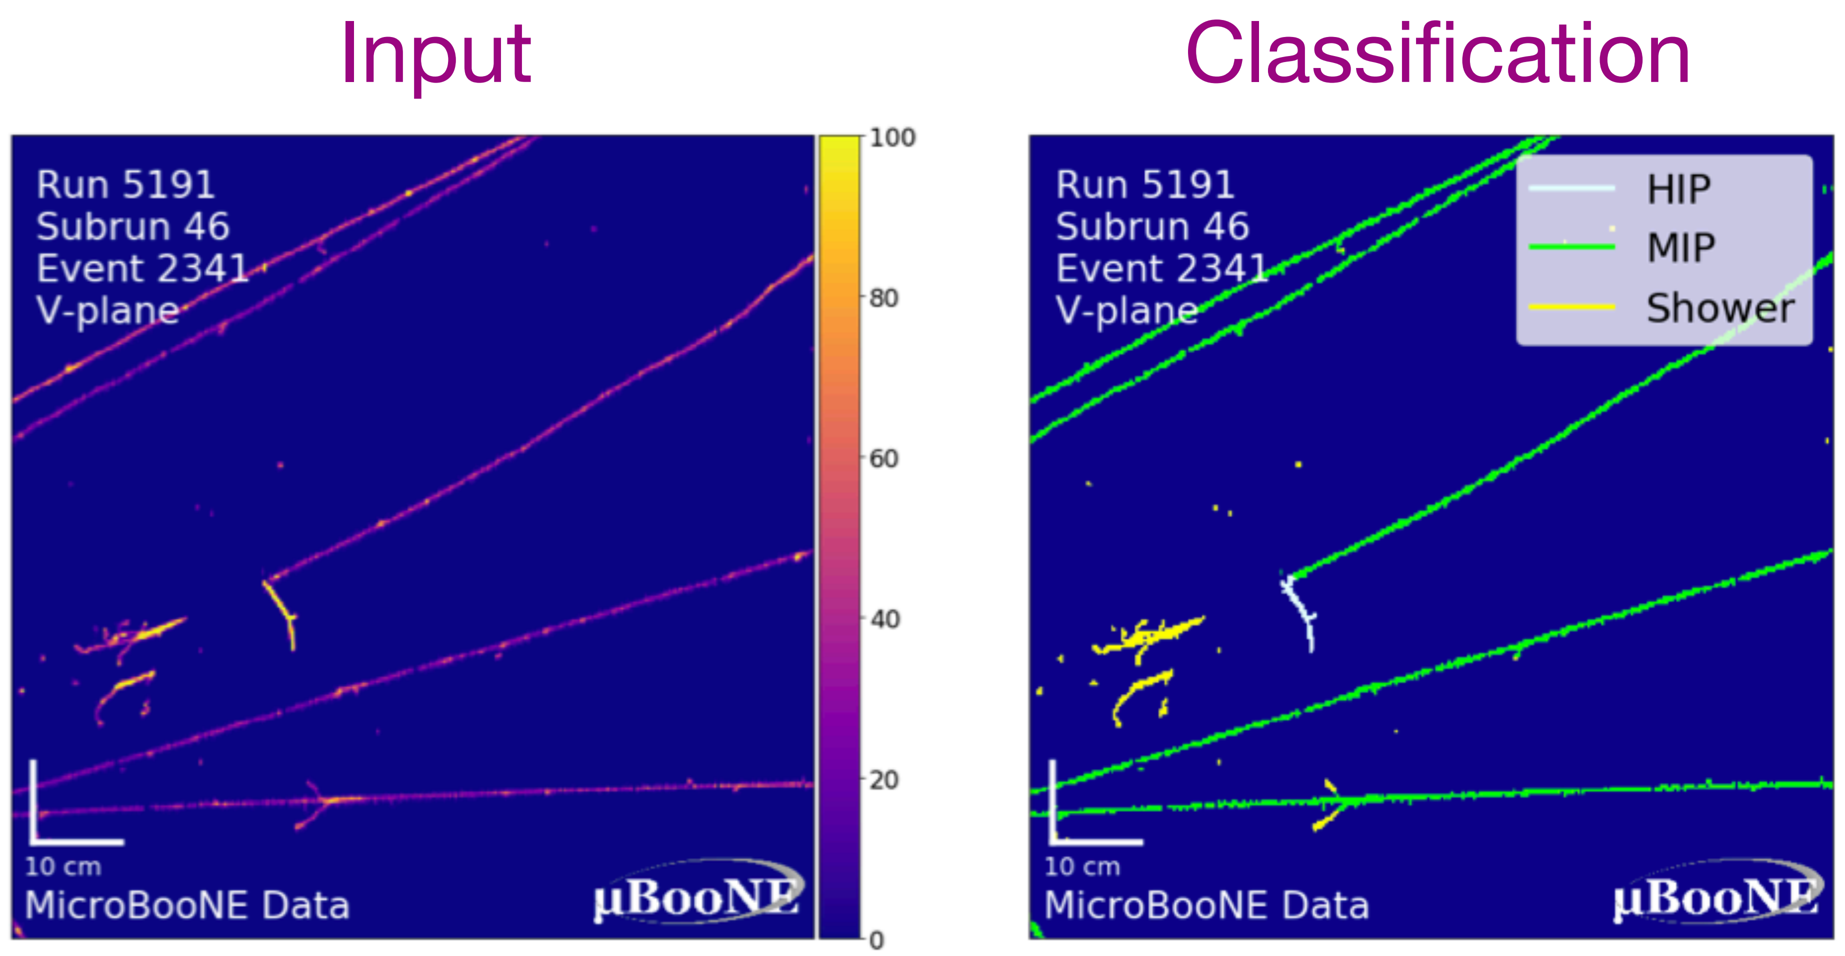

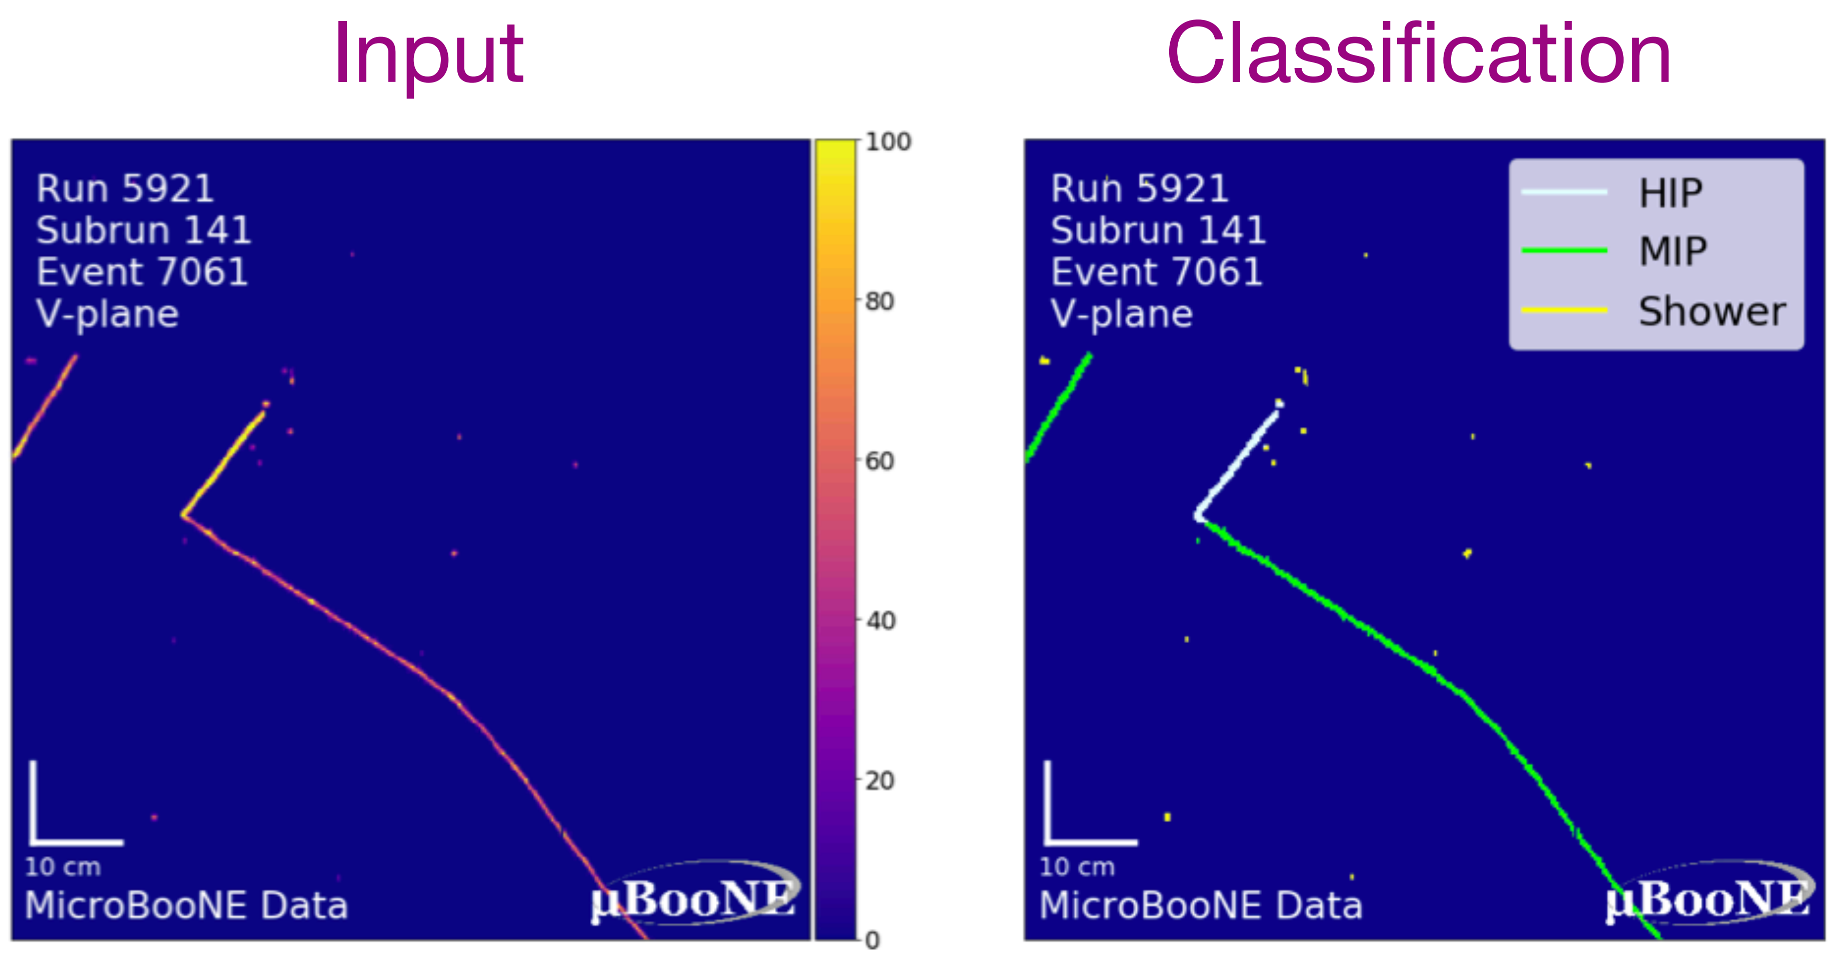

Figures reproduced from P. Abratenko *et al.*, Semantic segmentation with a sparse convolutional neural network for event reconstruction in MicroBooNE, Phys. Rev. D **103**, 052012 (2021)

ProtoDUNE-SP

- ProtoDUNE-SP was a large scale prototype for the DUNE FD
- It was located in a test-beam at CERN
  - It didn't collect neutrinos, but I include it here as a neutrino detector prototype
- Collected data from the test beam:
  - e^+, μ^+, π^+, p, K^+
  - Particles in the energy range 1 - 7 GeV/c
- Also exposed to a high rate of cosmic rays

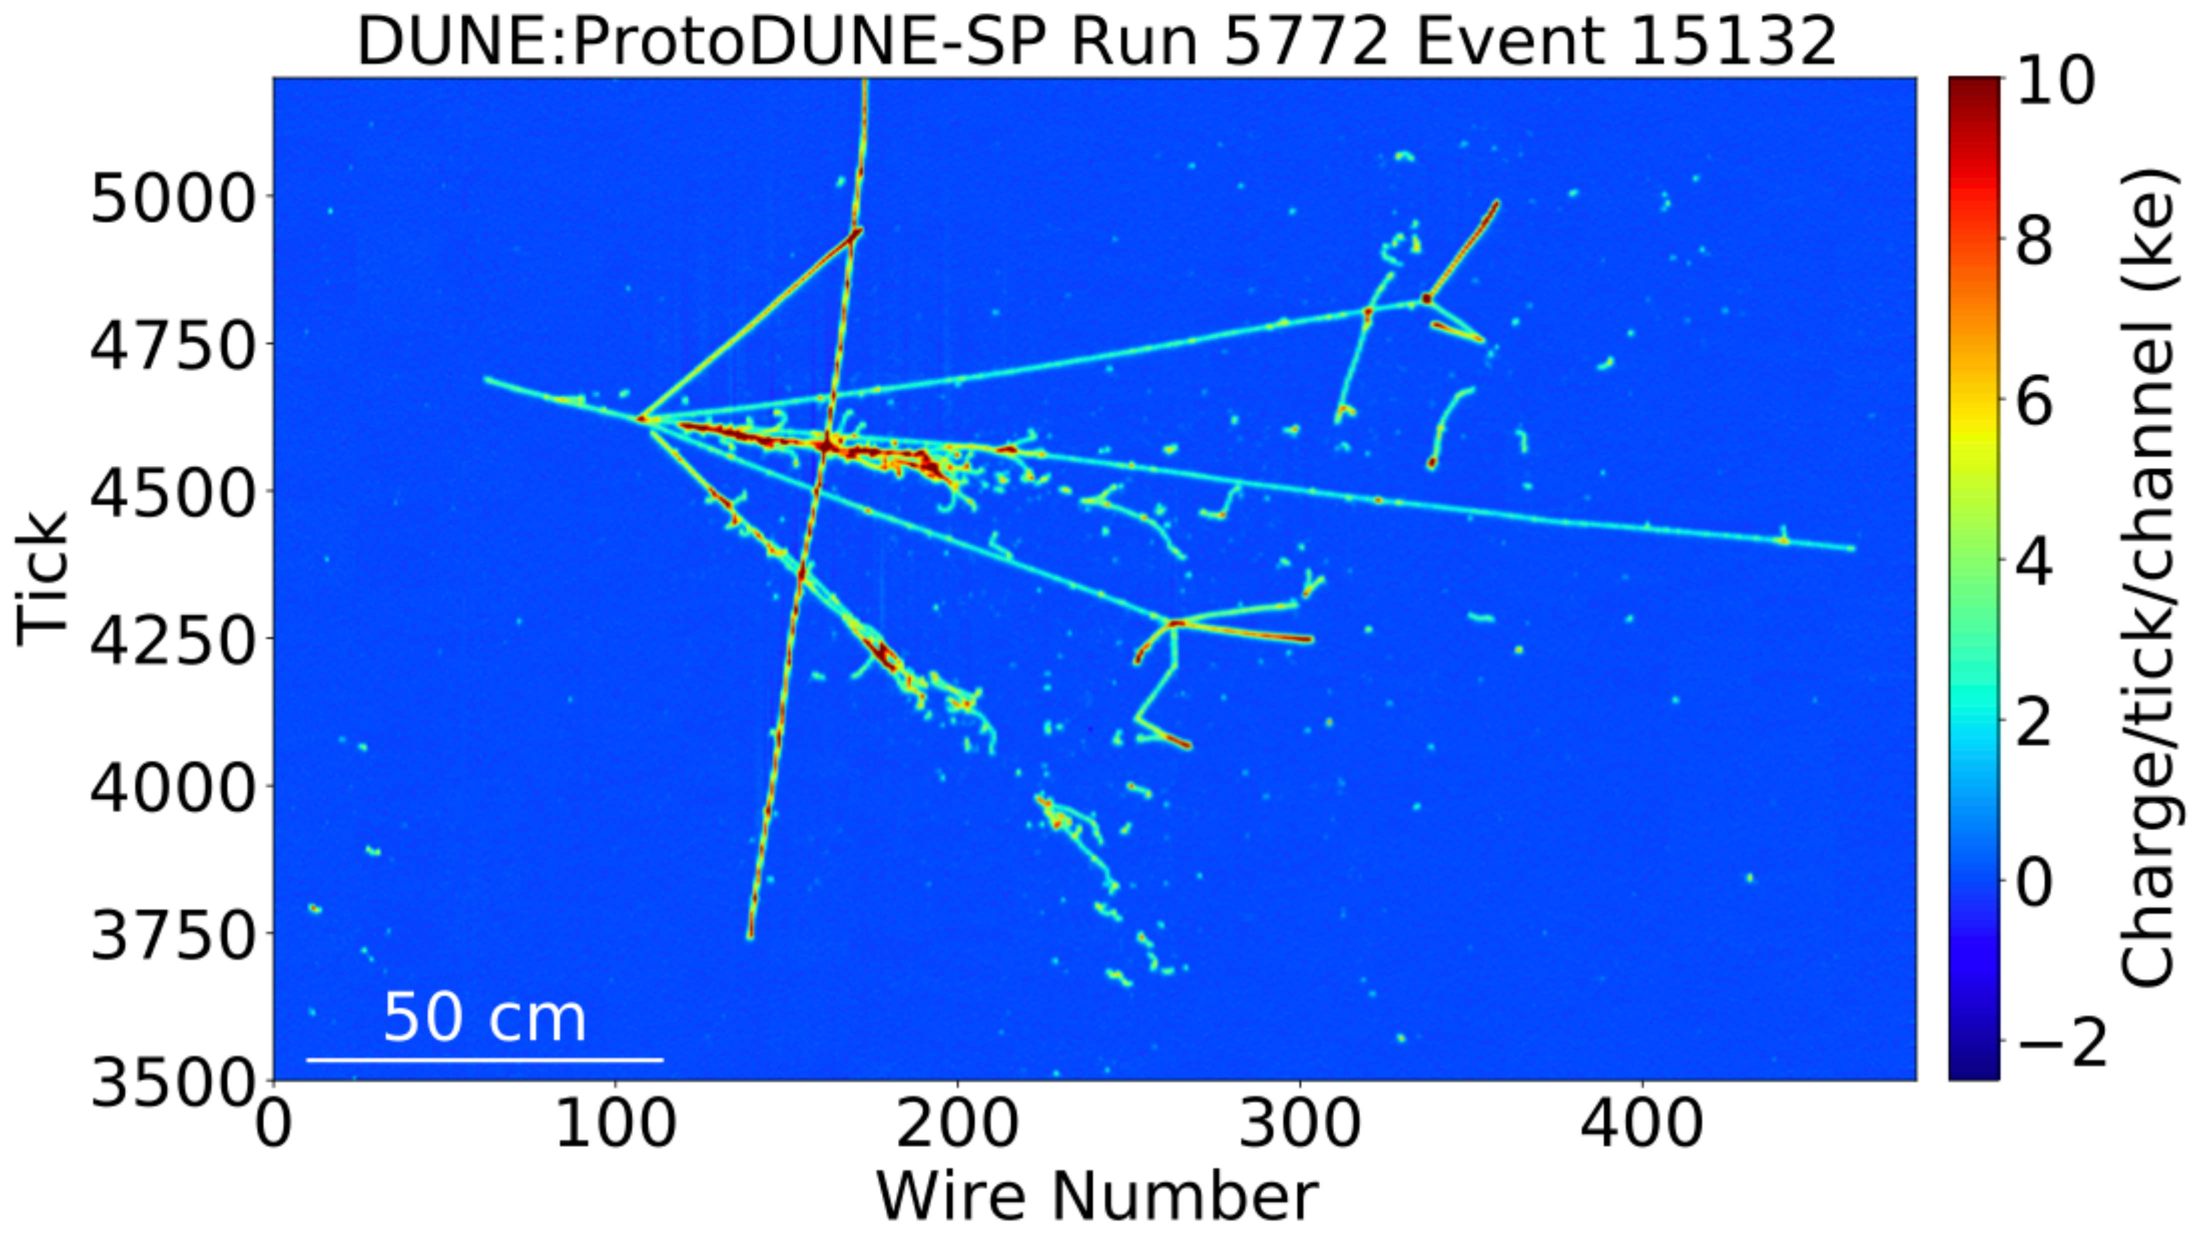

Figure reproduced from B. Abi, *et al.*, First results on ProtoDUNE-SP liquid argon time projection chamber performance from a beam test at the CERN Neutrino Platform, *JINST* **15** 12, P12004, (2020)

ProtoDUNE hit-tagging CNN

- Aim: to classify hits as either track-like or shower-like
  - Also, separately, if they are Michel-electron-like
- Architecture:
  - Very simple CNN with a single convolutional layer
  - Operates on small 48 x 48 patches of the events
- Very low memory usage and fast for CPU inference
  - Developed as an alternative to semantic segmentation

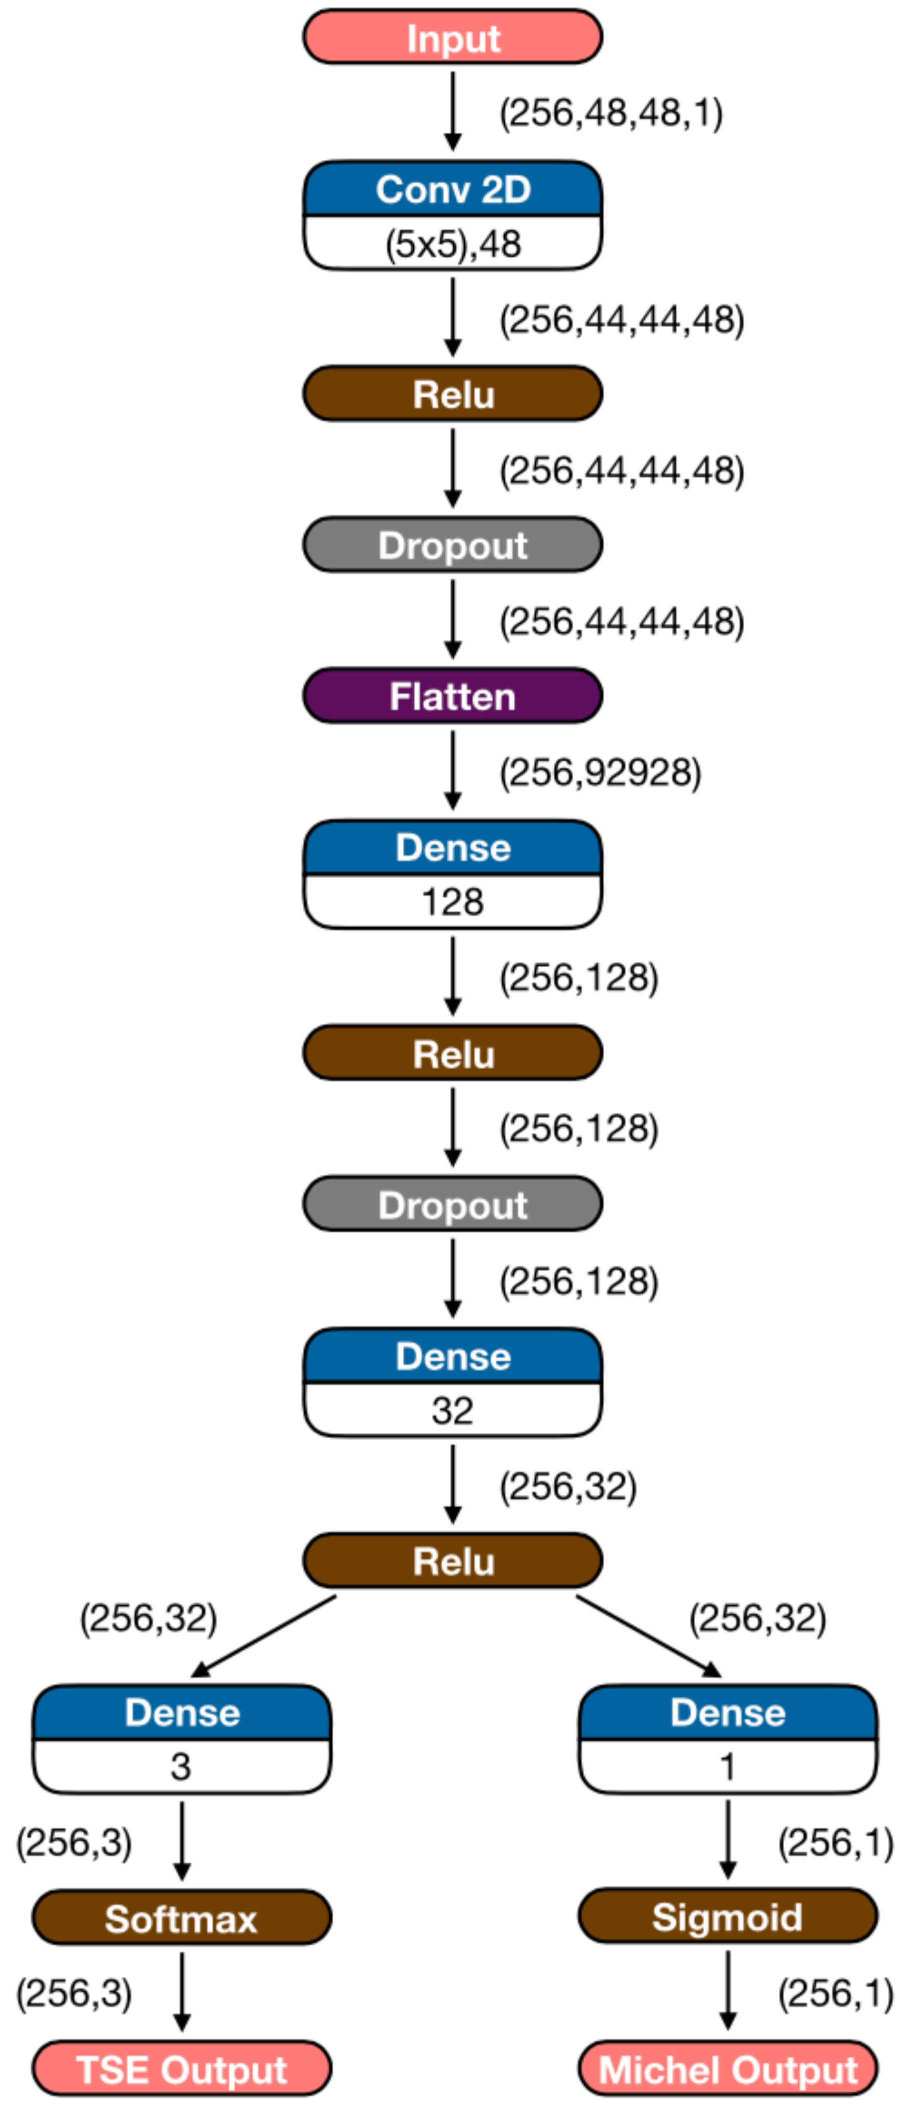

Figure reproduced from A. Abed Abud, *et al.*, Separation of track- and shower-like energy deposits in ProtoDUNE-SP using a convolutional neural network, *Eur.Phys.J.C* **82** 10, 903, (2022)

ProtoDUNE hit-tagging CNN

- It was tested at the hit-level and particle-level (by averaging scores per particle)

Cosmic-ray muon hits
(can include delta rays)

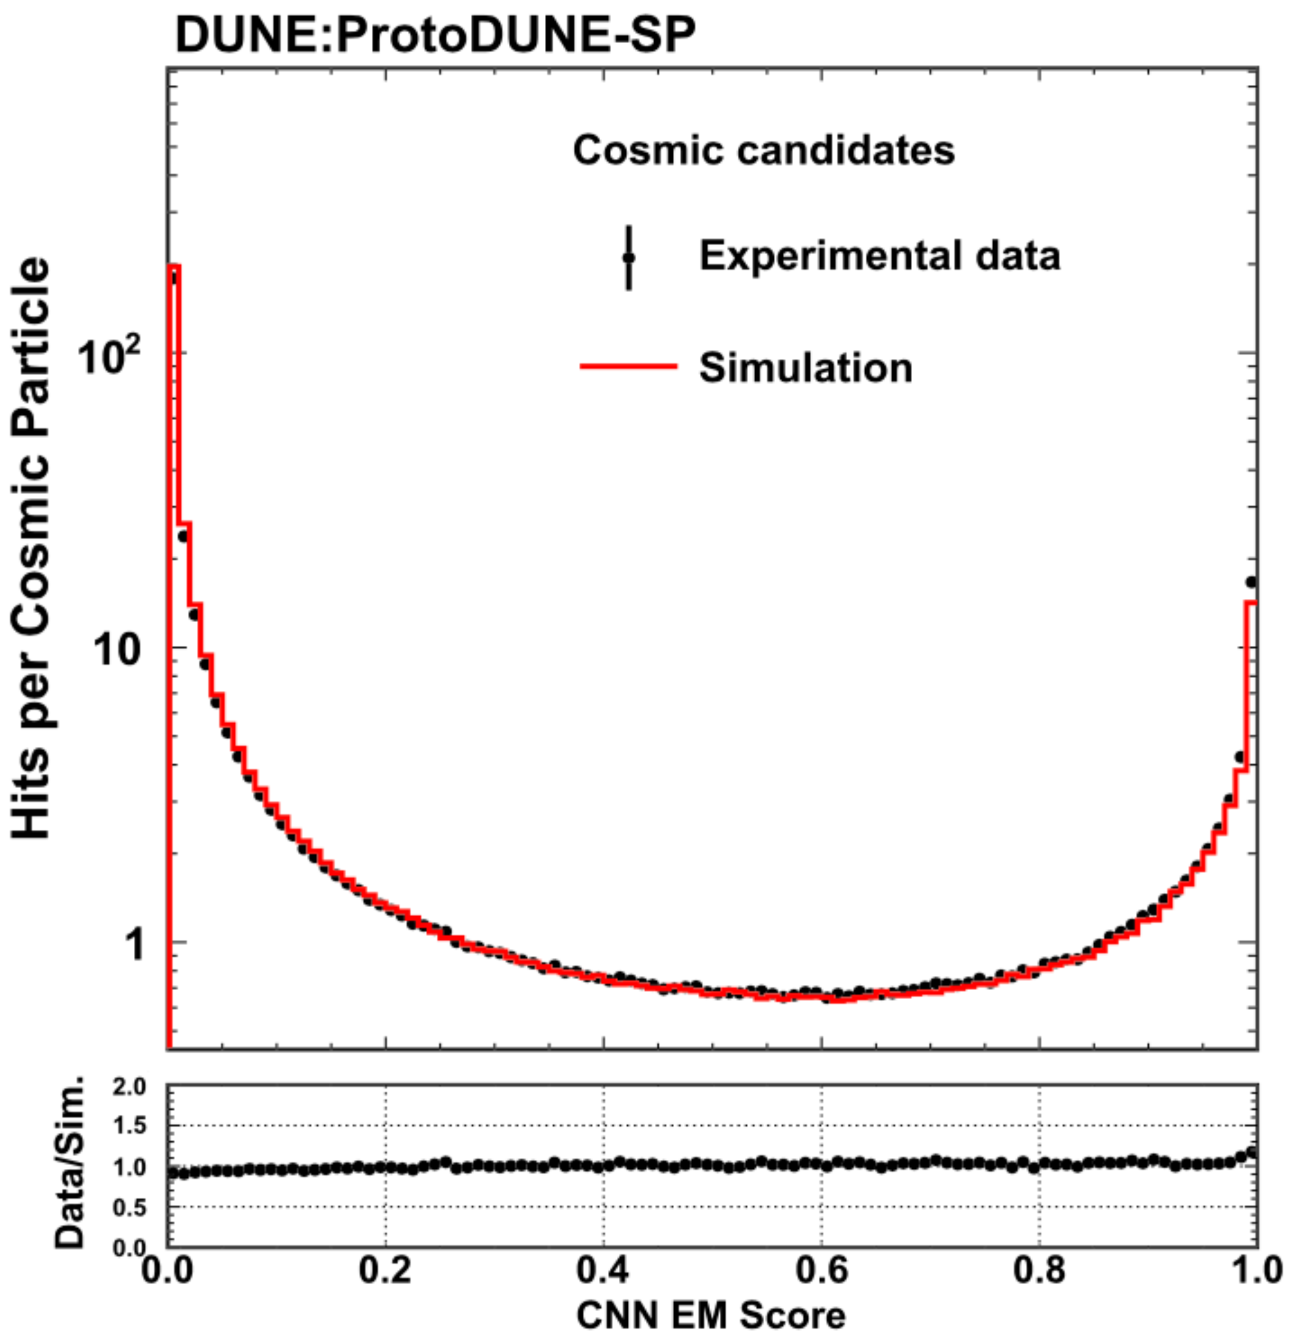

Example event display from data

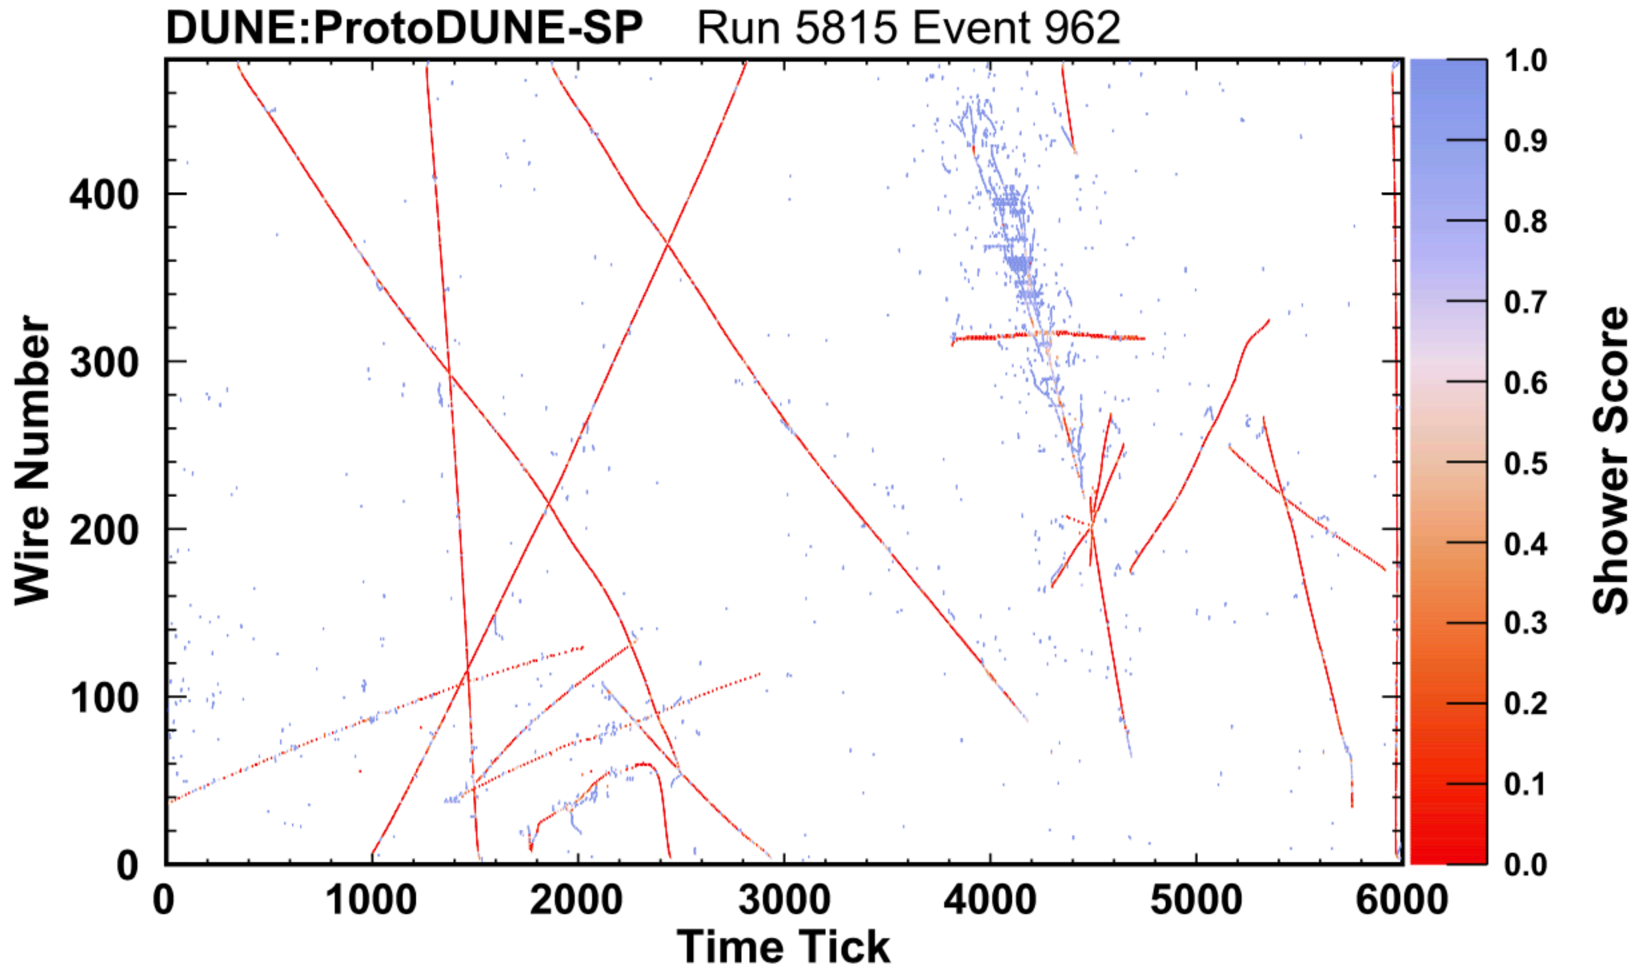

Test-beam particle scores

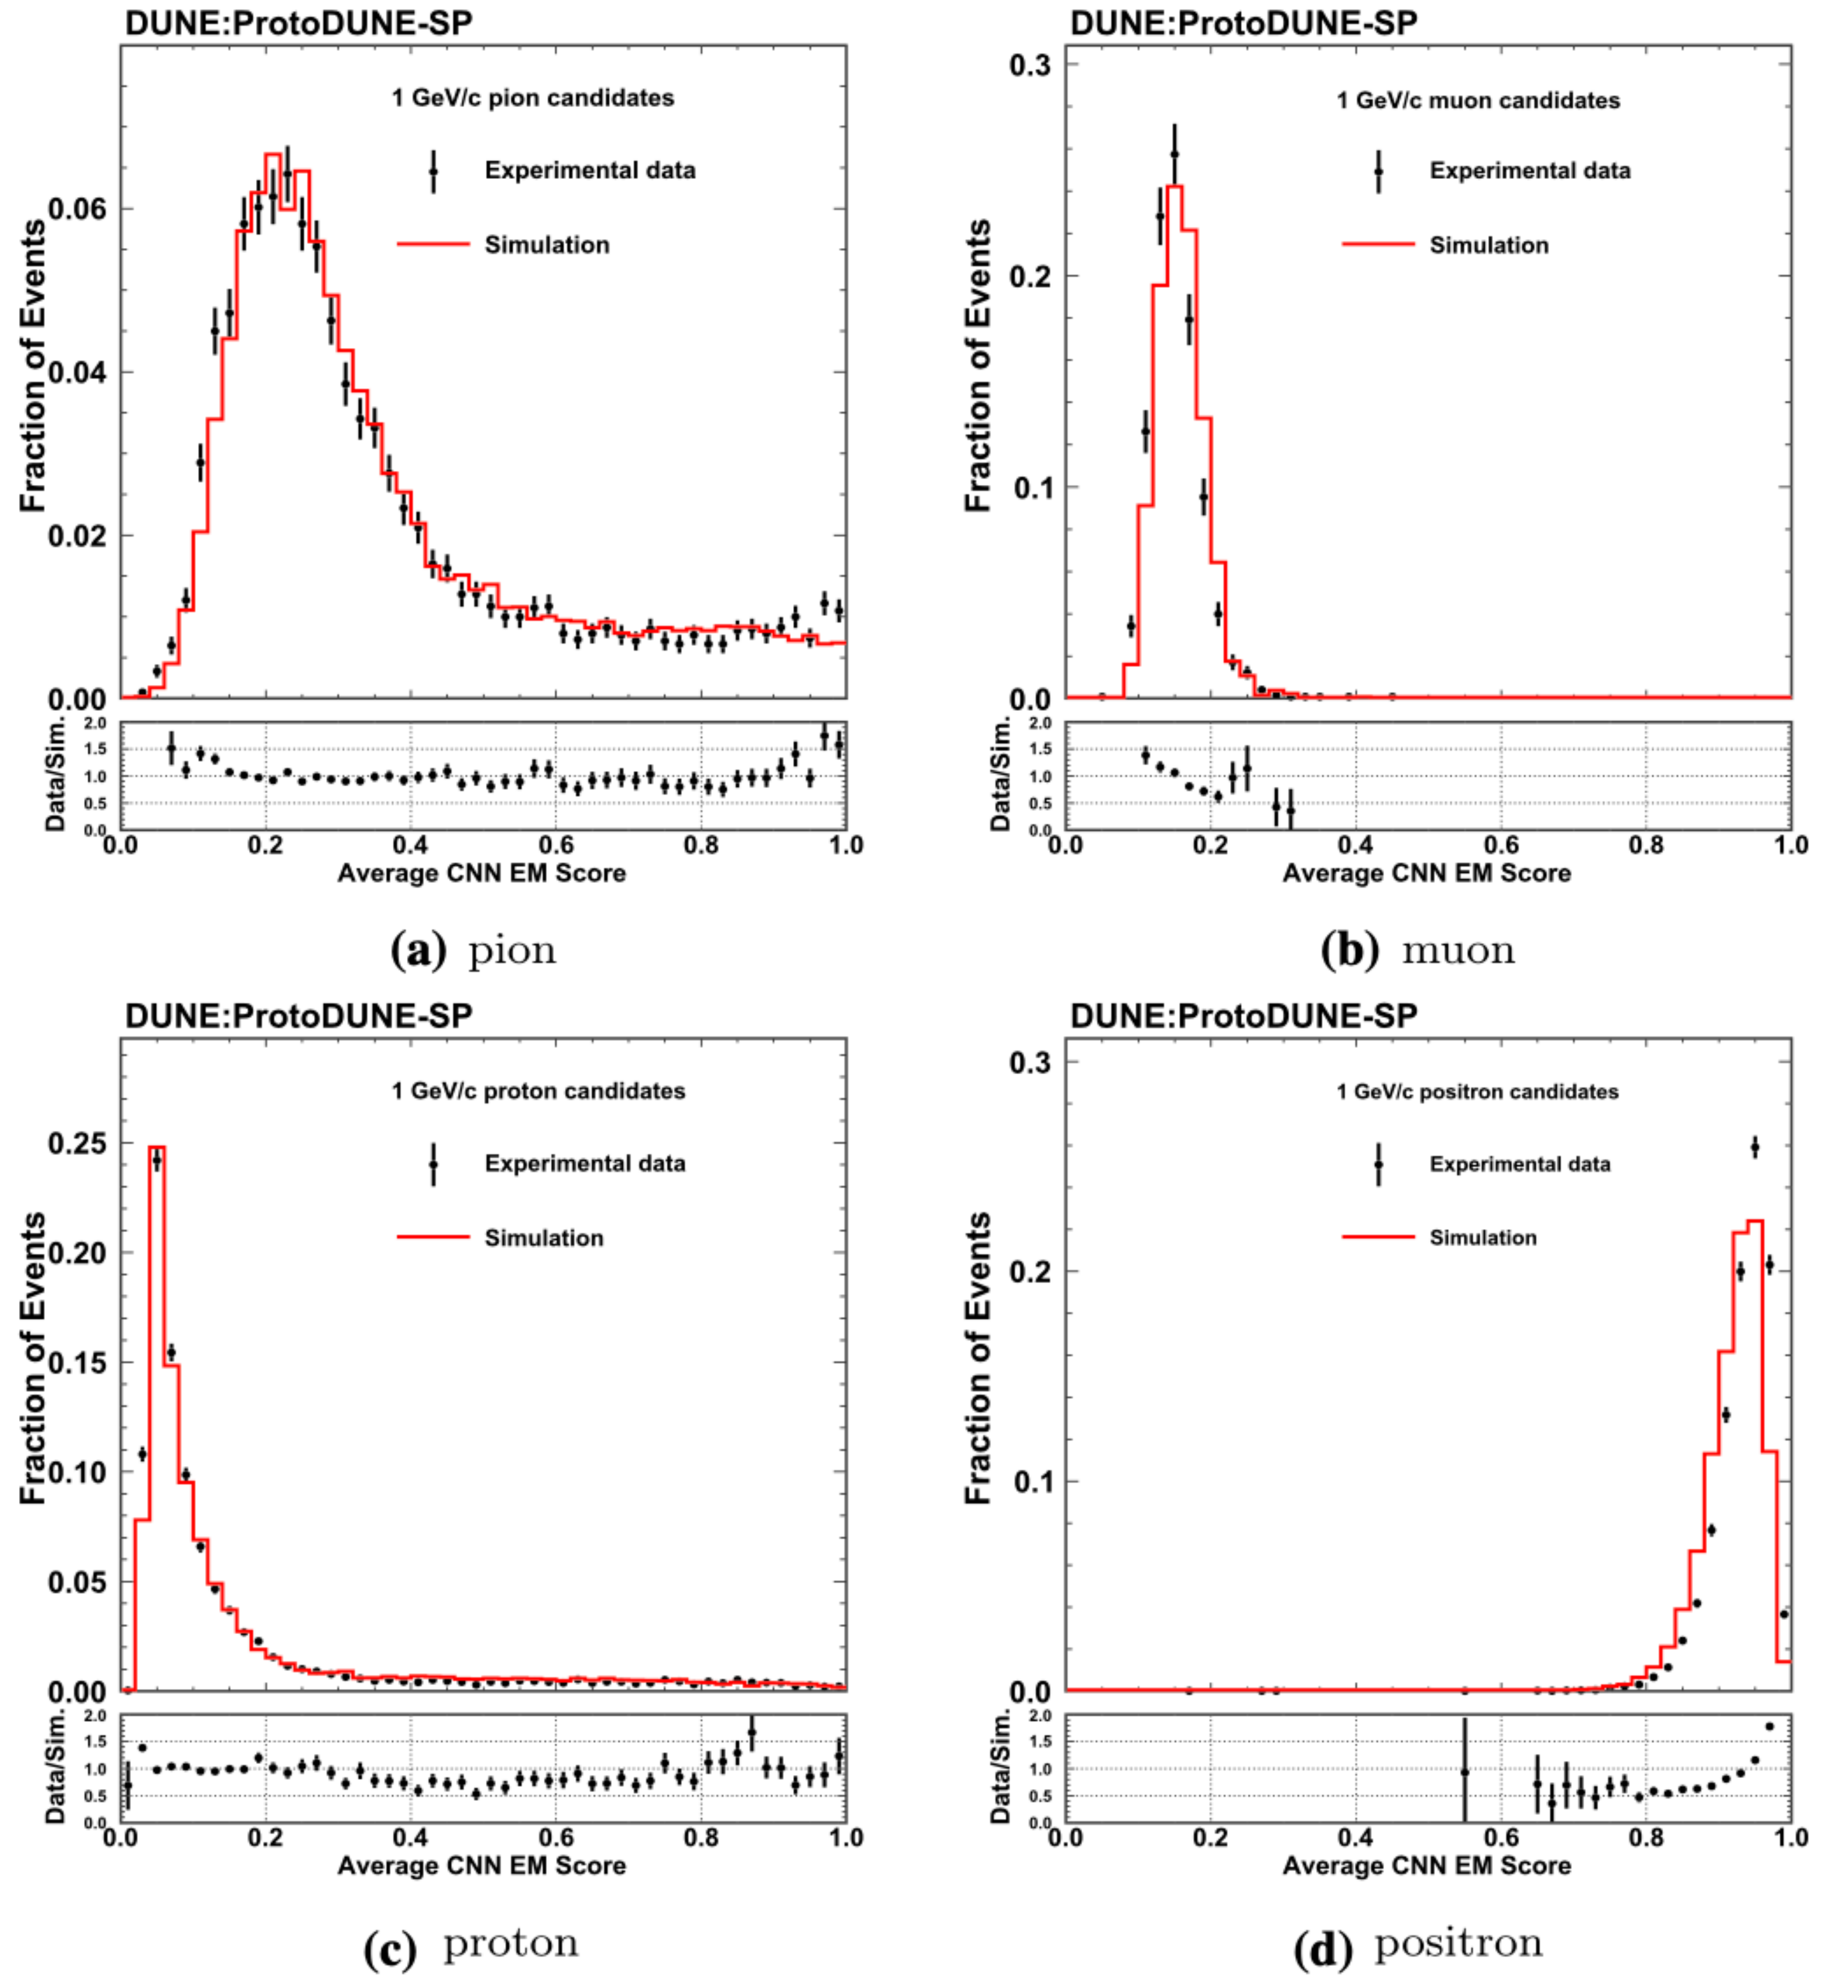

Figures reproduced from A. Abed Abud, *et al.*, Separation of track- and shower-like energy deposits in ProtoDUNE-SP using a convolutional neural network, *Eur.Phys.J.C* **82** 10, 903, (2022)

Black Boxes and Potential Biases

Occlusion tests

- One way to gain an understanding of what a CNN is looking for is to perform occlusion tests
  - This involves removing a small patch of pixels to see how the classification changes
  - Time consuming as we do it for each pixel in the image and rerun the inference
  - Output is a map of the change in classification score for each pixel
- Look at DUNE CVN example[1]
- This work was inspired by a study from NOvA[2]

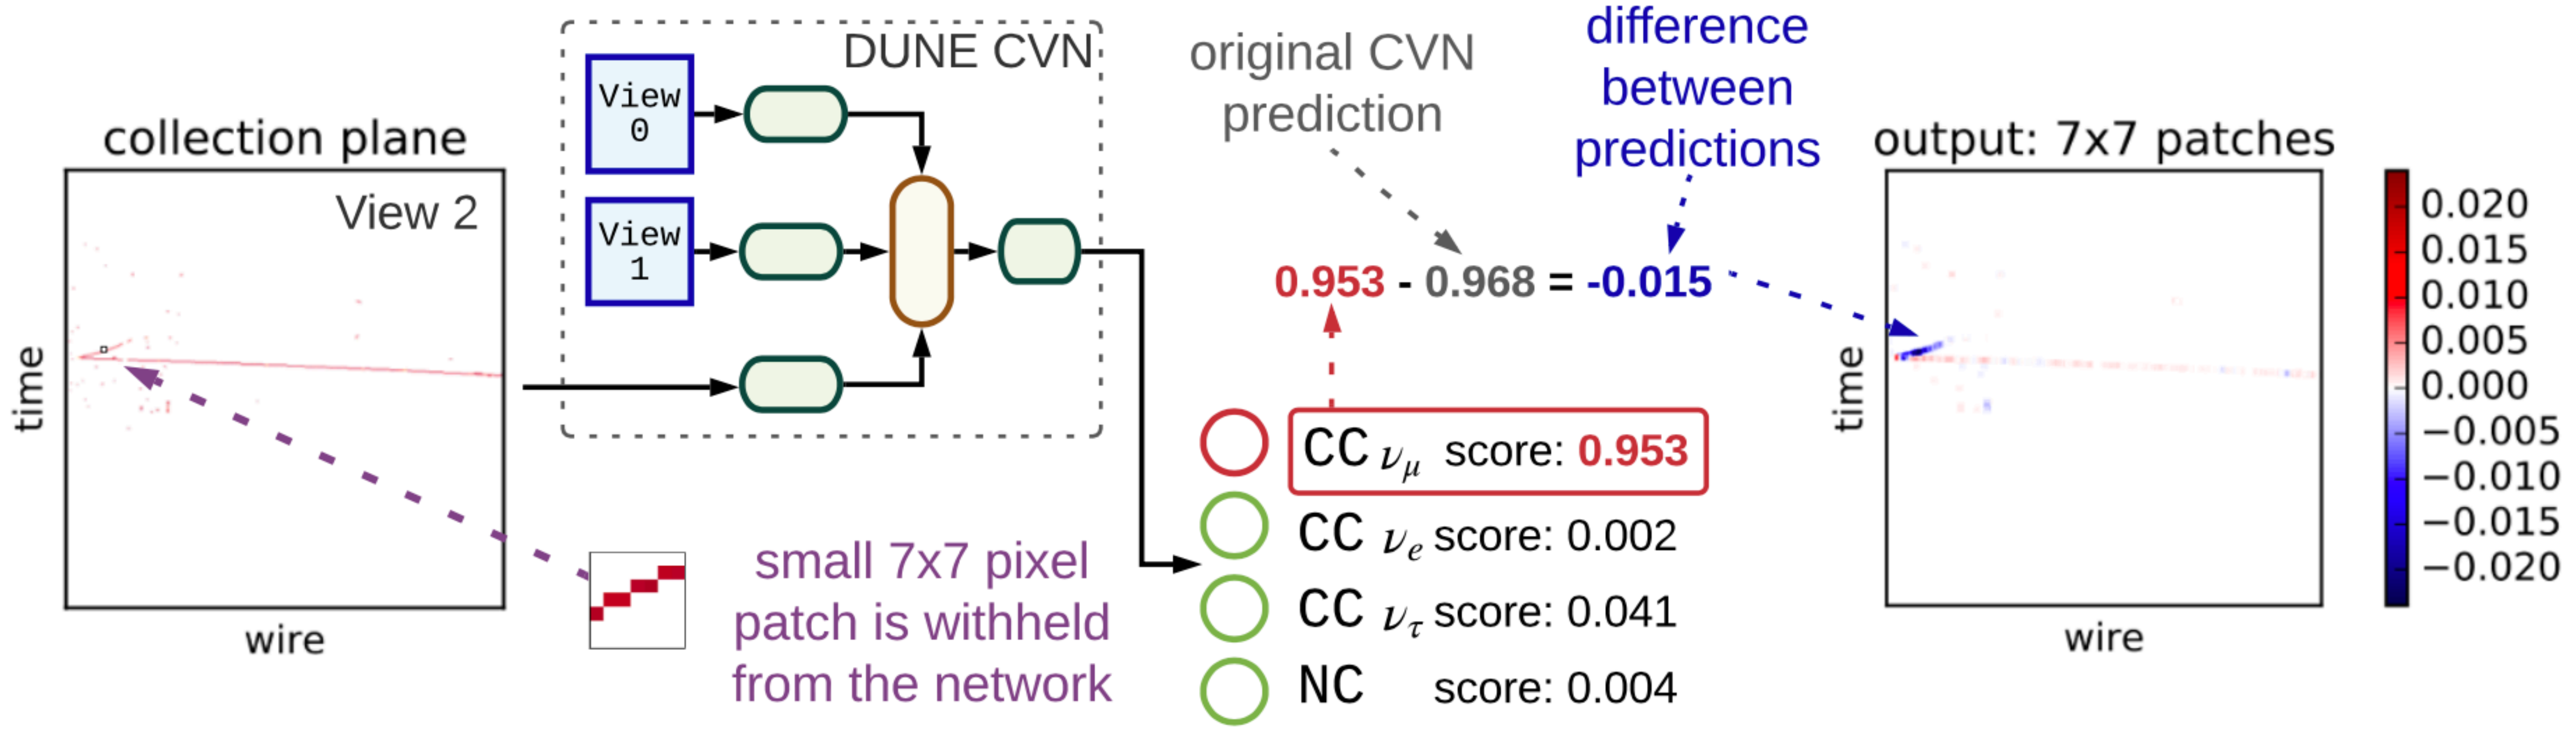

Figure reproduced from [1]

[1] S. Alonso Monsalve, *Novel usage of deep learning and high-performance computing in long-baseline neutrino oscillation experiments*, PhD Thesis, Universidad Carlos III Madrid (2021)
[2] B. L. Howard Jr, *Toward a Precision Era of Neutrino Oscillation Physics: Liquid Argon Scintillation Detector Development for DUNE and Neutrino Oscillation Studies with NOvA*, PhD Thesis, Indiana University (2019)

DUNE CVN - Occlusion tests

- Change in CC ν_e score for a true CC ν_e event occluding (5 x 5) pixel patches
  - We expect that removing pixels around the vertex will lower the score

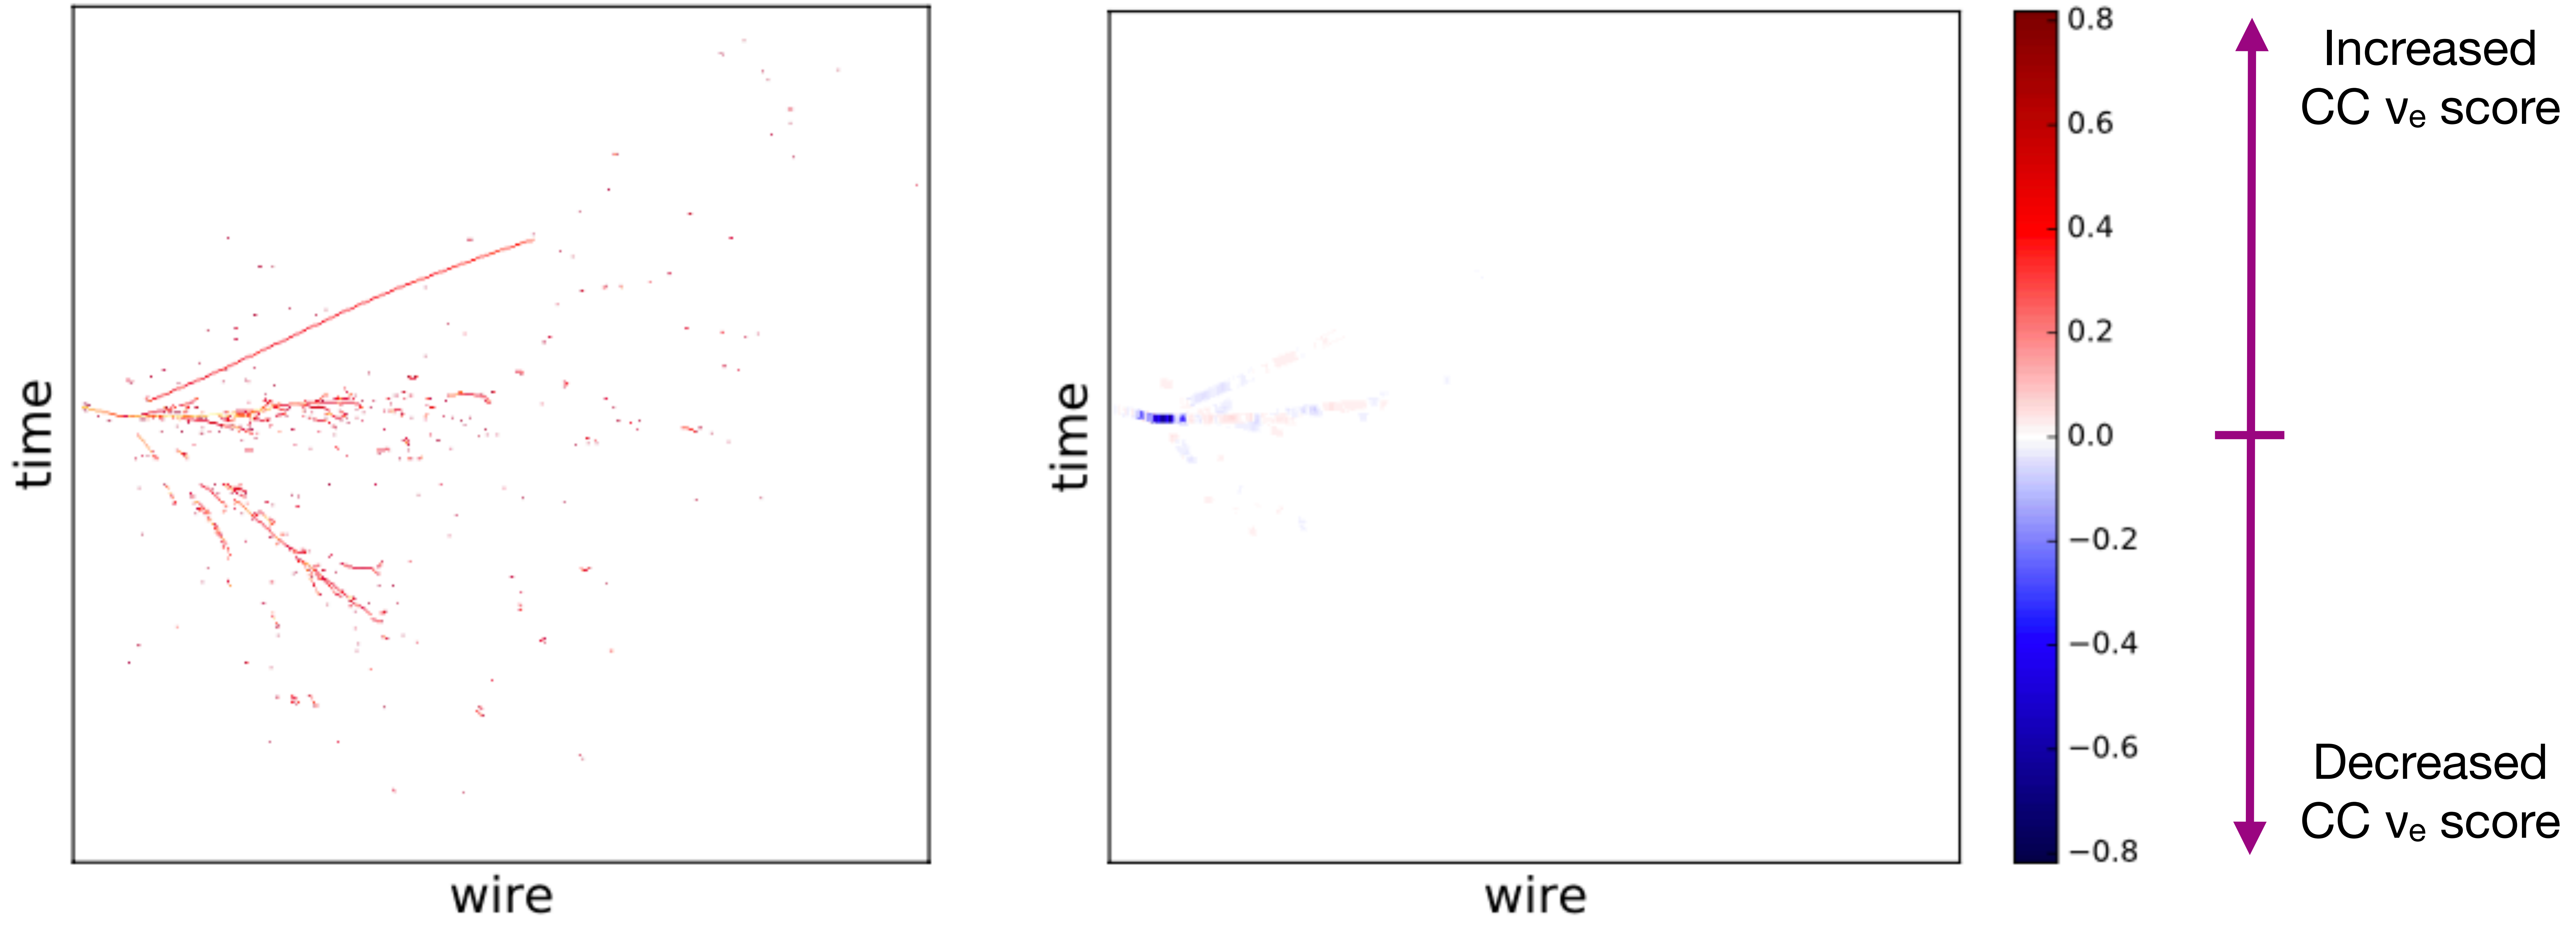

Figure reproduced from [1]

[1] S. Alonso Monsalve, *Novel usage of deep learning and high-performance computing in long-baseline neutrino oscillation experiments*, PhD Thesis, Universidad Carlos III Madrid (2021)

DUNE CVN - Occlusion tests

- Change in CC ν_e score for a true CC ν_e event occluding (5 x 5) pixel patches
  - We expect that removing pixels around the vertex will lower the score

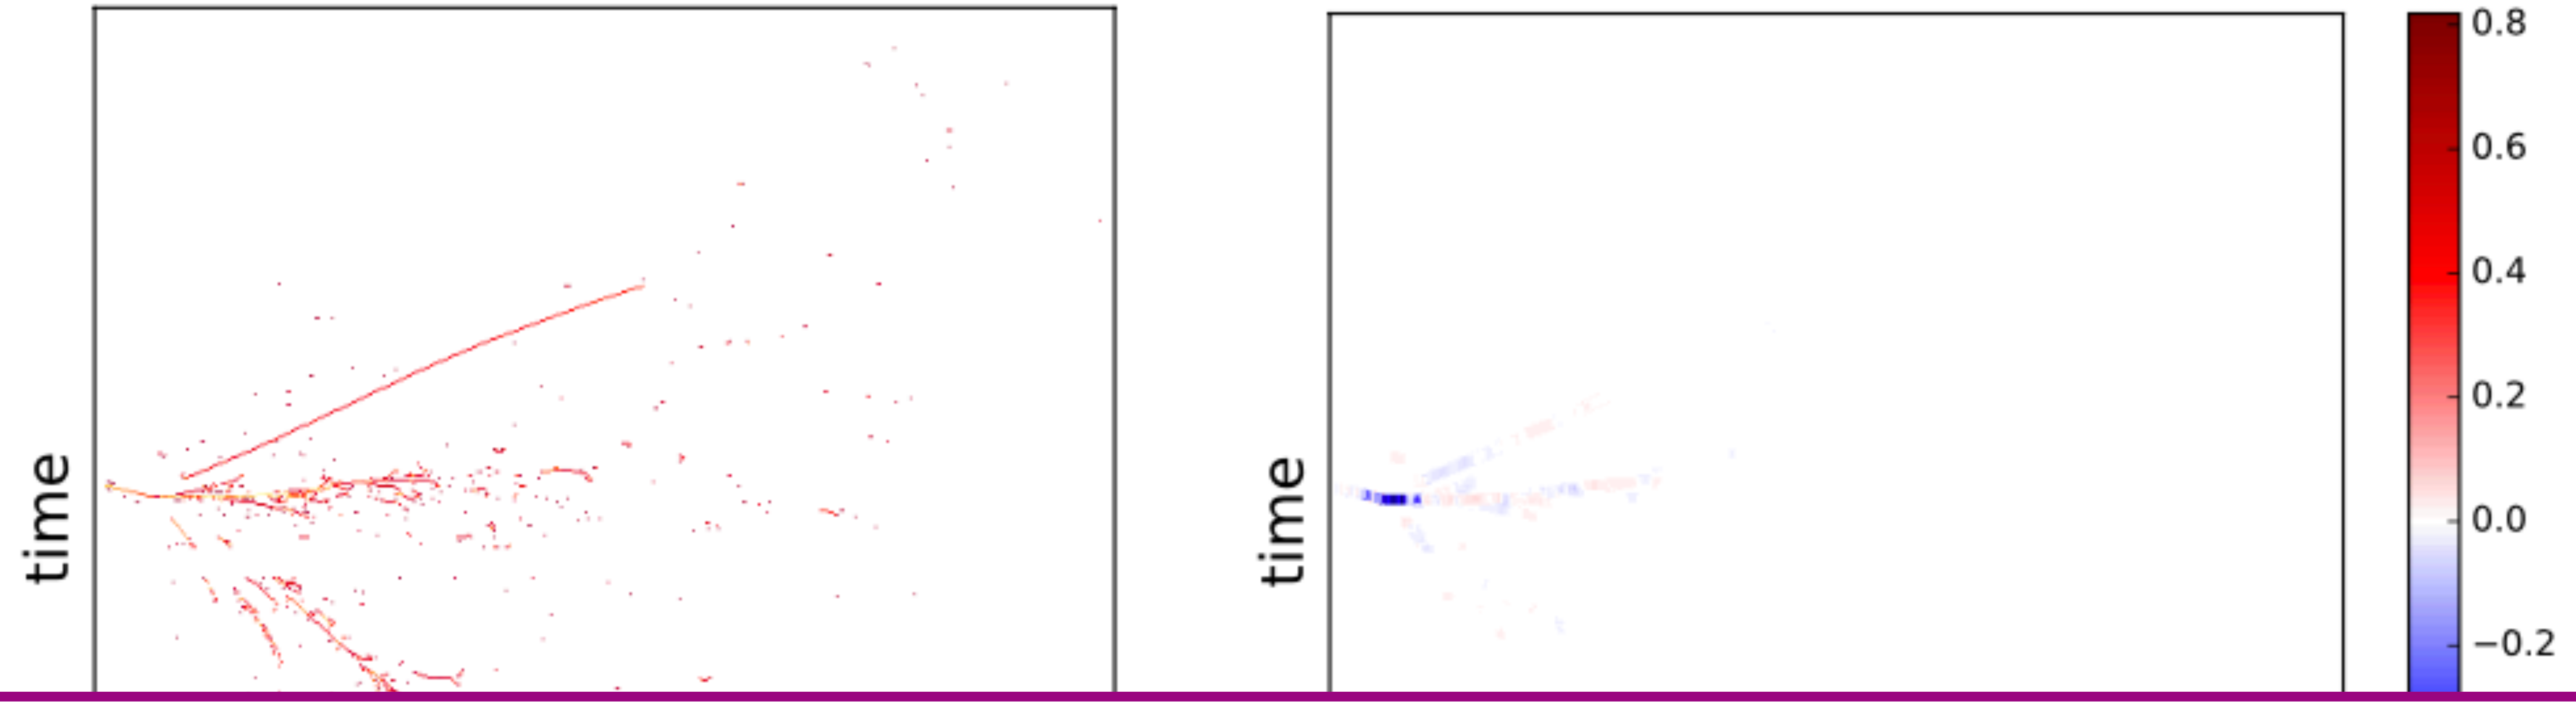

Increased CC ν_e score

Decreased CC ν_e score

	CC ν_μ	CC ν_e	CC ν_τ	NC
Default	0.0009	0.9184	0.0090	0.0717
Max Change	0.0015	0.1003	0.0098	0.8884

Figure reproduced from [1]

[1] S. Alonso Monsalve, *Novel usage of deep learning and high-performance computing in long-baseline neutrino oscillation experiments*, PhD Thesis, Universidad Carlos III Madrid (2021)

DUNE CVN - Occlusion tests

- Change in CC ν_μ score for a true CC ν_μ event occluding (5 x 5) pixel patches
  - This is a bit of a tricky event to classify

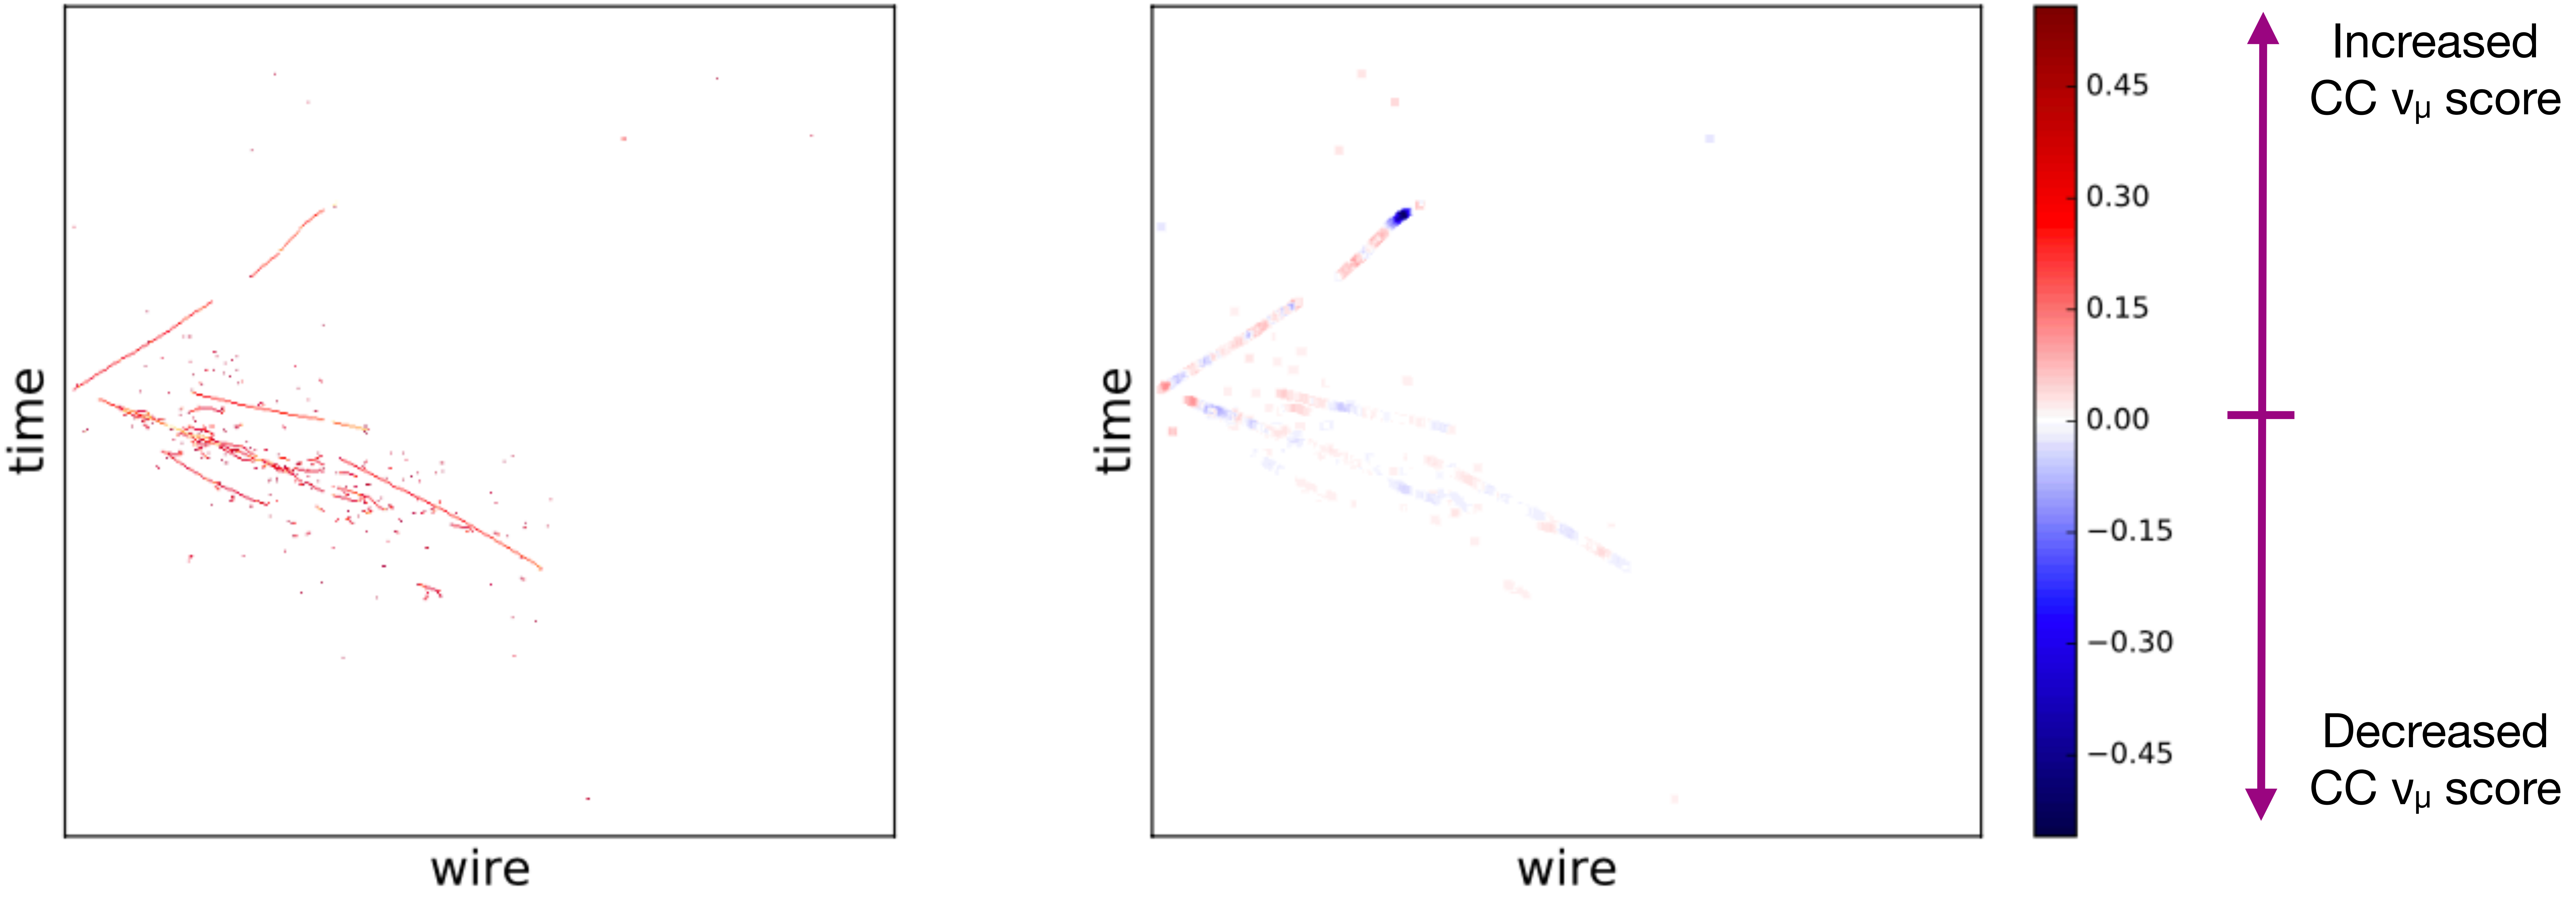

Figure reproduced from [1]

[1] S. Alonso Monsalve, *Novel usage of deep learning and high-performance computing in long-baseline neutrino oscillation experiments*, PhD Thesis, Universidad Carlos III Madrid (2021)

DUNE CVN - Occlusion tests

- Change in CC ν_μ score for a true CC ν_μ event occluding (5 x 5) pixel patches
  - This is a bit of a tricky event to classify

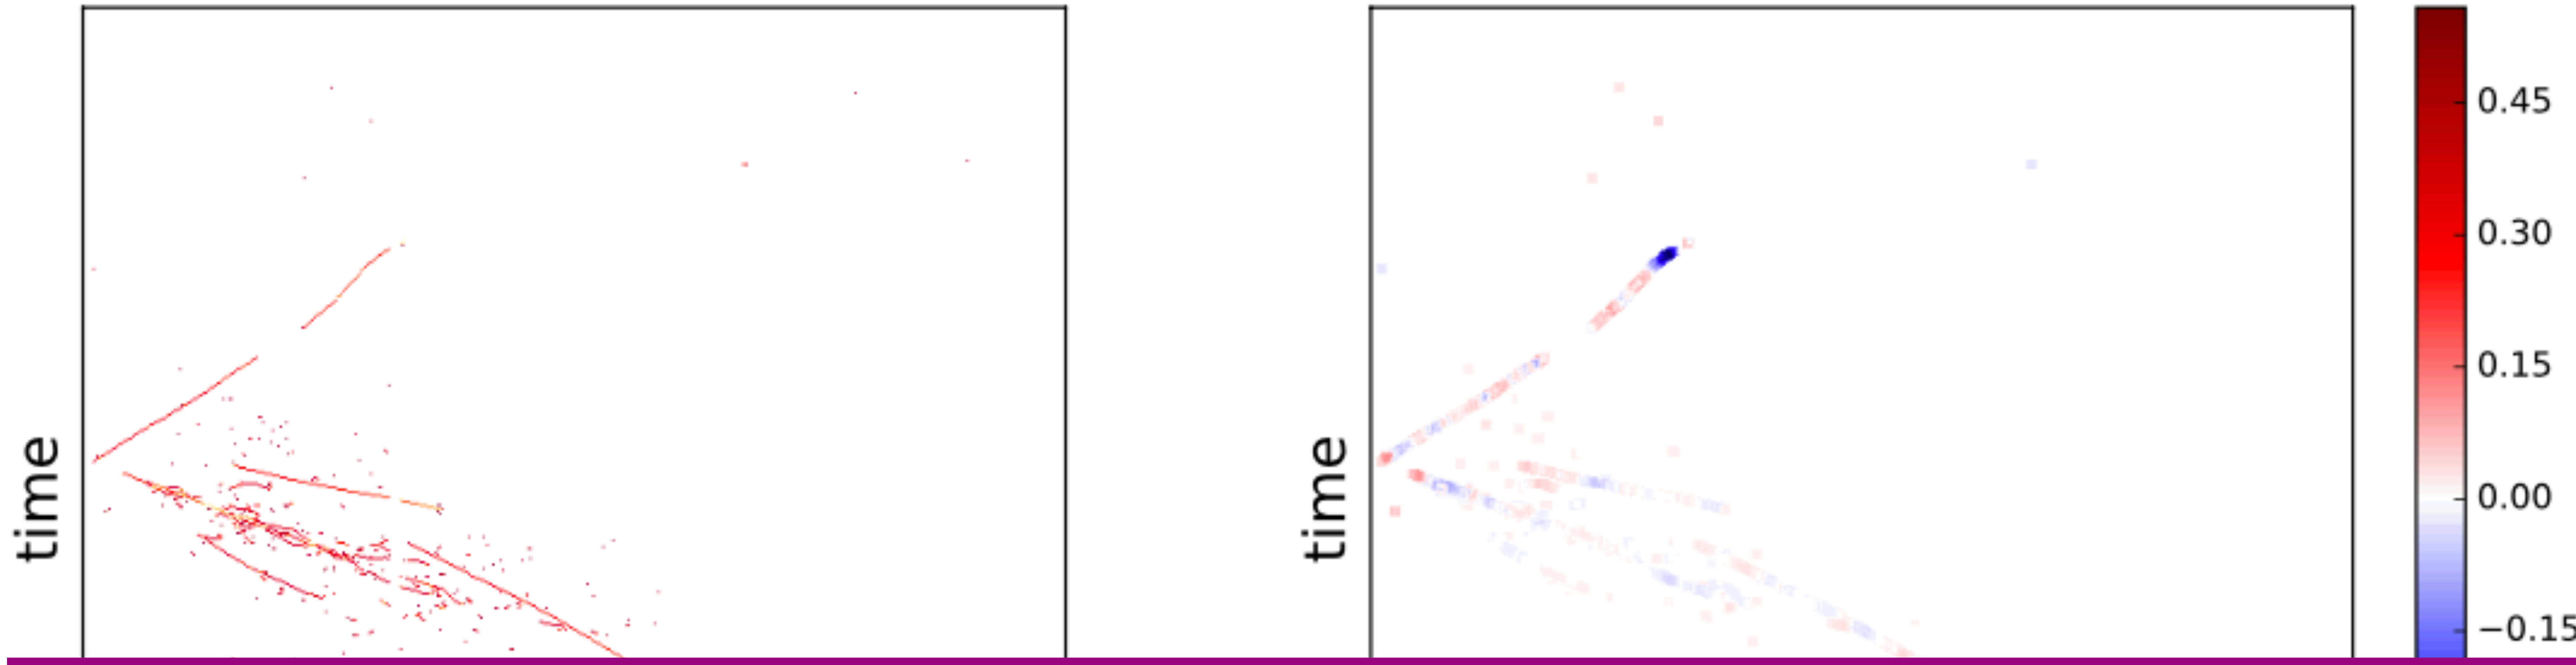

	CC ν_μ	CC ν_e	CC ν_τ	NC
Default	0.07142	0.0007	0.0750	0.2101
Max Change	0.1551	0.0011	0.1152	0.6886

Increased CC ν_μ score

Decreased CC ν_μ score

Figure reproduced from [1]

[1] S. Alonso Monsalve, *Novel usage of deep learning and high-performance computing in long-baseline neutrino oscillation experiments*, PhD Thesis, Universidad Carlos III Madrid (2021)

Potential biases

- One common concern on the use of deep learning is the choice of which MC training sample to use
  - Simulations are never perfect
  - Will it have biases when tested on real data given it is trained on MC?
- The MINERvA experiment performed a nice study to investigate this[1]
- They have a CNN used for vertex finding
  - It is trained with a domain adversarial neural network (DANN)[2]
  - The DANN tries to distinguish between different simulation samples

[1] G.N. Perdue, *et al.*, Reducing model bias in a deep learning classifier using domain adversarial neural networks in the MINERvA experiment, JINST 13 P11020 (2018)
[2] Y. Ganin, *et al.*, Domain-Adversarial Training of Neural Networks, Journal of Machine Learning Research 17 59 (2016)

MINERvA vertex finding

- I don't have time to go into the details of DANNs, but they look a bit like this:
- In MINERvA's use case:
  - Green: feature extractor CNN
  - Blue: vertex location prediction
  - Pink: "which MC sample is this?"
- The role of the domain classifier is to allow the CNN to only learn generic features common to the different domains

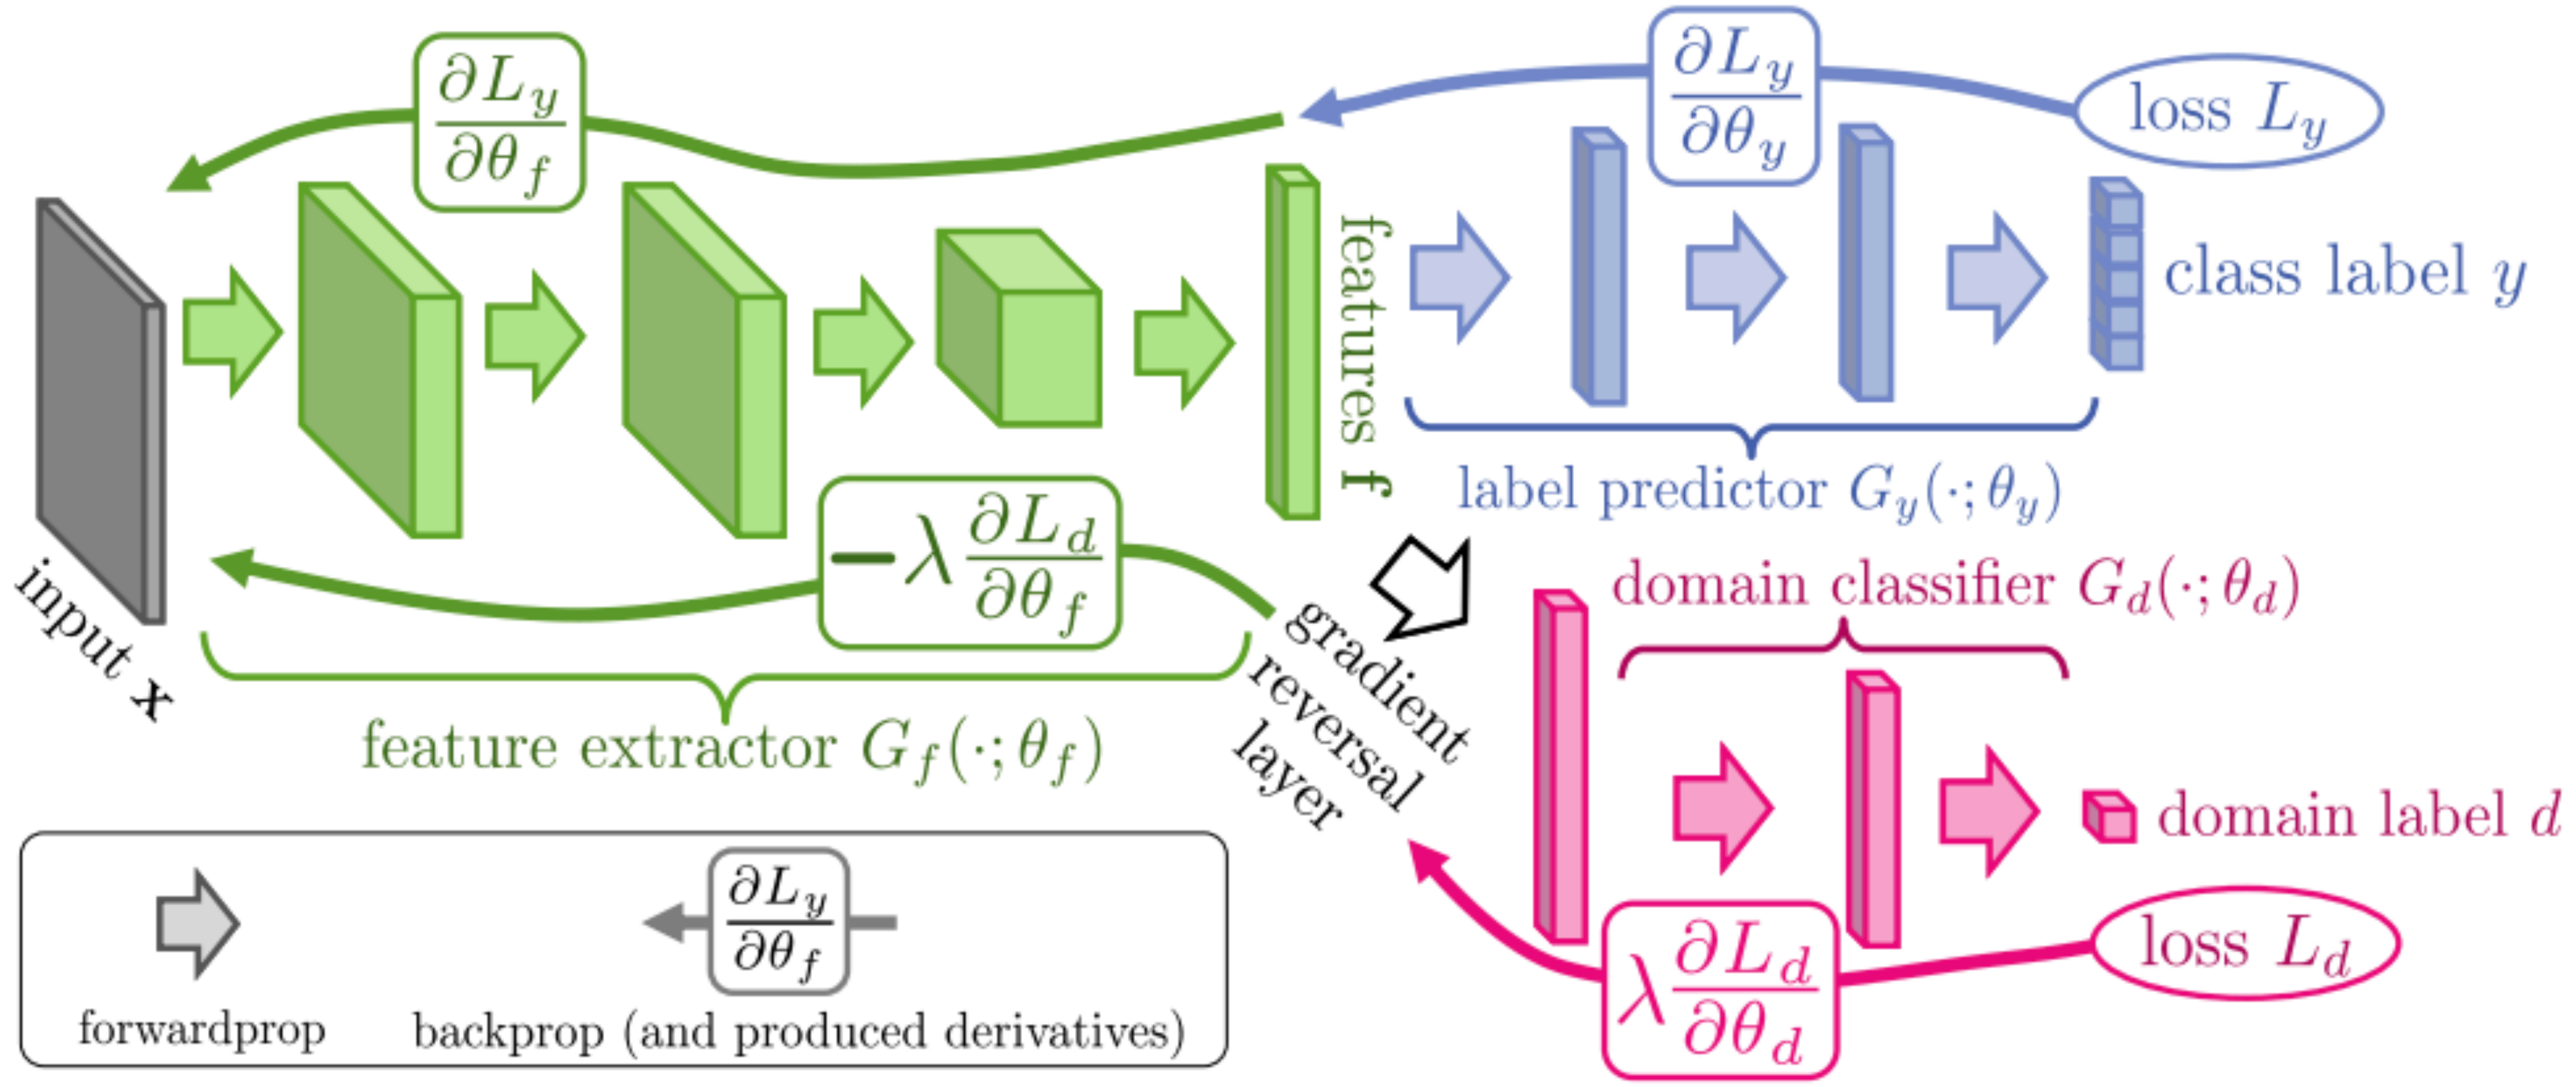

Figure from [2]

[1] G.N. Perdue, *et al.*, Reducing model bias in a deep learning classifier using domain adversarial neural networks in the MINERvA experiment, JINST 13 P11020 (2018)
[2] Y. Ganin, *et al.*, Domain-Adversarial Training of Neural Networks, Journal of Machine Learning Research 17 59 (2016)

MINERvA vertex finding

- The accuracy for different trainings is given below
- Blue: train and test CNN on the same sample
- Black: train CNN on one sample but test on another
- Green: train on one sample but test on another with the DANN
- Red: as above but with more statistics

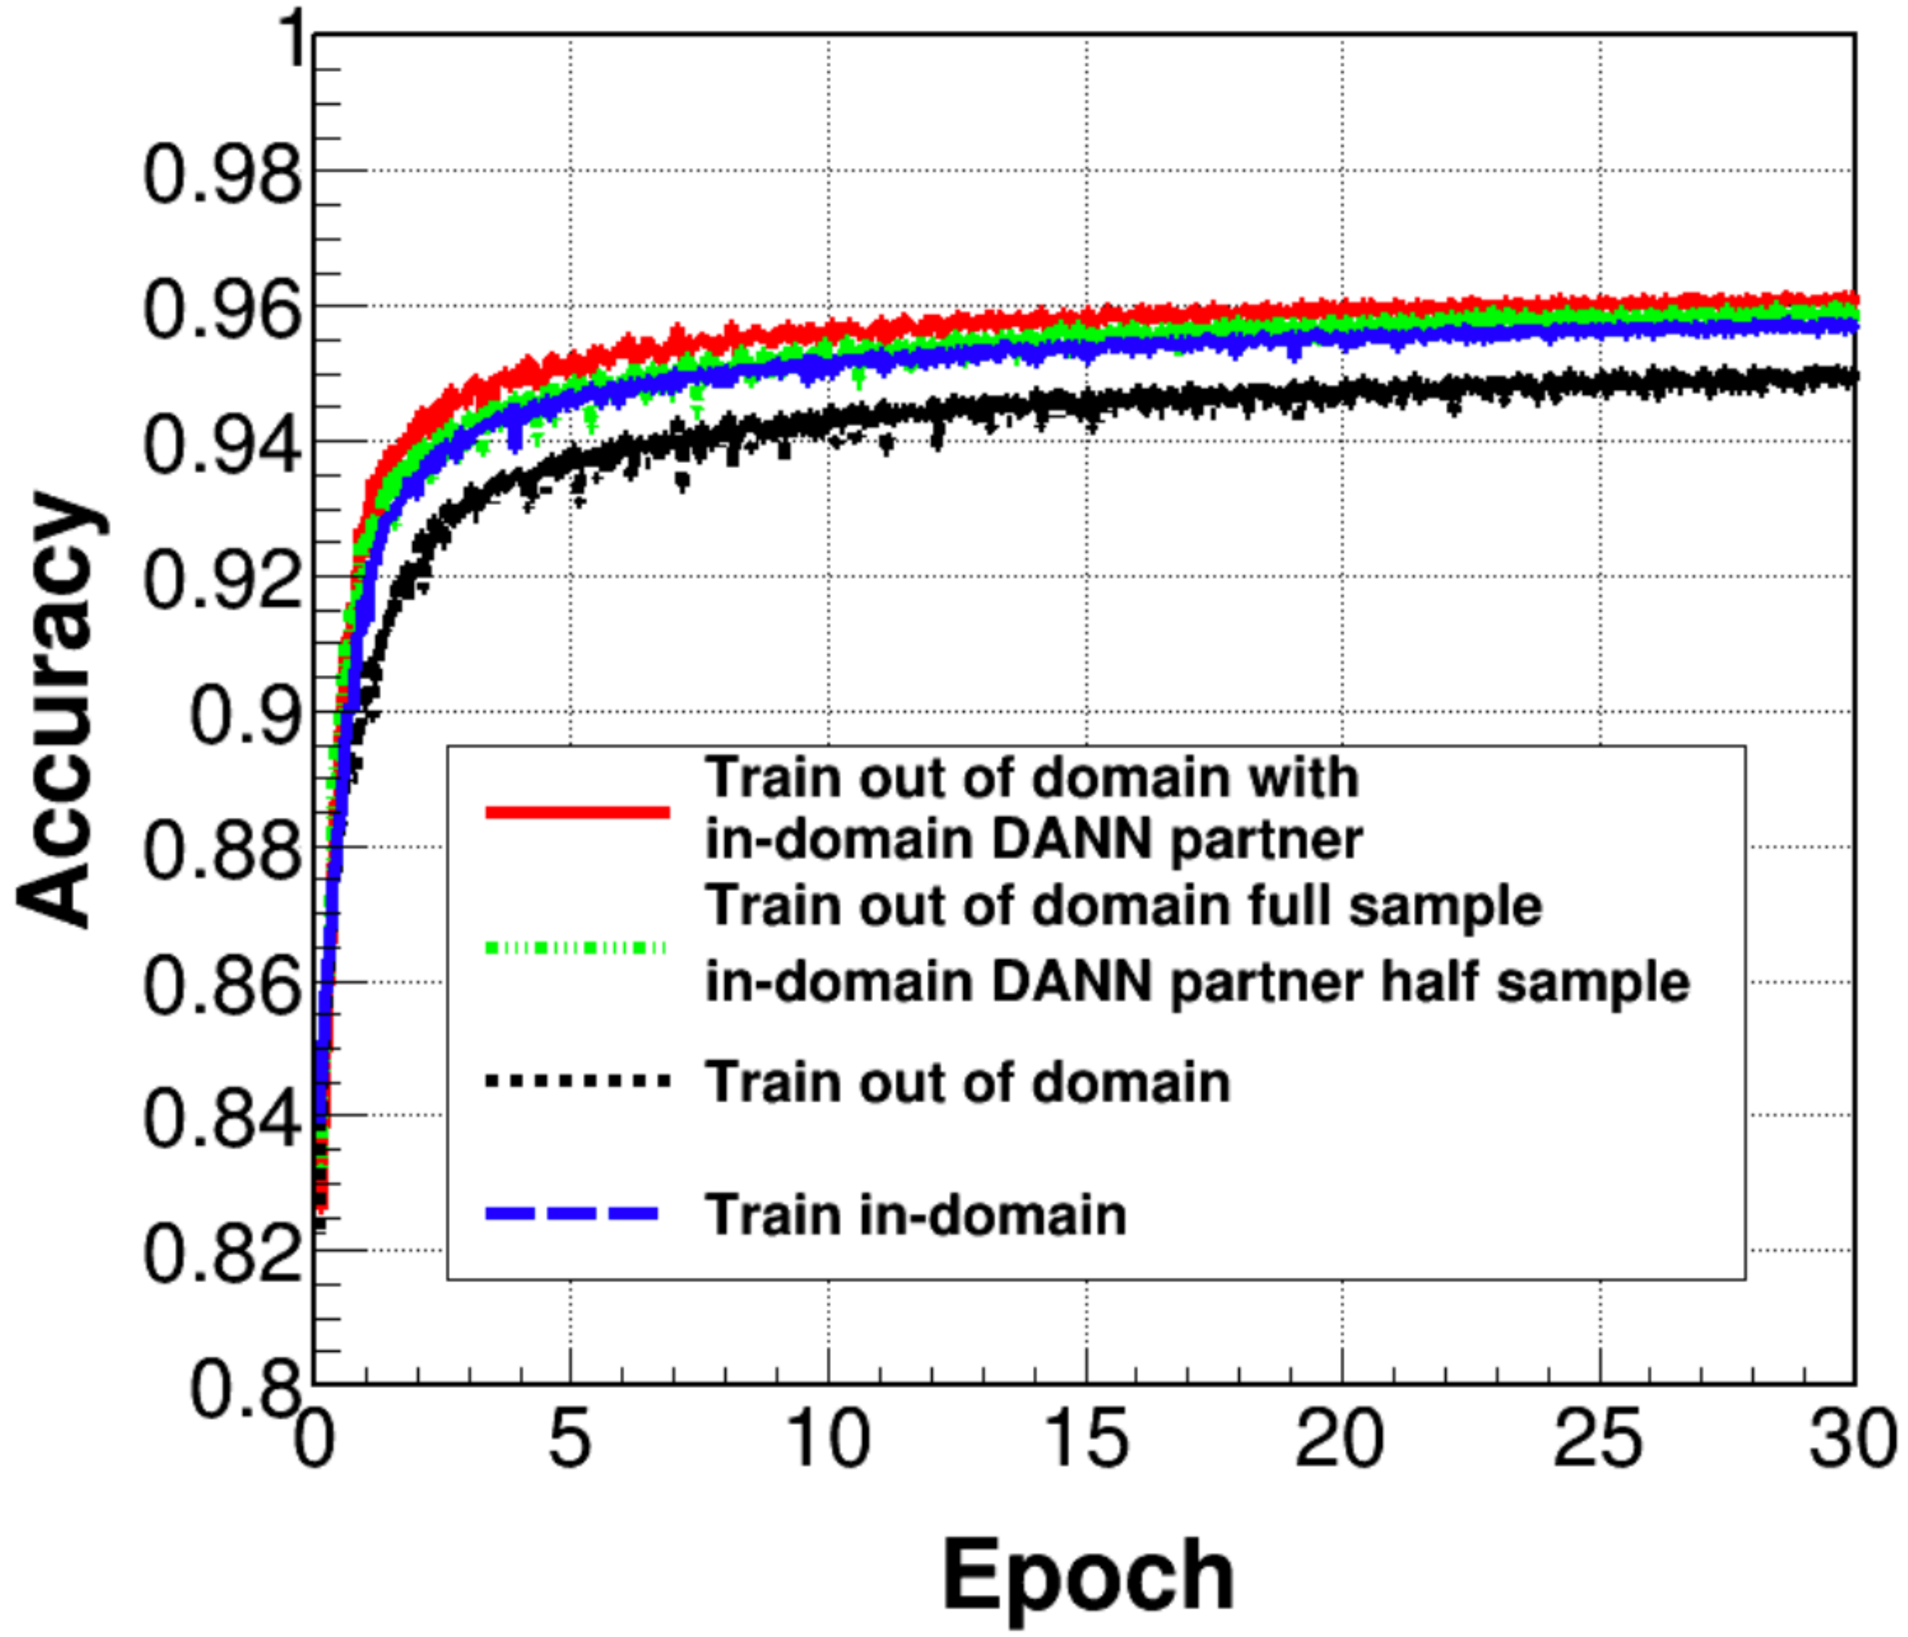

Figure from [1]

[1] G.N. Perdue, *et al.*, Reducing model bias in a deep learning classifier using domain adversarial neural networks in the MINERvA experiment, JINST 13 P11020 (2018)

Transfer Learning

What is transfer learning?

- Transfer learning makes use of previously trained networks
  - Allows you to fine tune a pre-trained network for your task
  - Can be useful if you don't have much data
  - The idea dates back to the early days of perceptrons[1]

- I will discuss a recent study we performed on using transfer learning in neutrino event classification

Eur. Phys. J. C (2022) 82:1099
https://doi.org/10.1140/epjc/s10052-022-11066-6

THE EUROPEAN PHYSICAL JOURNAL C

Check for updates

Regular Article - Experimental Physics

Application of transfer learning to neutrino interaction classification

Andrew Chappell[2,a], **Leigh H. Whitehead**[1,b]

[1] Department of Physics, University of Cambridge, Cambridge CB3 0HE, UK
[2] Department of Physics, University of Warwick, Coventry CV4 7AL, UK

https://link.springer.com/article/10.1140/epjc/s10052-022-11066-6

[1] S. Bozinovski, A. Fulgosi, *The influence of pattern similarity and transfer learning upon the training of a base perceptron* b2. In: Proceedings of Symposium Informatica, Bled, Slovenia (1976) p. 3–1215.

Transfer Learning in Physics

- I was only able to find once example of transfer learning in a related field when we started this work
- The AT-TPC[1] was a nuclear physics experiment
- Used transfer learning due to a small simulation dataset
- Also used some hand-labelled data due to poor simulation quality

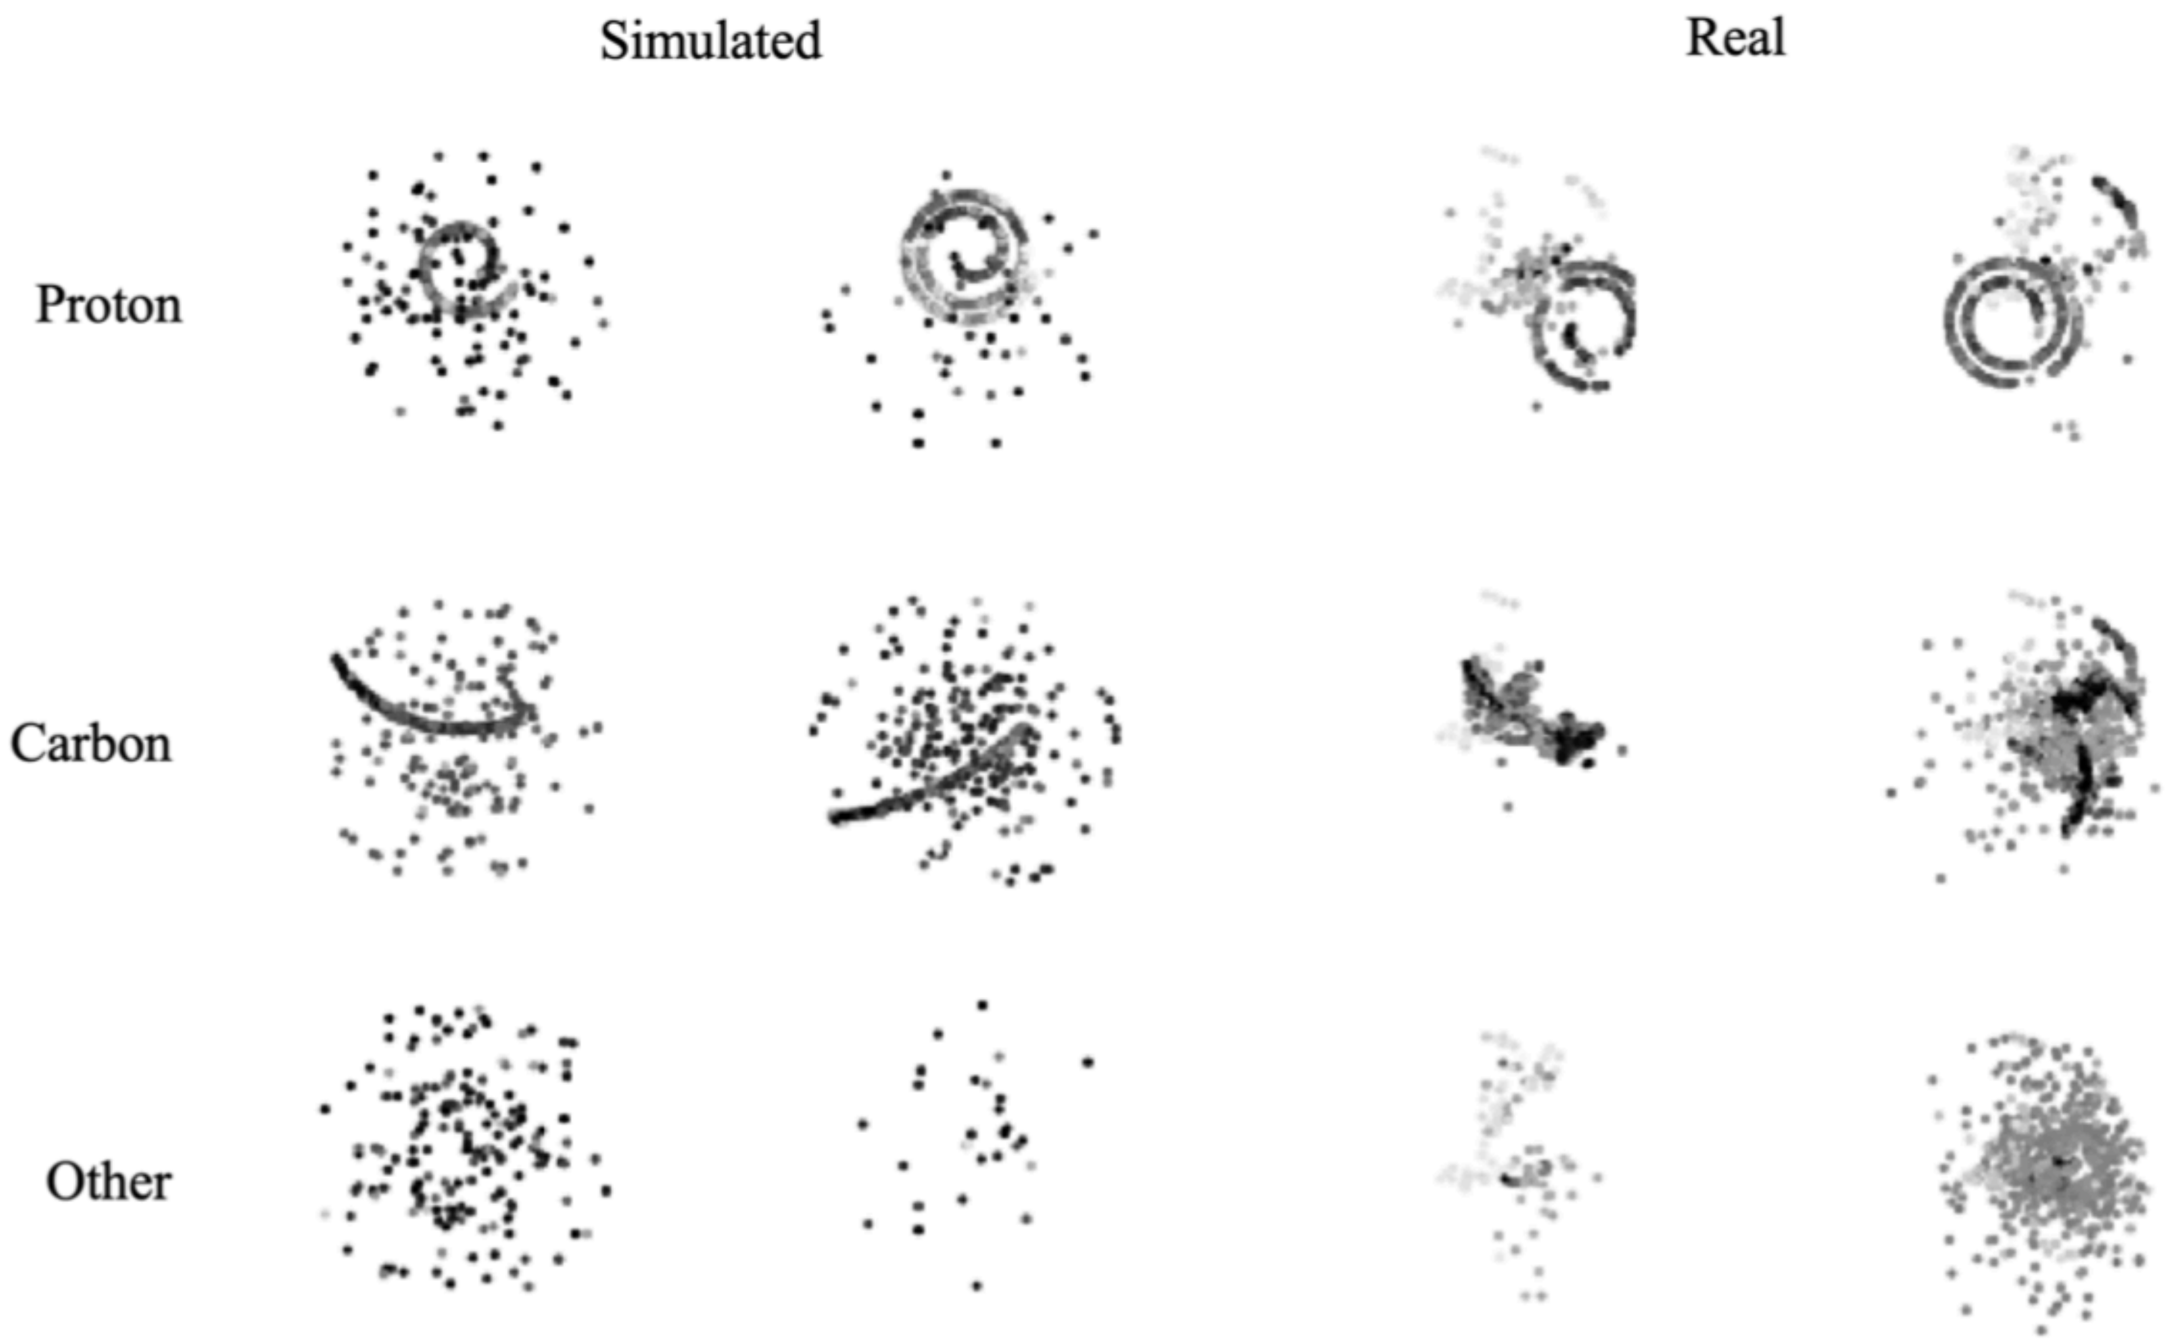

[1] M. P. Kuchera, et al., Machine Learning Methods for Track Classification in the AT-TPC, NIM A 940 (2019) 156-167, 1810.10350

Transfer Learning in LArTPCs

- Can we use transfer learning to reduce the number of training examples?
  - Simulations are time consuming and GPUs need a lot of power
- Conveniently, LArTPC detectors, such as DUNE, have three readout planes
  - We get three images of a given interaction
  - Photographic images have depth three (red, green and blue channels)
- Can we use a network trained on photographs for our event classification?
  - There are plenty of networks trained on photograph-based challenges
  - Use these networks as a starting point and fine tune the weights

TL: Event Sample

- GENIE neutrino events:
  - CC ν_μ, CC ν_e and NC
  - 50,000 of each type

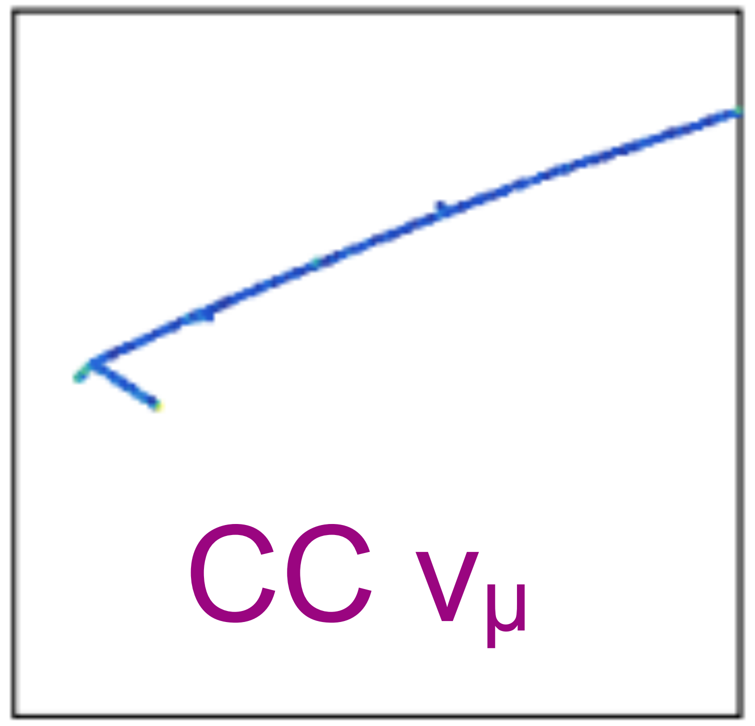

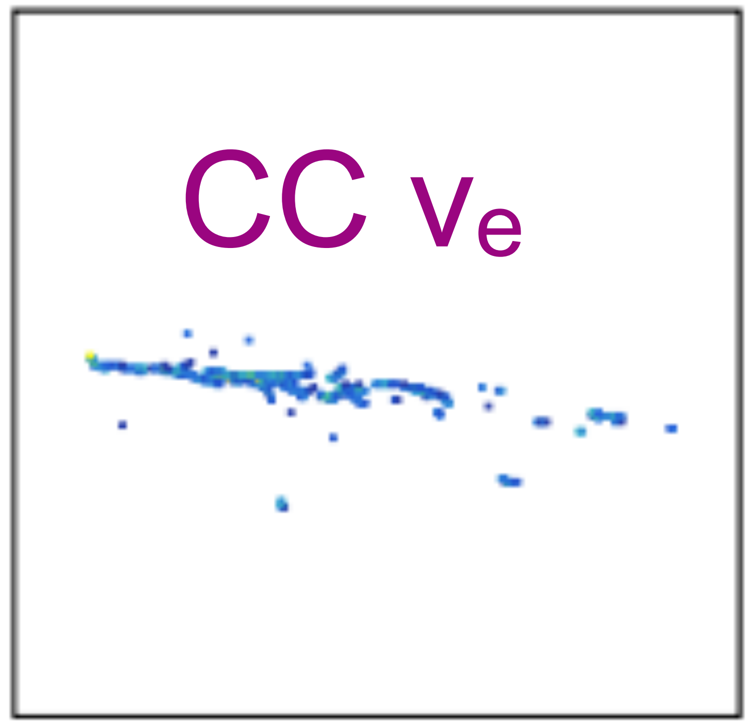

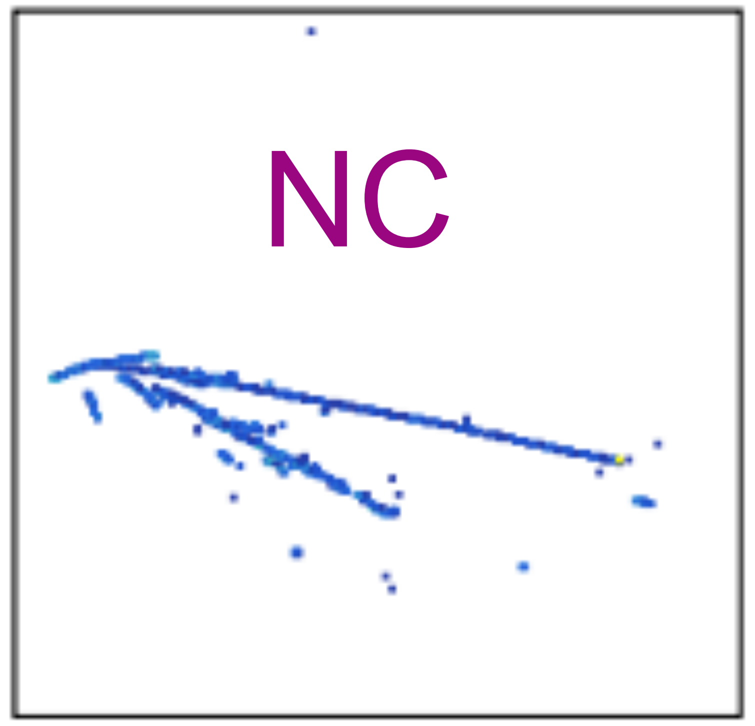

- Events passed through simple LArTPC simulation
  - Outputs three images of each event
  - Three projections of the (y,z) plane

CC ν_μ event with the three views overlaid as RGB channels

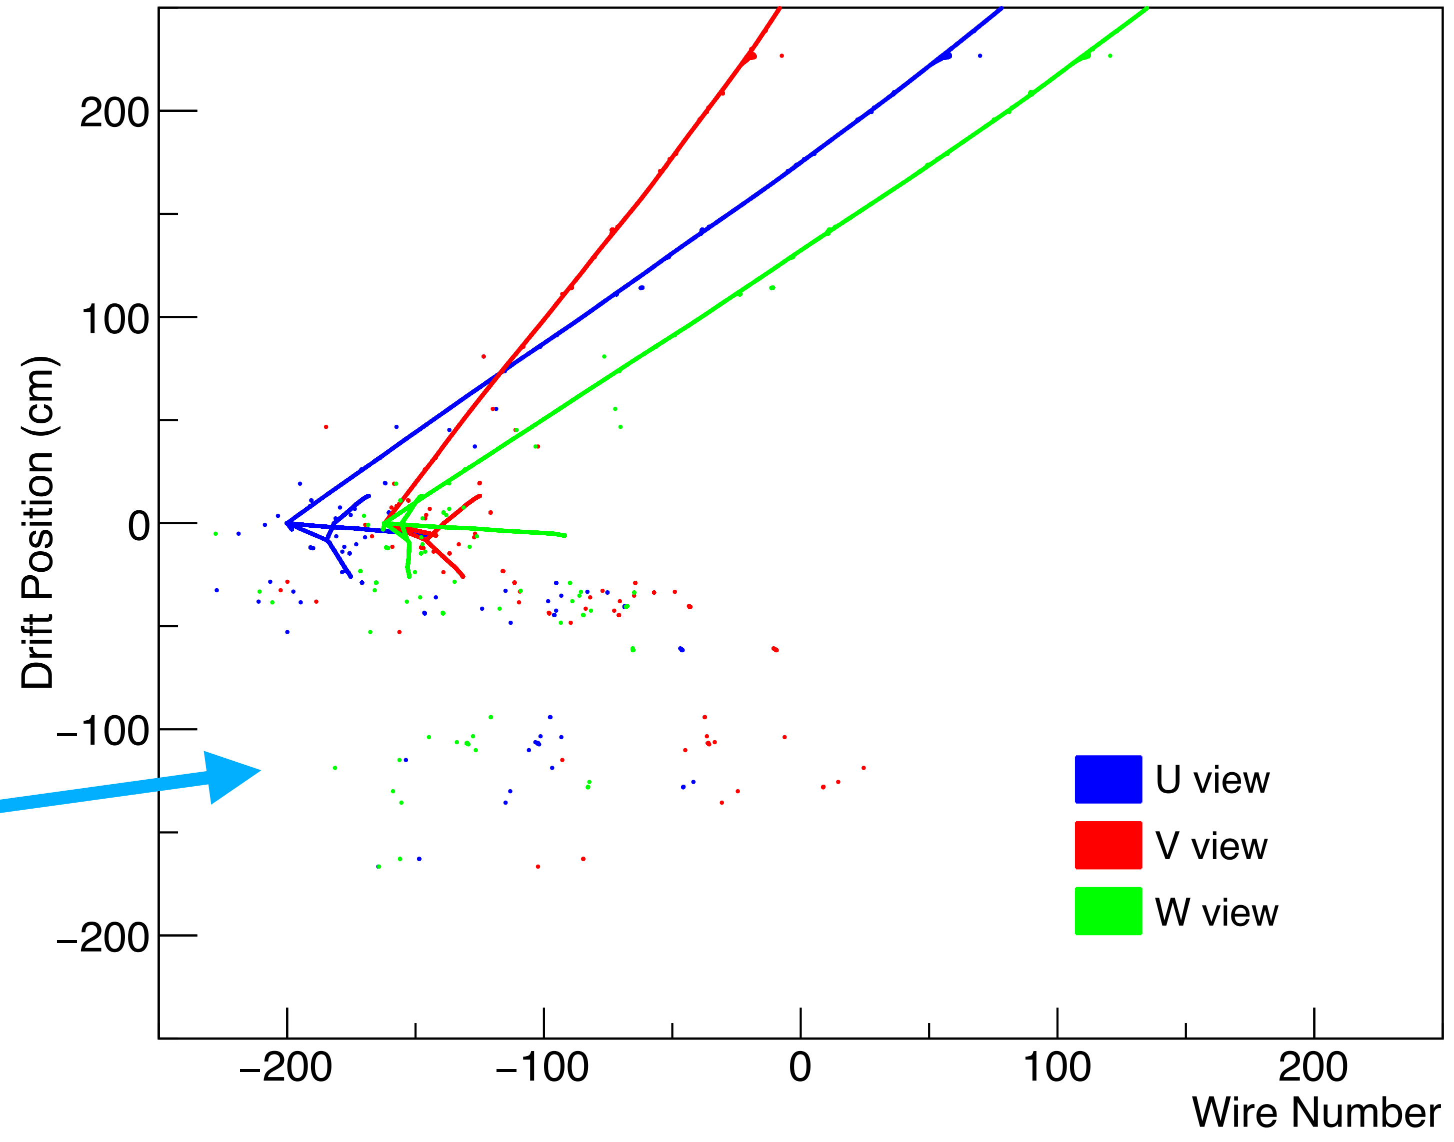

TL: Architecture and Training

- We chose to use the Pytorch implementation of ResNet18
  - Small depth was chosen since this study involved training over 1000 networks
- The pre-trained version of ResNet18 was trained on ImageNet
  - We had to change the final layer from 1000 to 3 classes: CC ν_μ, CC ν_e and NC
- Trained ensembles of 25 networks with:
  - Either:
    - Kaiming (He) randomly initialised weights ← Standard initialisation for ResNets
    - Weights from the pre-trained ImageNet network
  - Training samples ranging from 1,000 to 100,000 events

Results: TF vs random initialisation

- Compared the F1 score from the transfer-learned networks fine-tuned with 1k to 100k images against the Kaiming-initialised network with 50k and 100k events

- Transfer-learned network out-performs the Kaiming-initialised network with 100k training images
  - For 7k training images and above

- Event fine-tuning just the final layer works surprising well
  - F1 score = 0.79

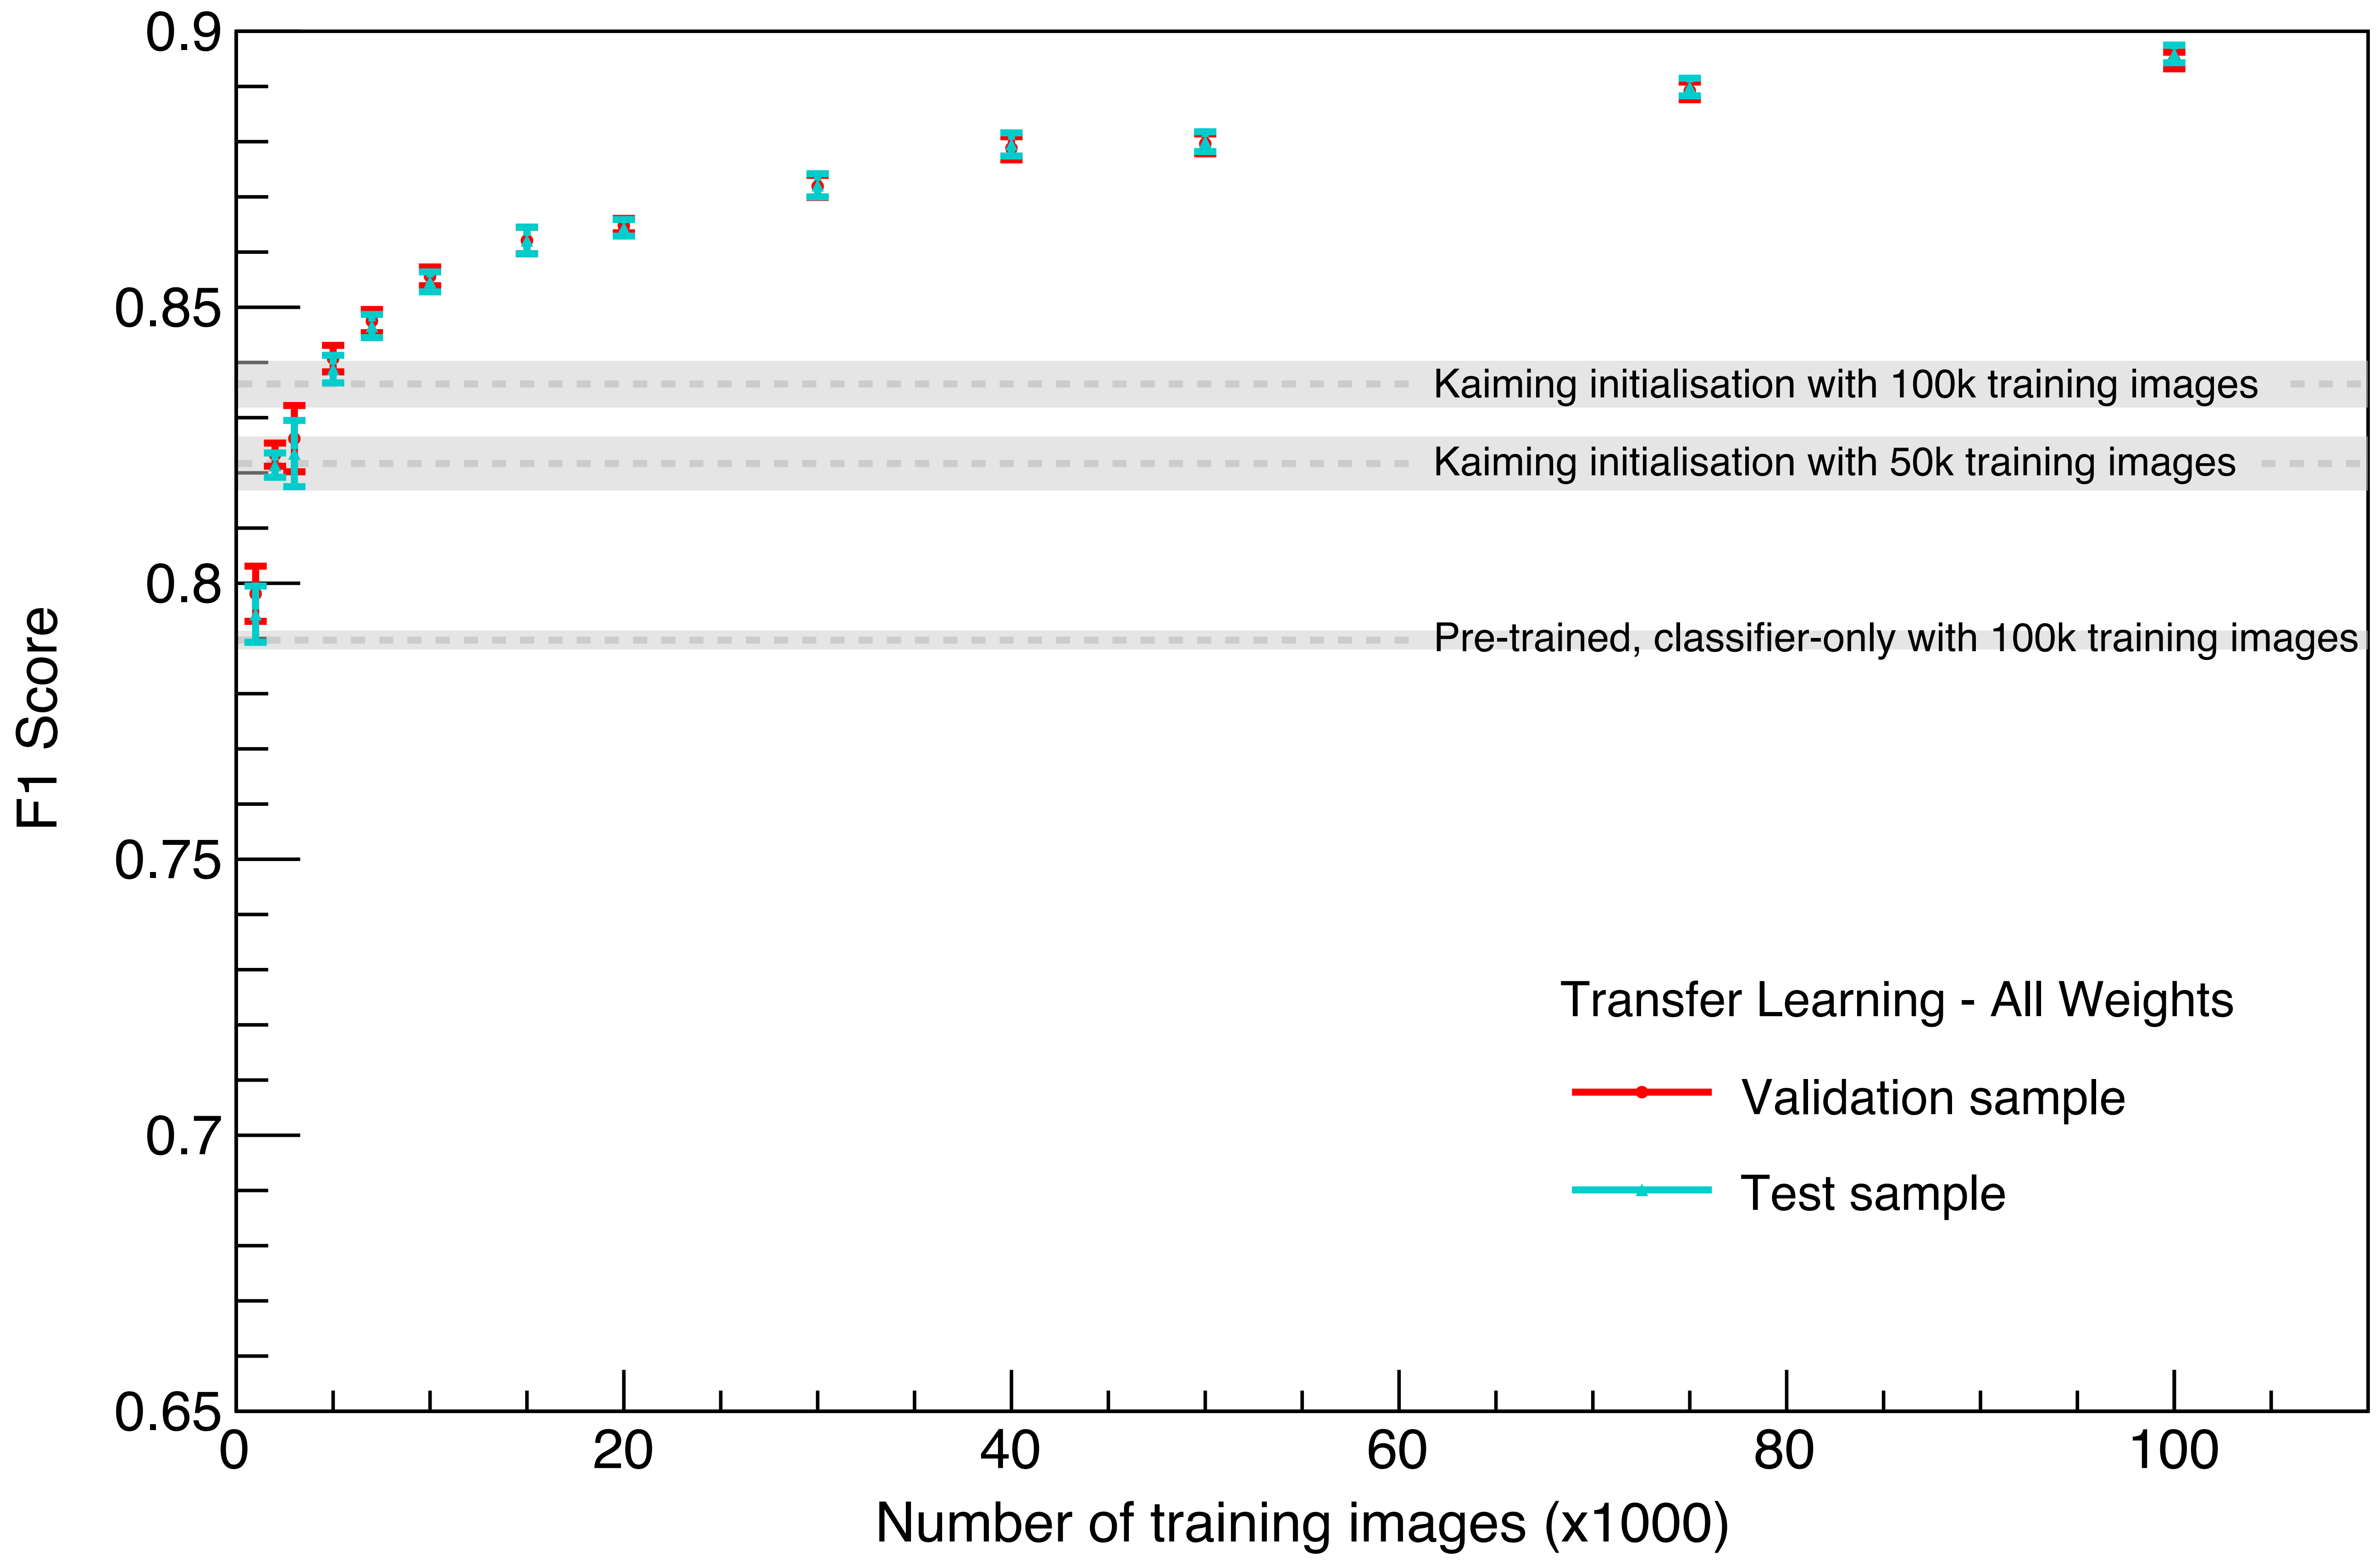

Results: TF vs random initialisation

- Also looked at the accuracy per class
  - We see improvements in each class individually

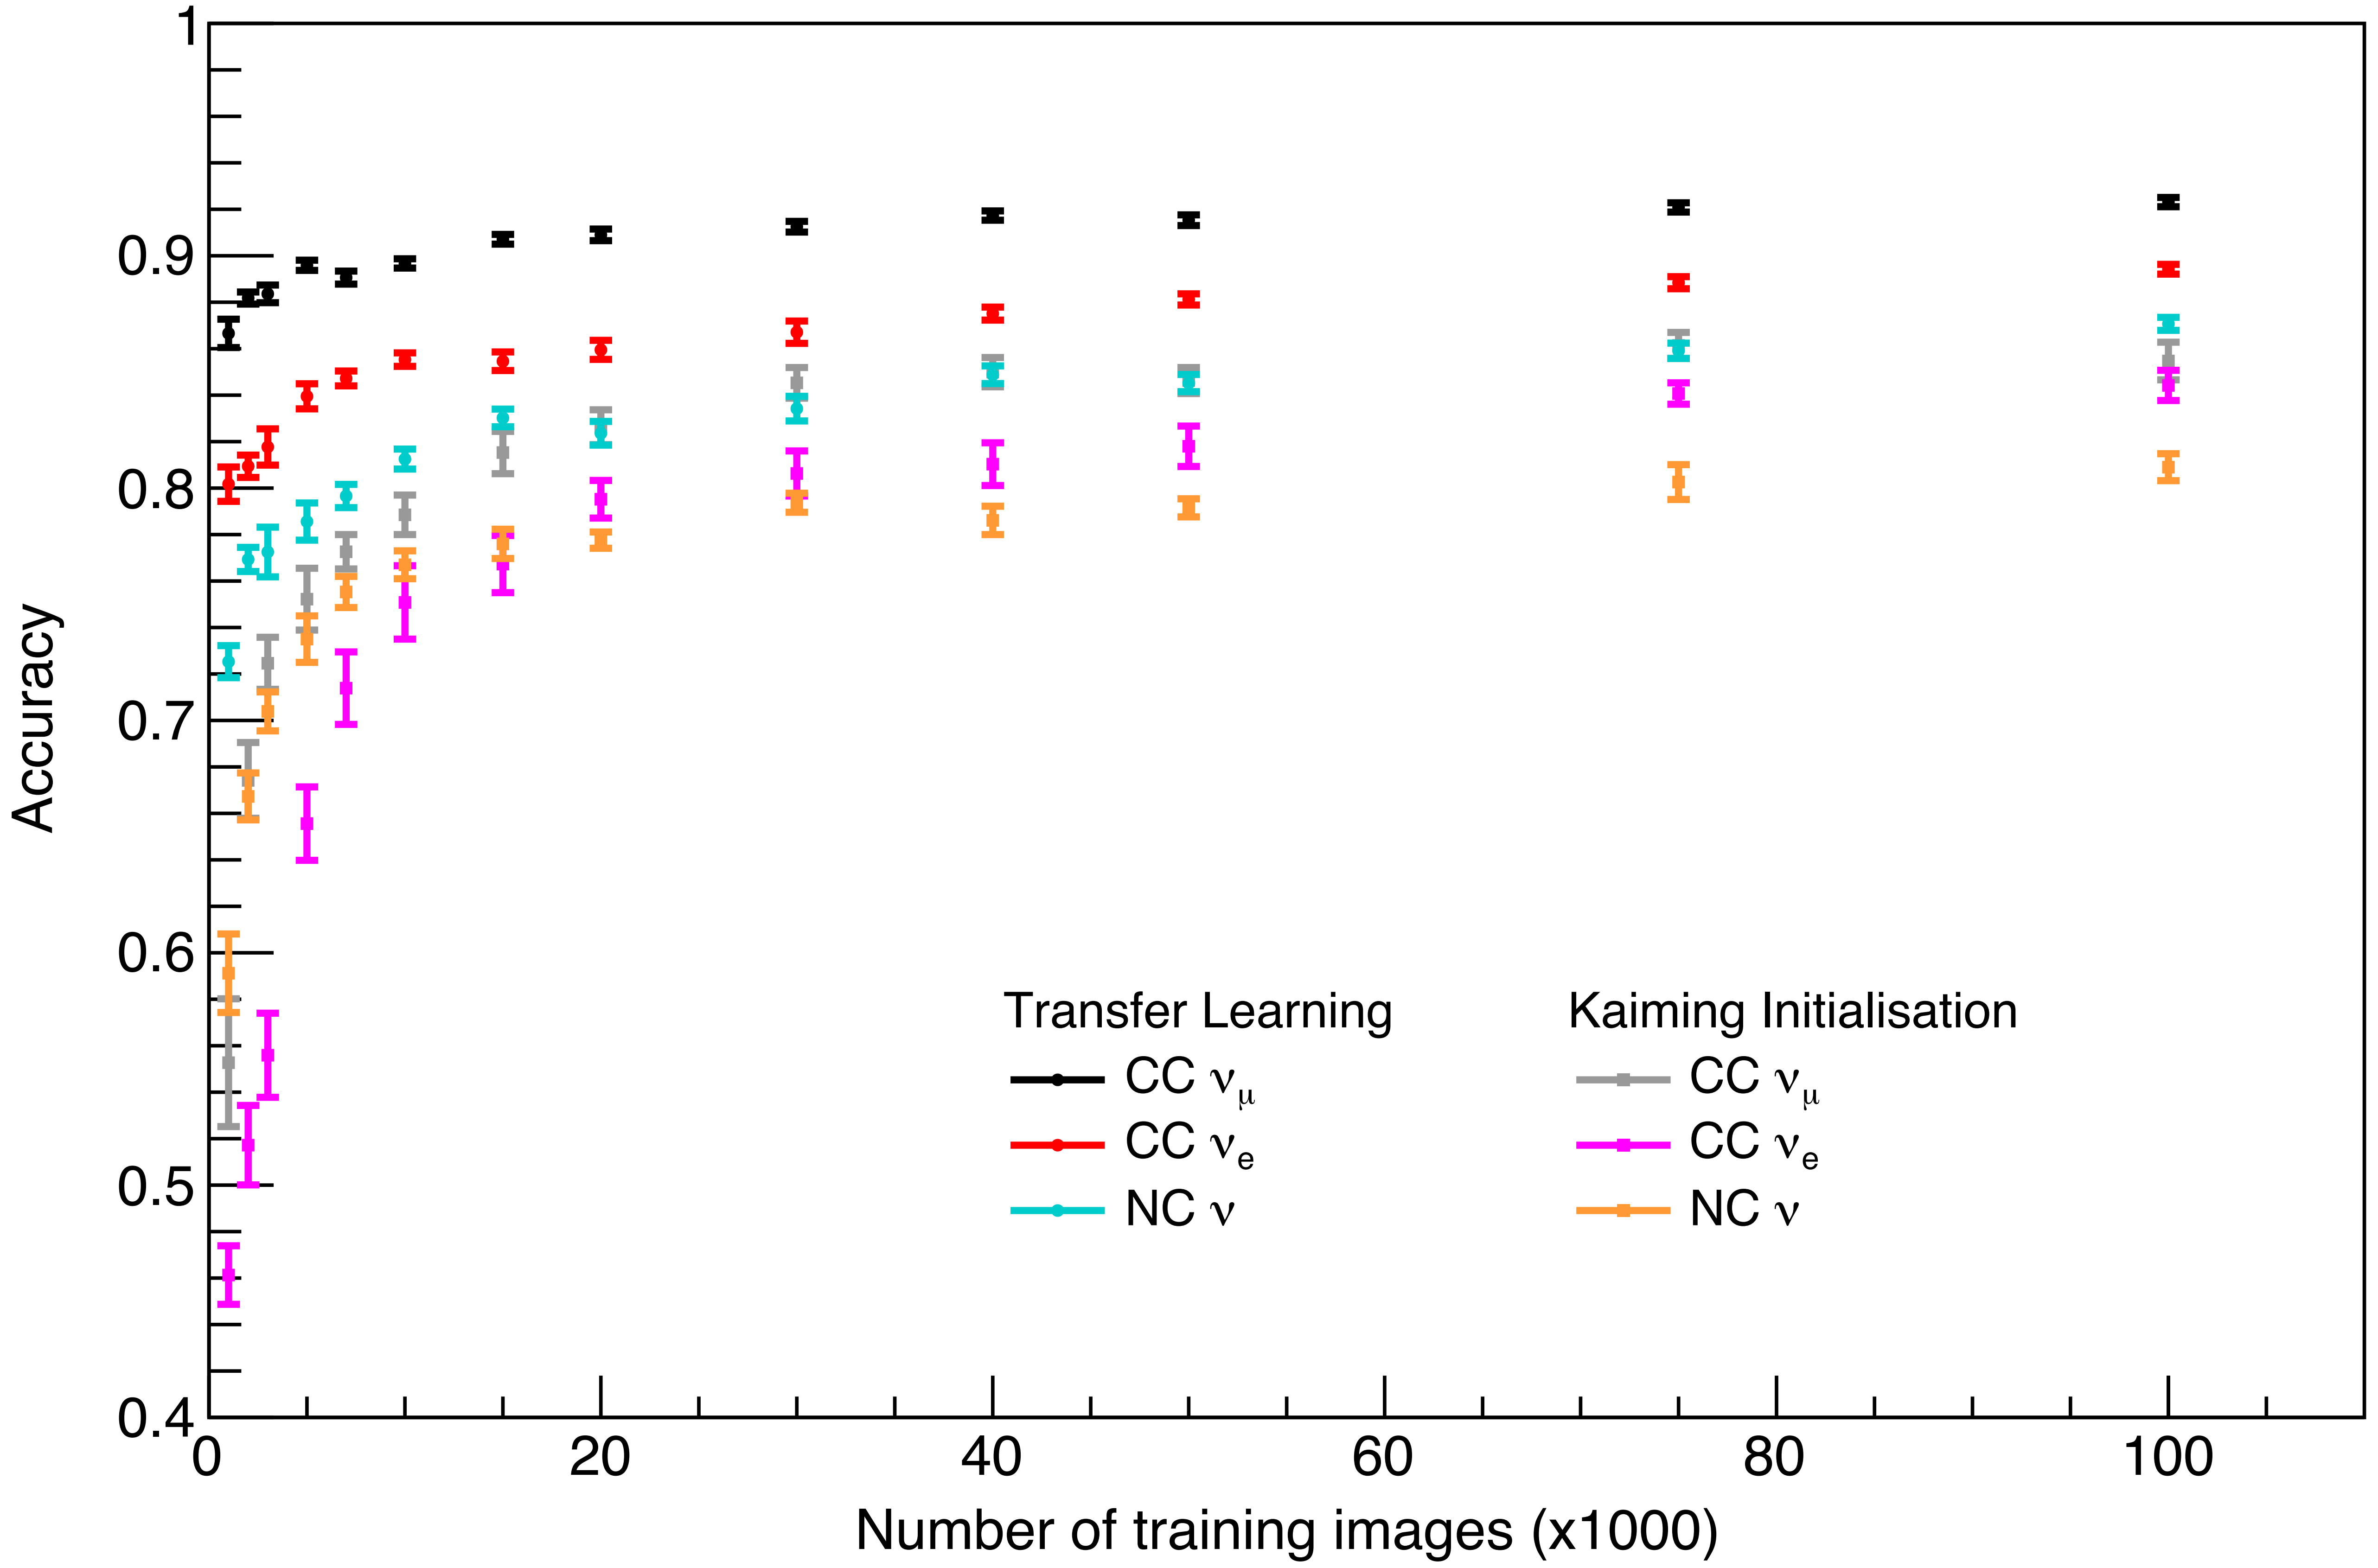

Transfer Learning in LArTPCs

- We also looked for potential biases between classes and as a function of energy
  - See reduced bias in both cases using transfer learning
  - Plots show examples from training with 100k events

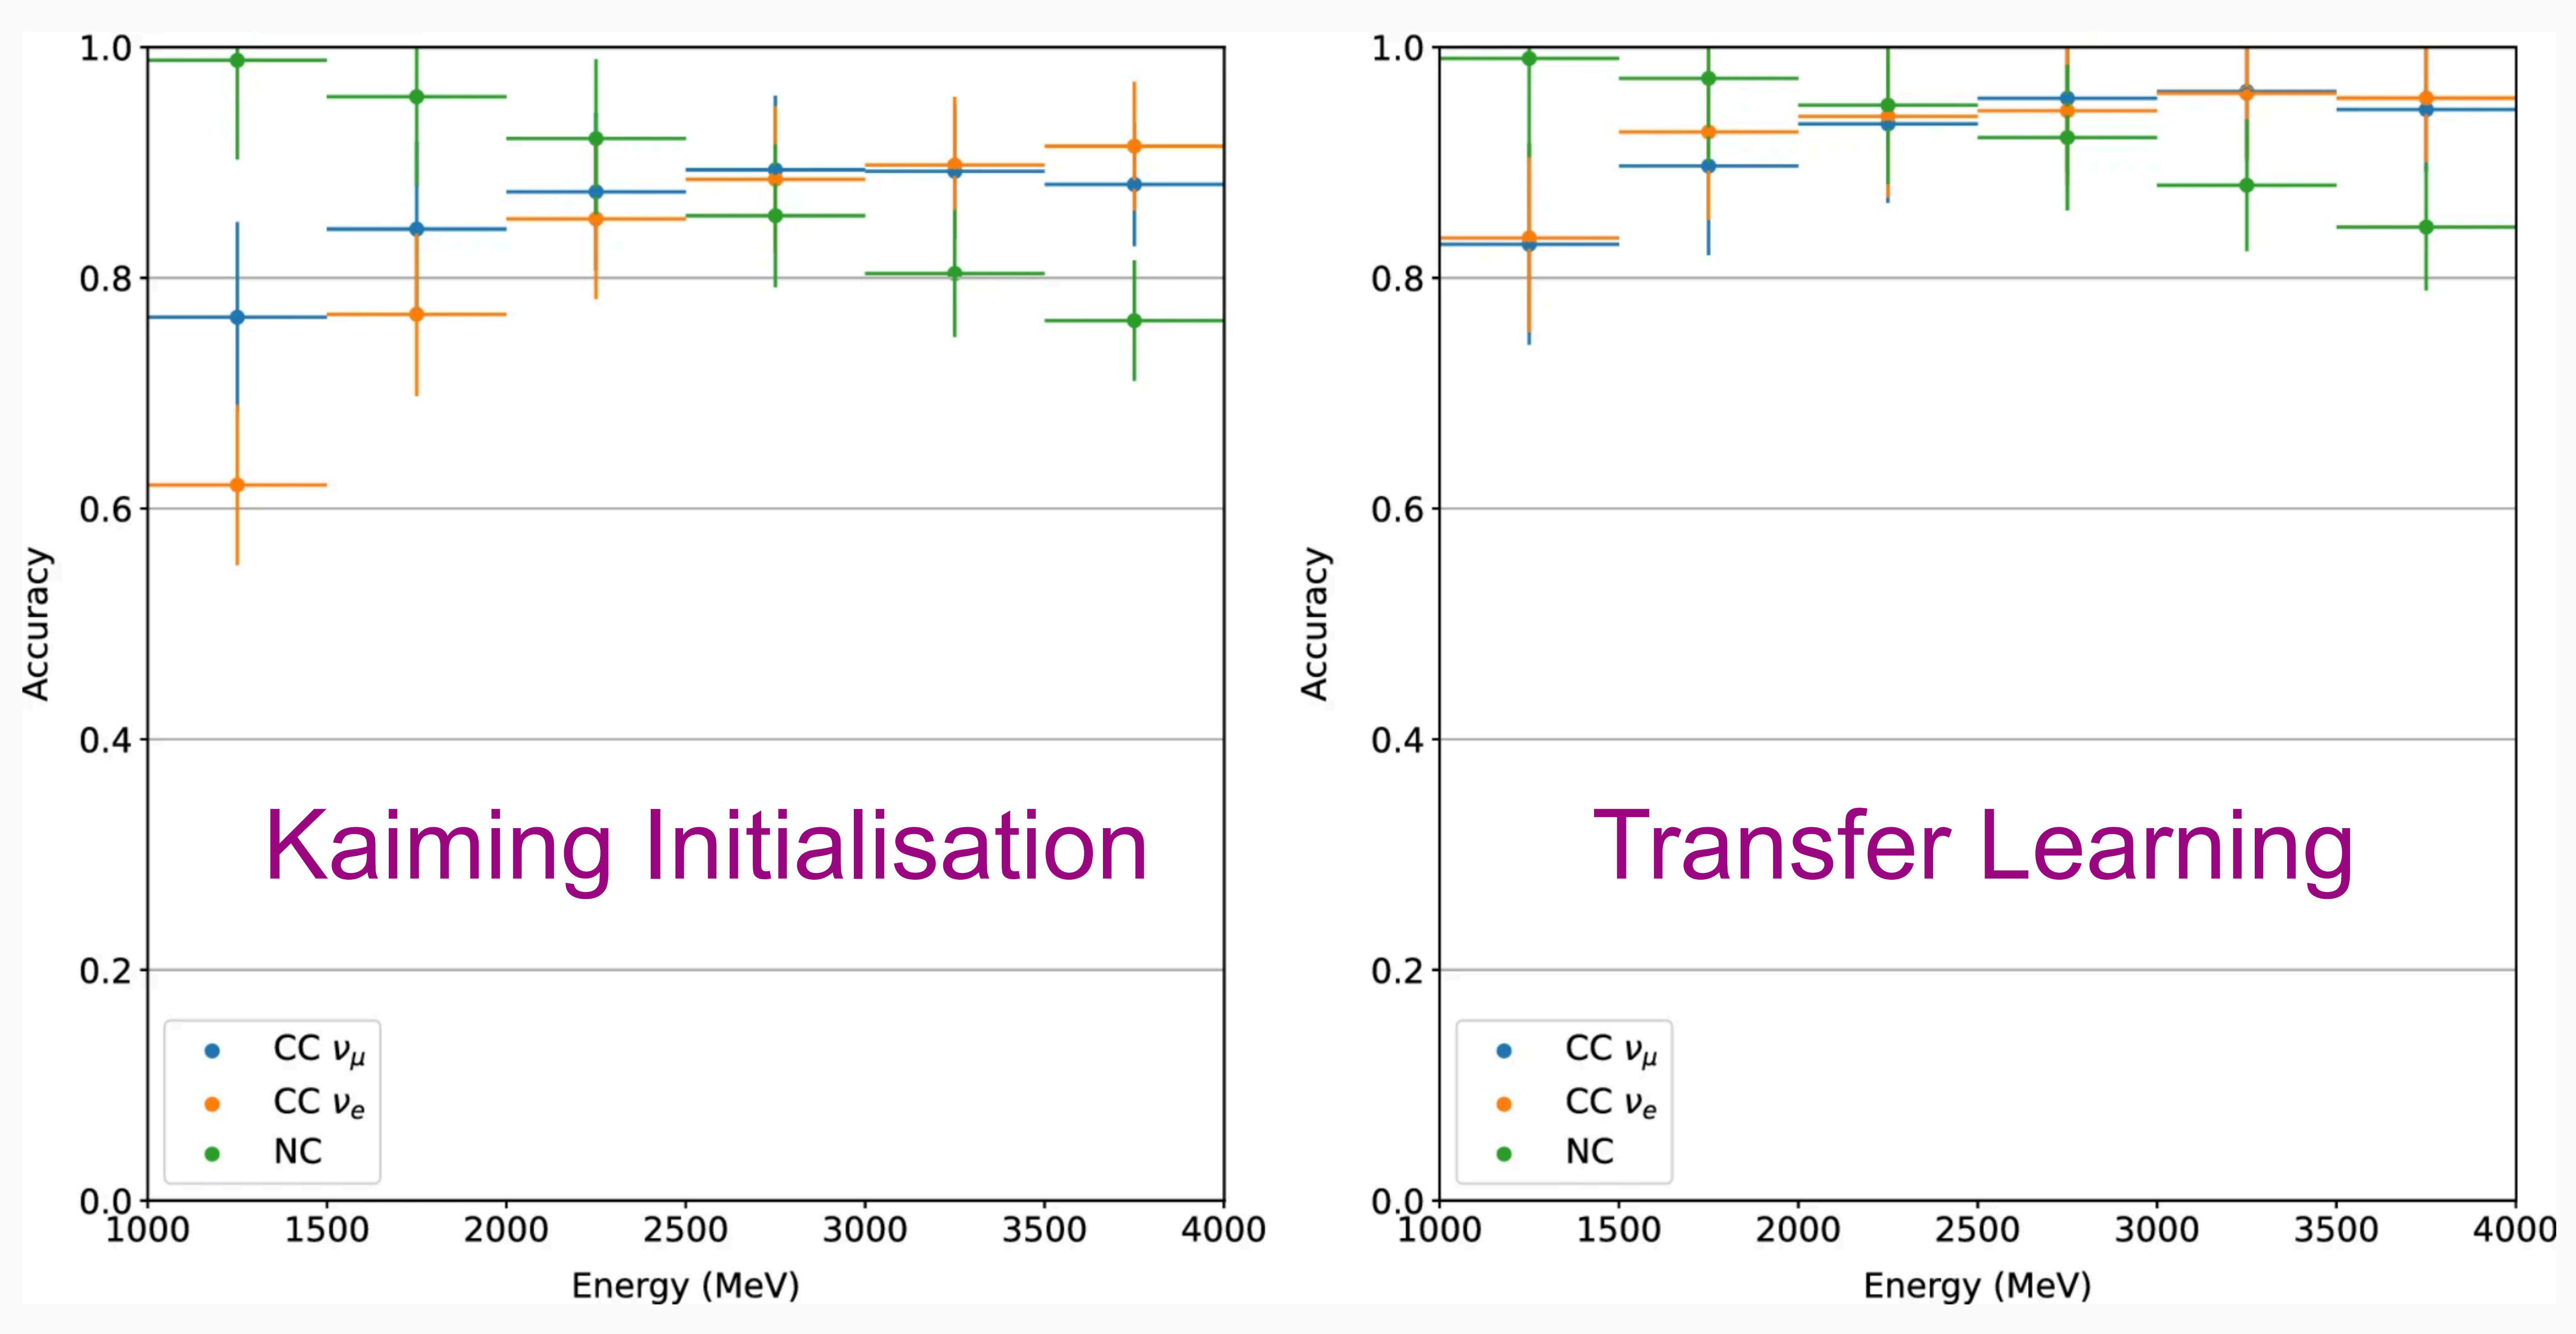

Transfer Learning in LArTPCs

- Also looked at the effect of freezing different layer weights
  - Layers 1 to 4 here correspond to the ResNet blocks
  - As a minimum we have to train the classifier (dense layer)
  - The difference between Layer 1 and All Weights is the first convolutional layer
    - No difference in performance is seen when the first layer weights can be fine-tuned
    - The ImageNet-trained first convolutional layer extracts all the information needed to classify our neutrino events

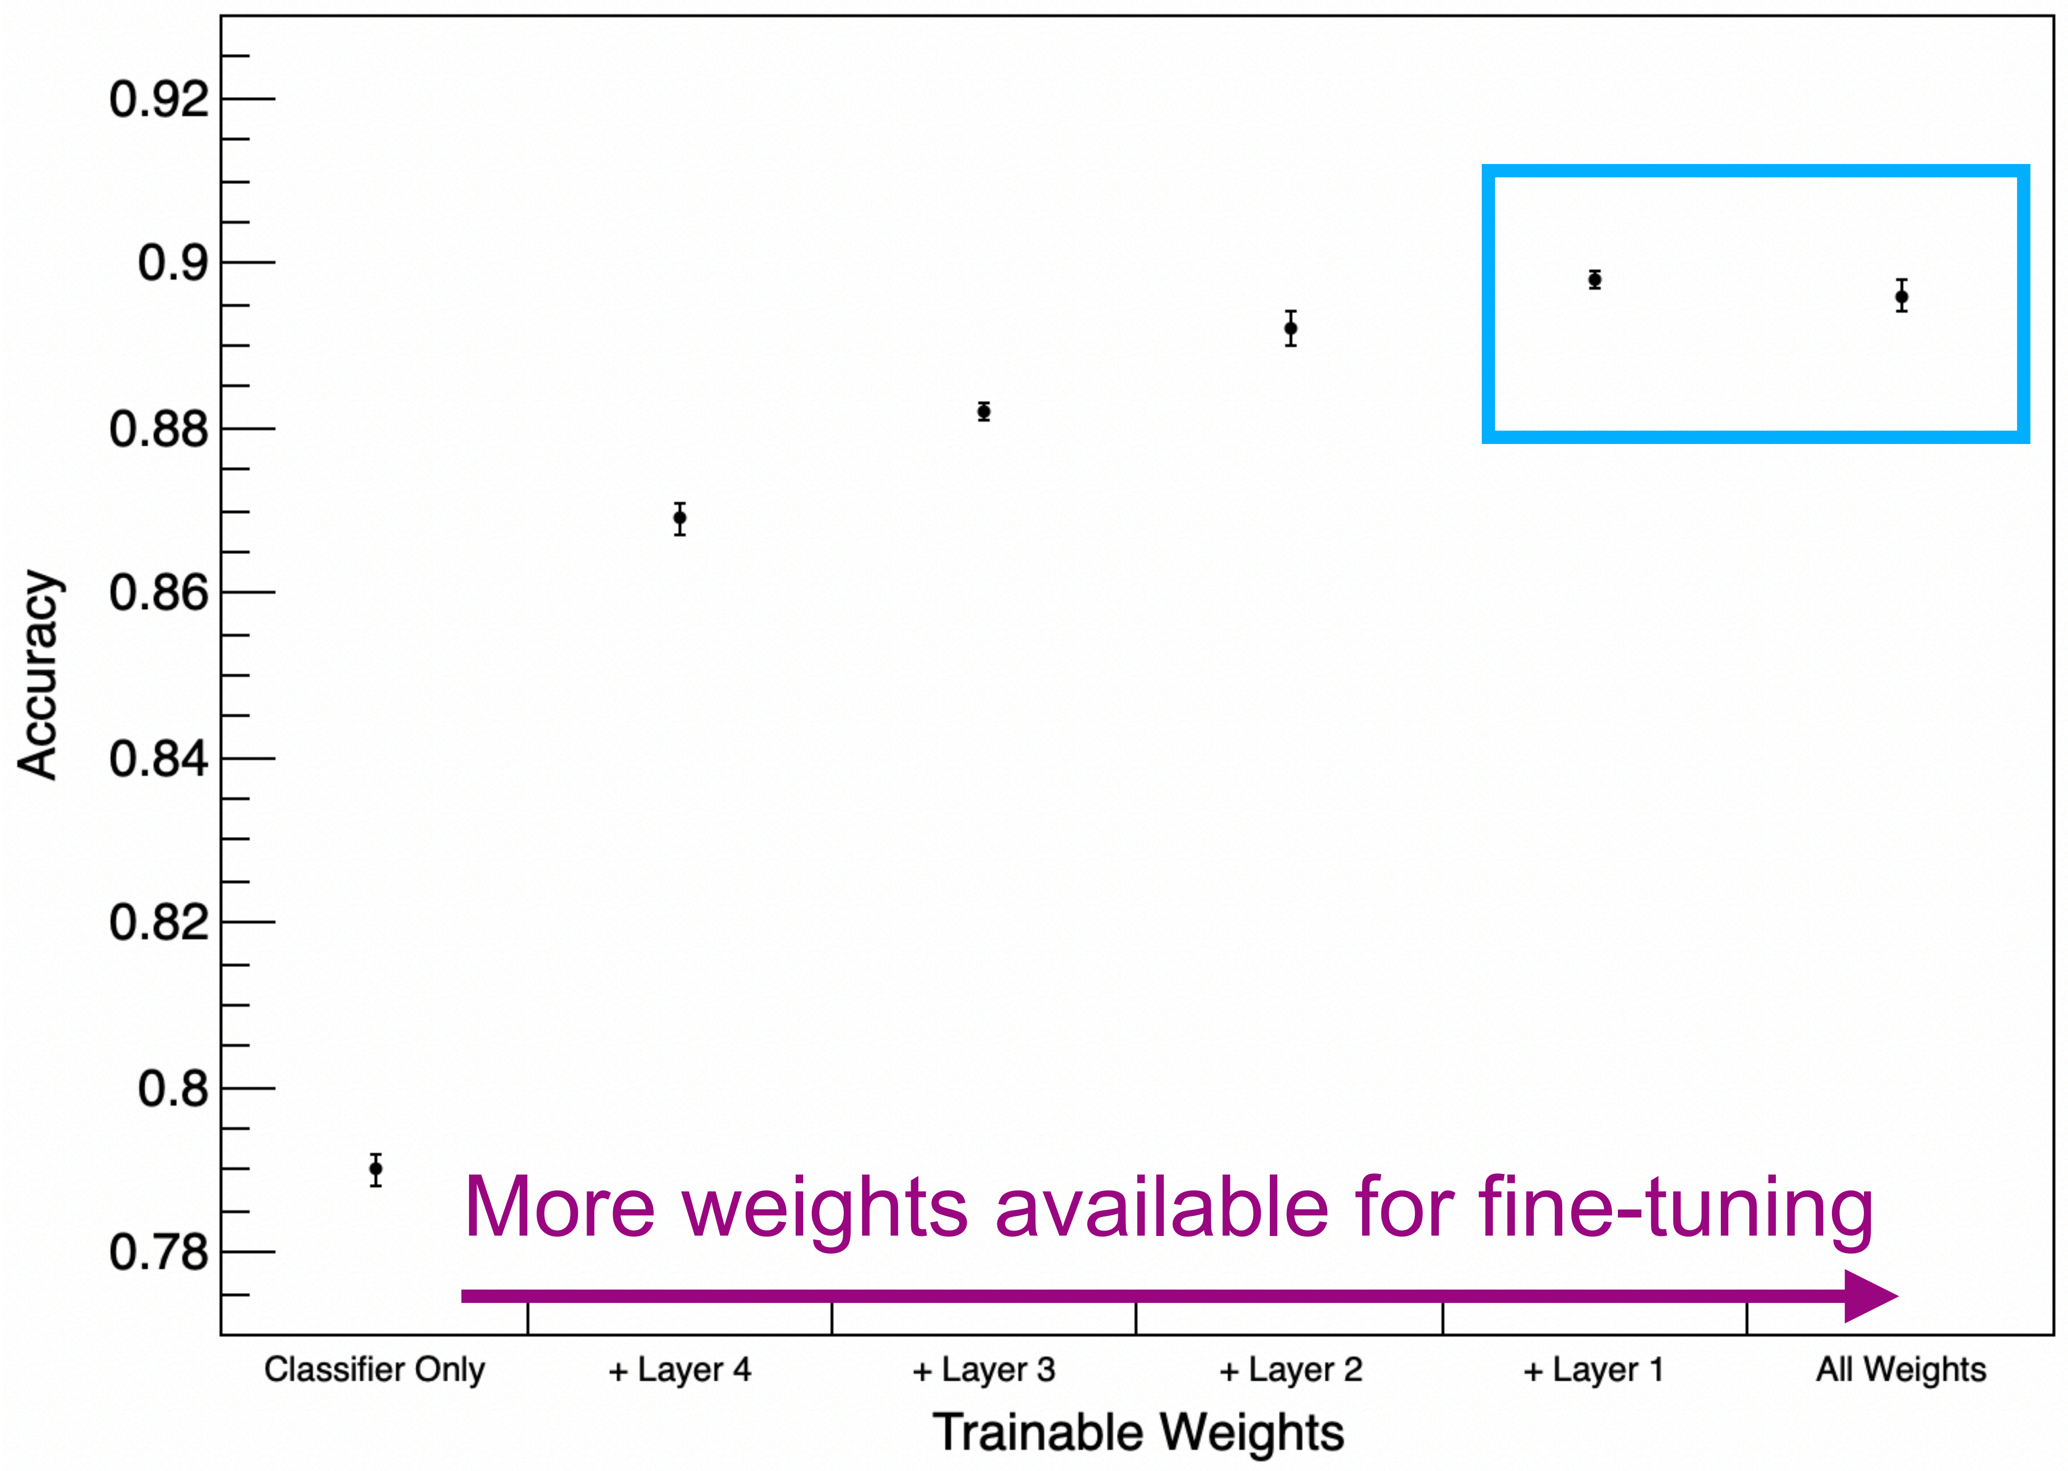

Conclusions

- Use of CNNs in neutrino physics is now well-established
  - Used for many use cases including event classification and semantic segmentation
- I think the coming years will show a focus on robustness
  - There isn't much physics gain in going from 95% to 96% efficiency
  - There is a lot of impact understanding your analysis at the 10% level to the 5% level
- Transfer learning looks to be a promising approach in some cases
  - Good performance with low numbers of training examples
  - Can help reduce computational burdens

Conclusions

- One note of caution: CNNs aren't always the right tool for the job!
  - If you find yourself needing to make complex projections to format your data as an image then using a CNN might not be the best approach
  - Tomorrow you'll see other approaches for differently structured data

Thank you… any questions?

Bonus Picture: A doe and her two fawns outside my window in the Stanford Guest House at 6:30am. A perk of jet lag, perhaps?

Additional references

- Study of using a DANN to reduce bias in ICARUS event filtering: https://doi.org/10.1103/PhysRevD.105.112009

- Some reviews:
  - P. Calafiura, D. Rousseau, K. Terao, Artificial Intelligence for High Energy Physics, World Scientific, 2022
  - A. Radovic et al. Machine learning at the energy and intensity frontiers of particle physics, Nature volume 560, pages 41–48 (2018)