Run 3493 Event 41075, Oct. 23rd, 2015

CONVOLUTIONAL NEURAL NETWORKS IN

MICROBOONE

Taritree Wongjirad
Tufts/MIT

DPF 2017

Outline

- Convolutional neural networks (CNNs) are a type of deep, feed-forward neural networks that have been successfully applied to a wide range of problems
- Discuss the ways MicroBooNE - a LArTPC detector - has been exploring the use of CNNs
- Three applications
  - Classification
  - Object detection
  - Semantic Segmentation

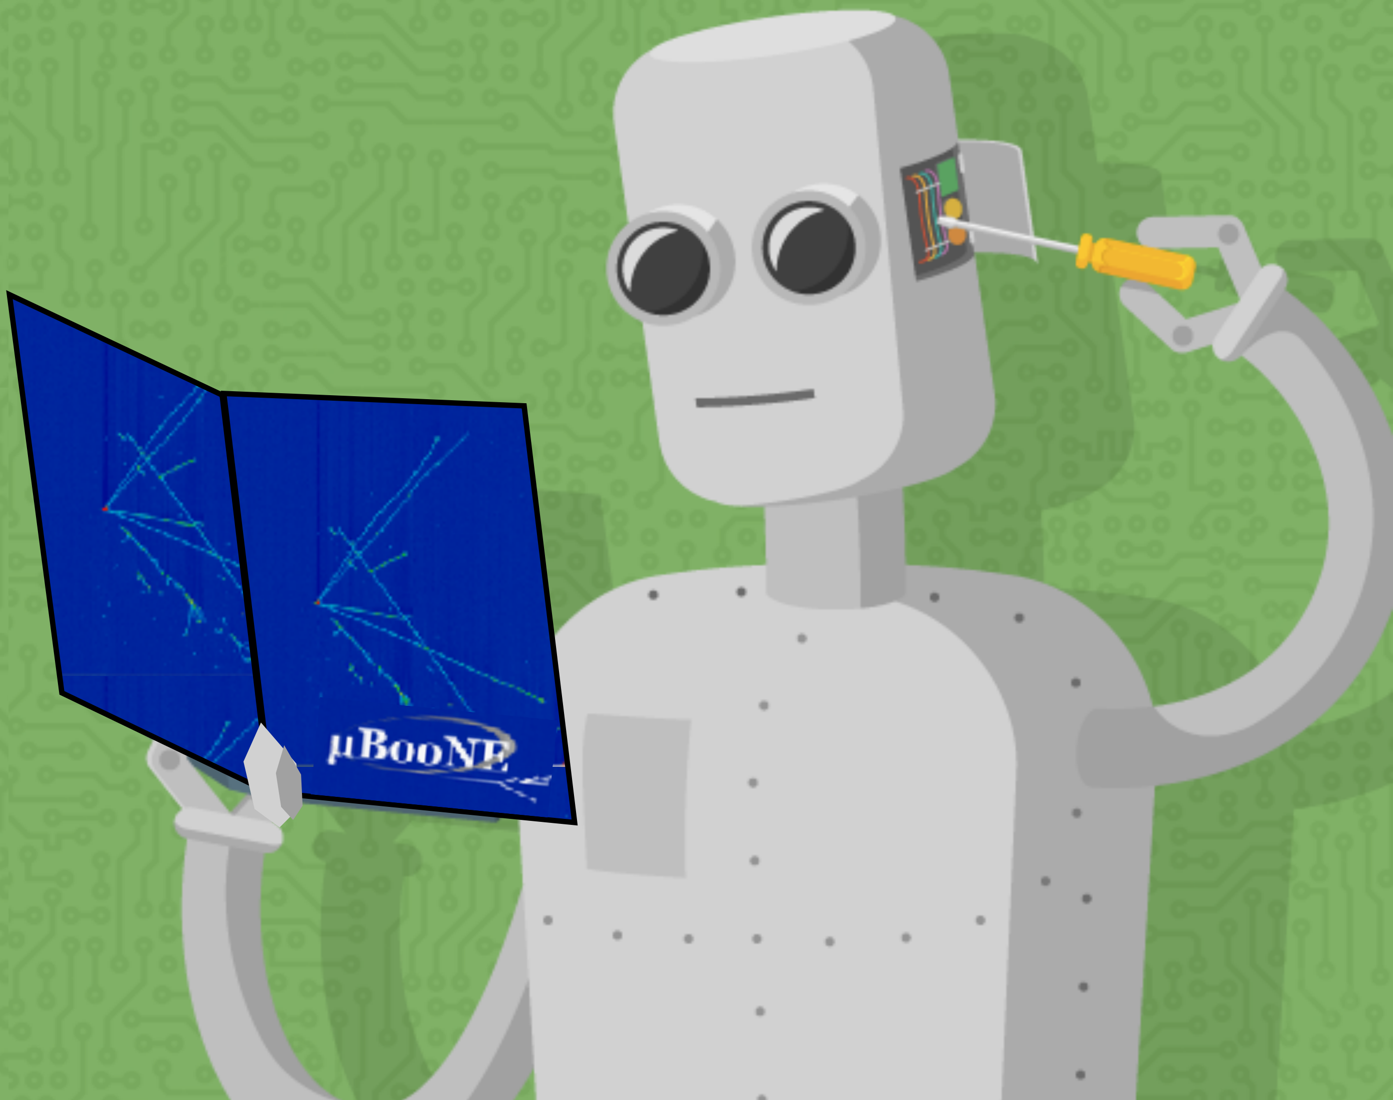

MICROBOONE GOALS

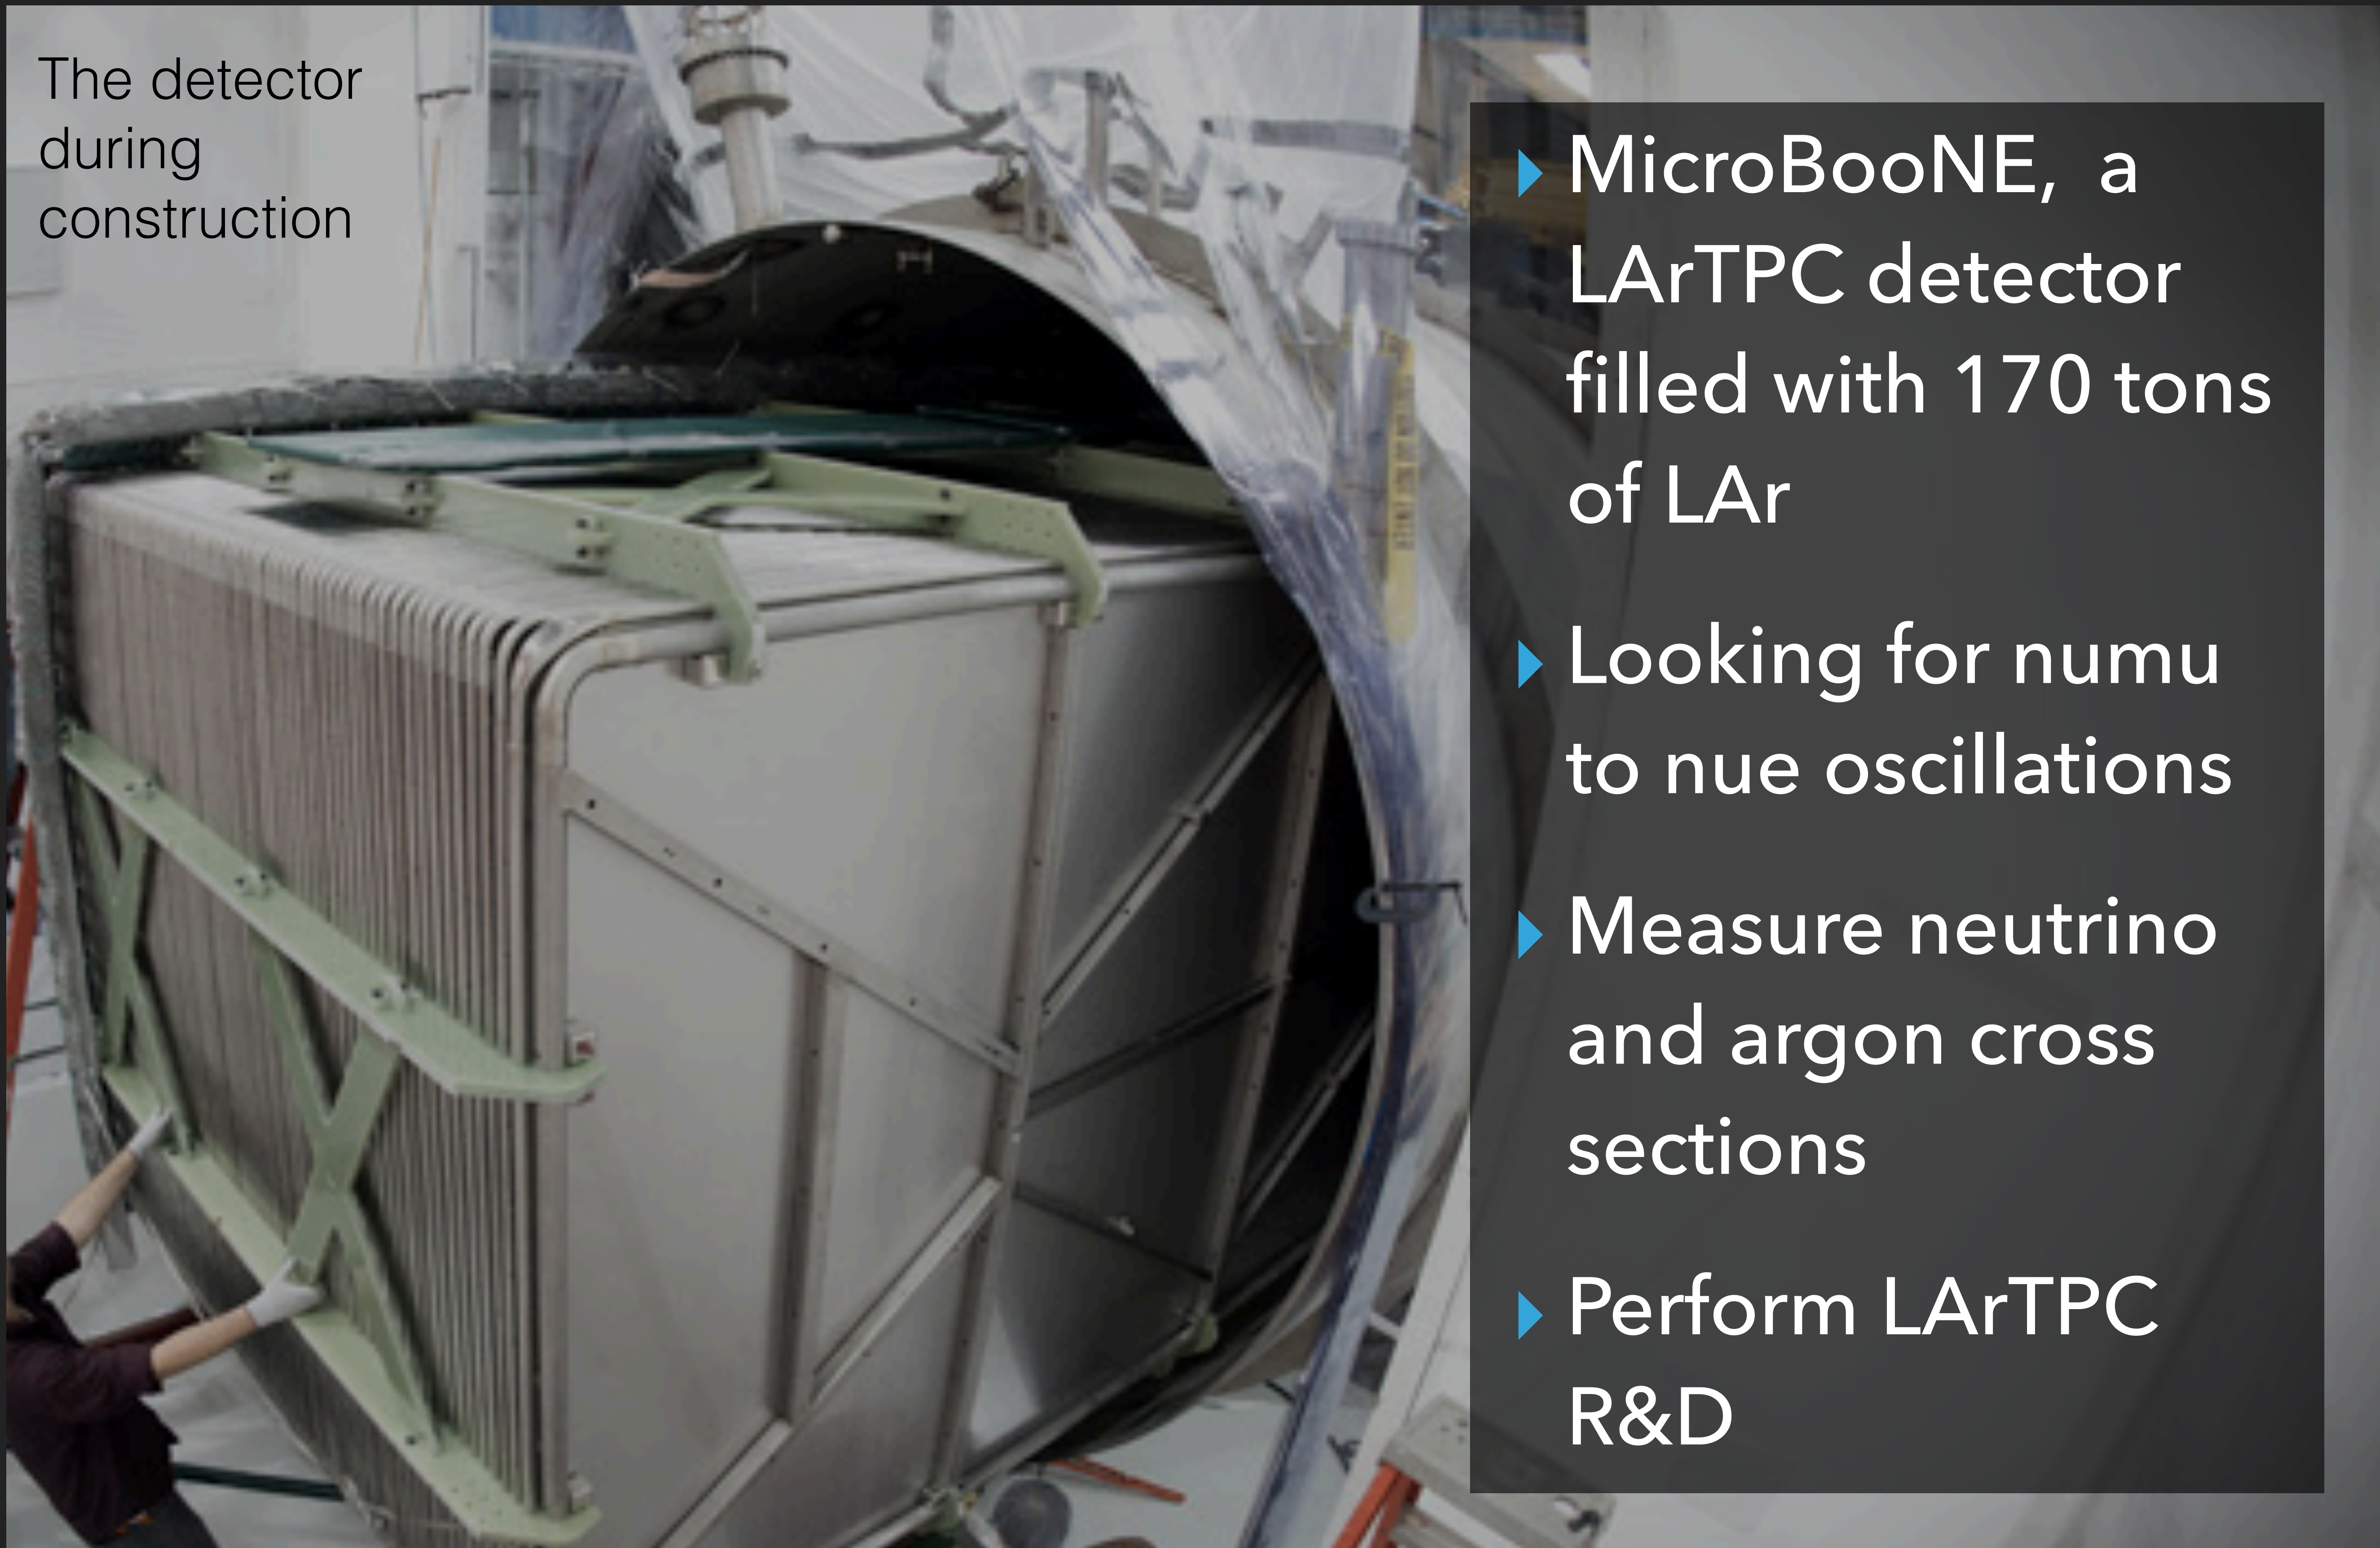

The detector during construction

- MicroBooNE, a LArTPC detector filled with 170 tons of LAr
- Looking for numu to nue oscillations
- Measure neutrino and argon cross sections
- Perform LArTPC R&D

MICROBOONE

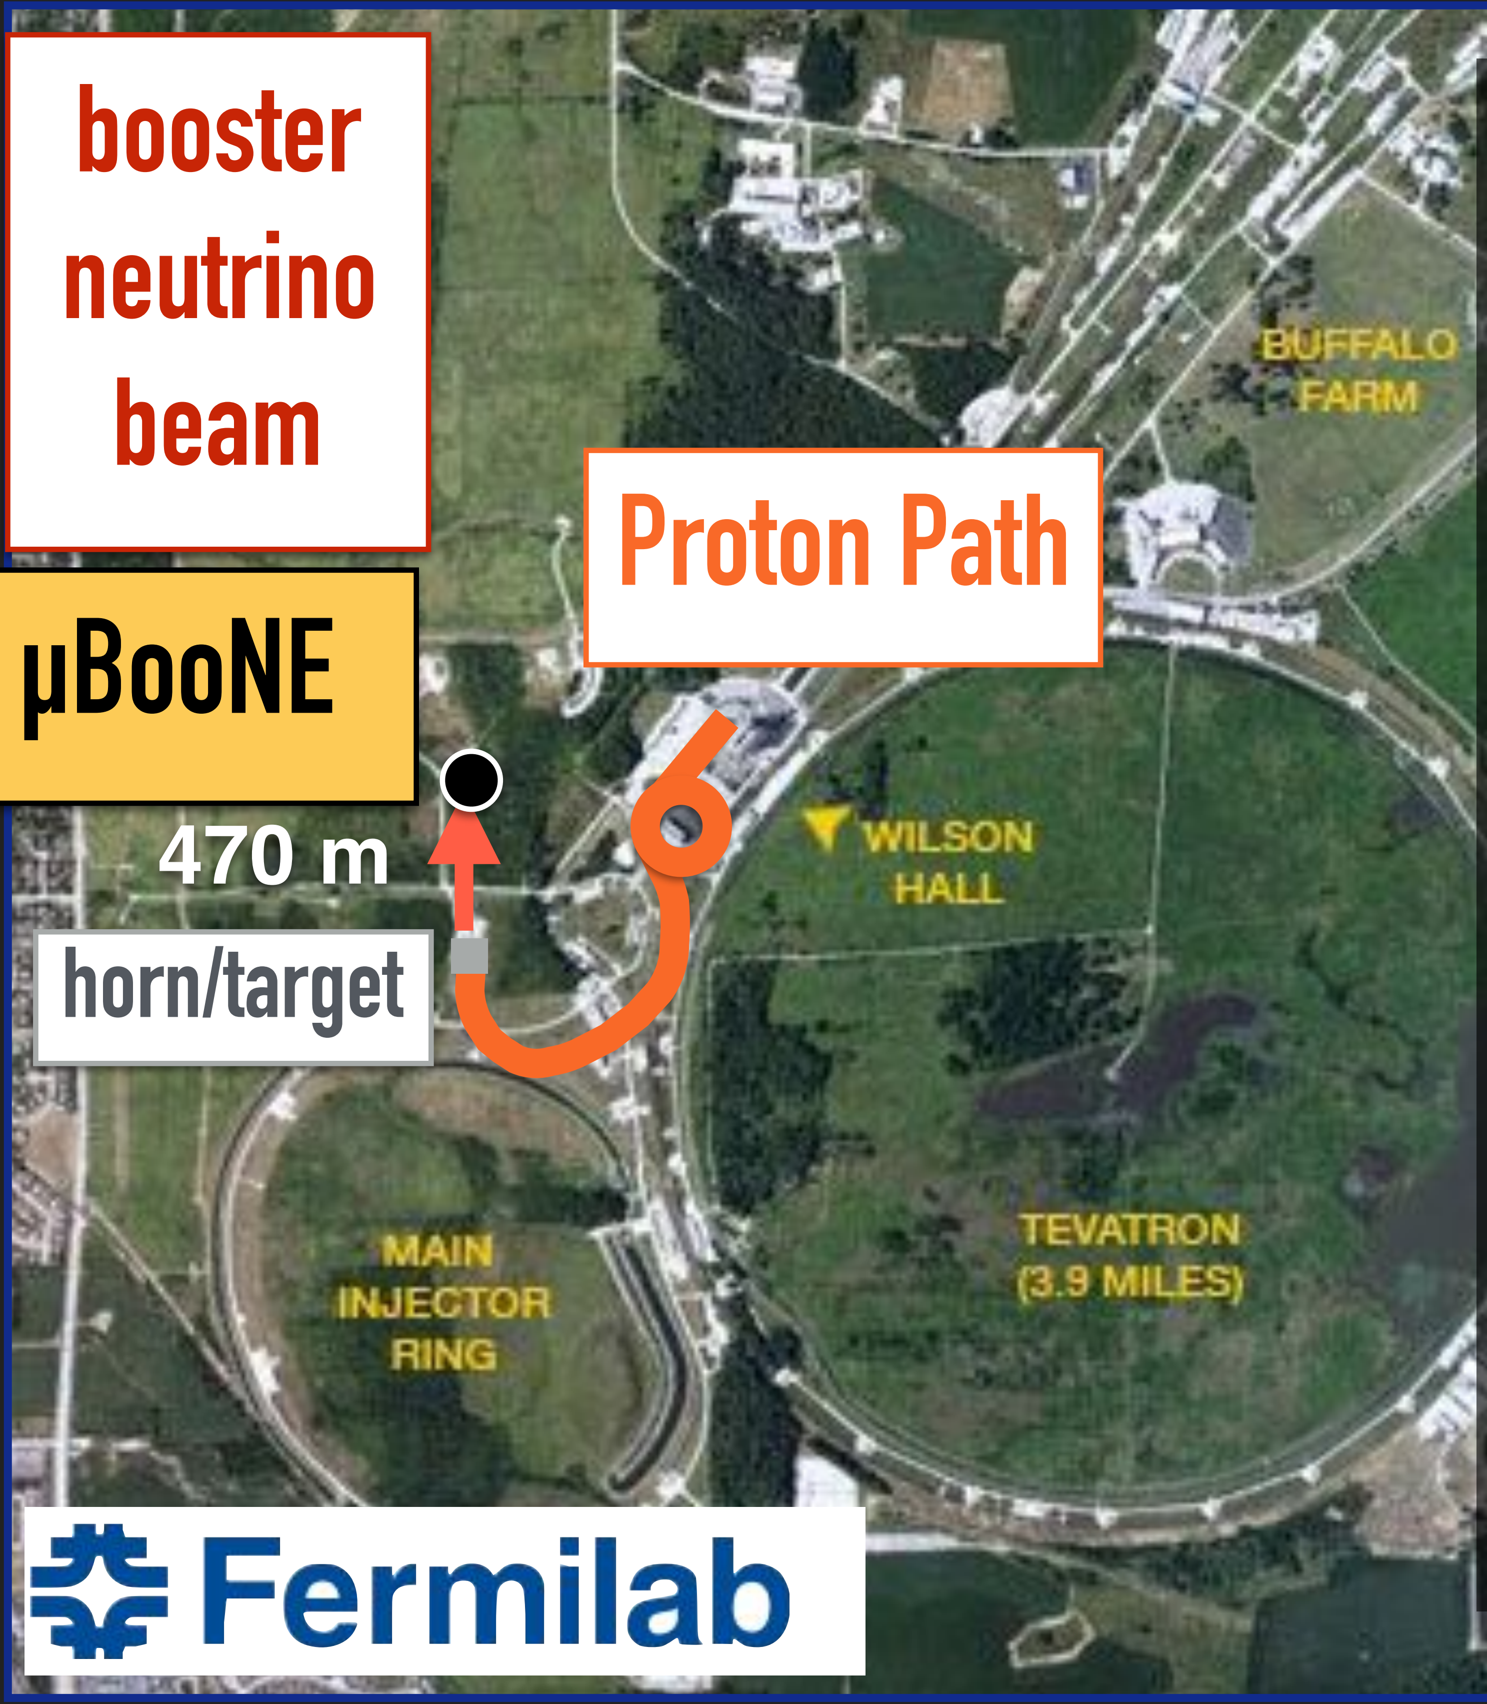

- MicroBooNE located here at FNAL
- Sits 470 m from the start of the Booster Neutrino Beam – produces mostly muon neutrinos

MICROBOONE EVENT

- Example neutrino event from the beam
- Lots of detail on location and amount of charge created in detector

RECONSTRUCTION

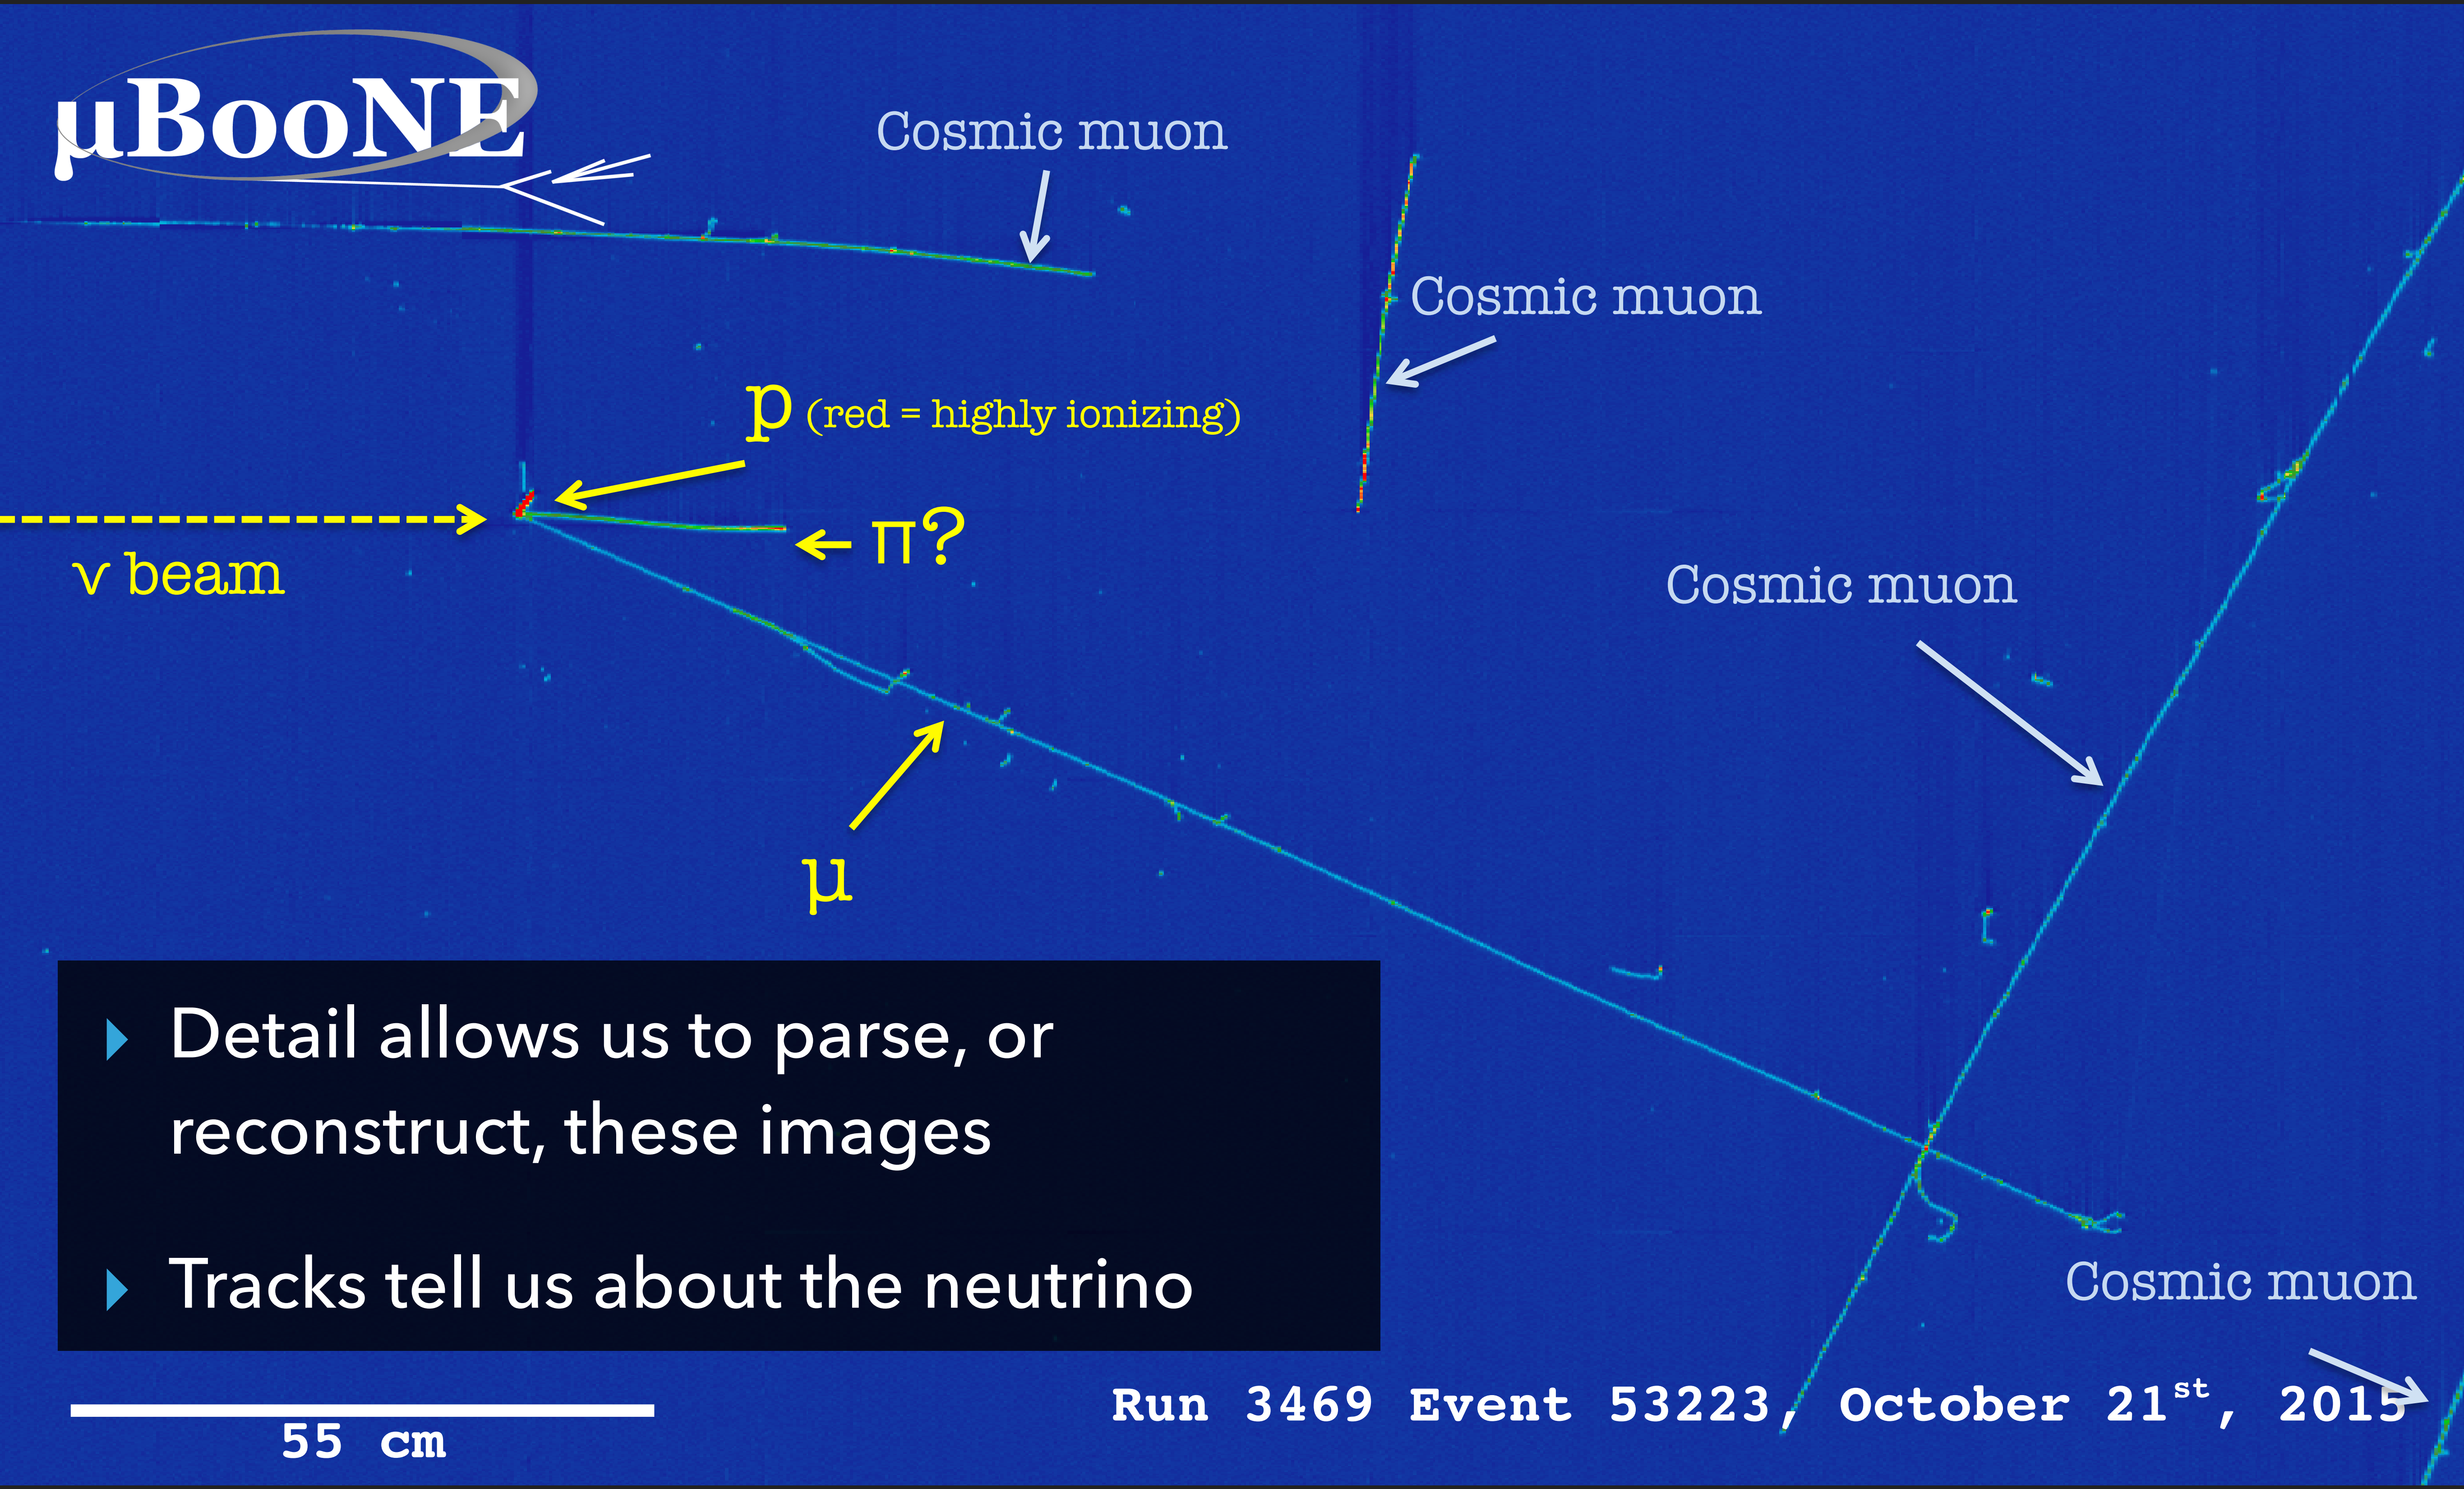

- Detail allows us to parse, or reconstruct, these images
- Tracks tell us about the neutrino

CHALLENGES

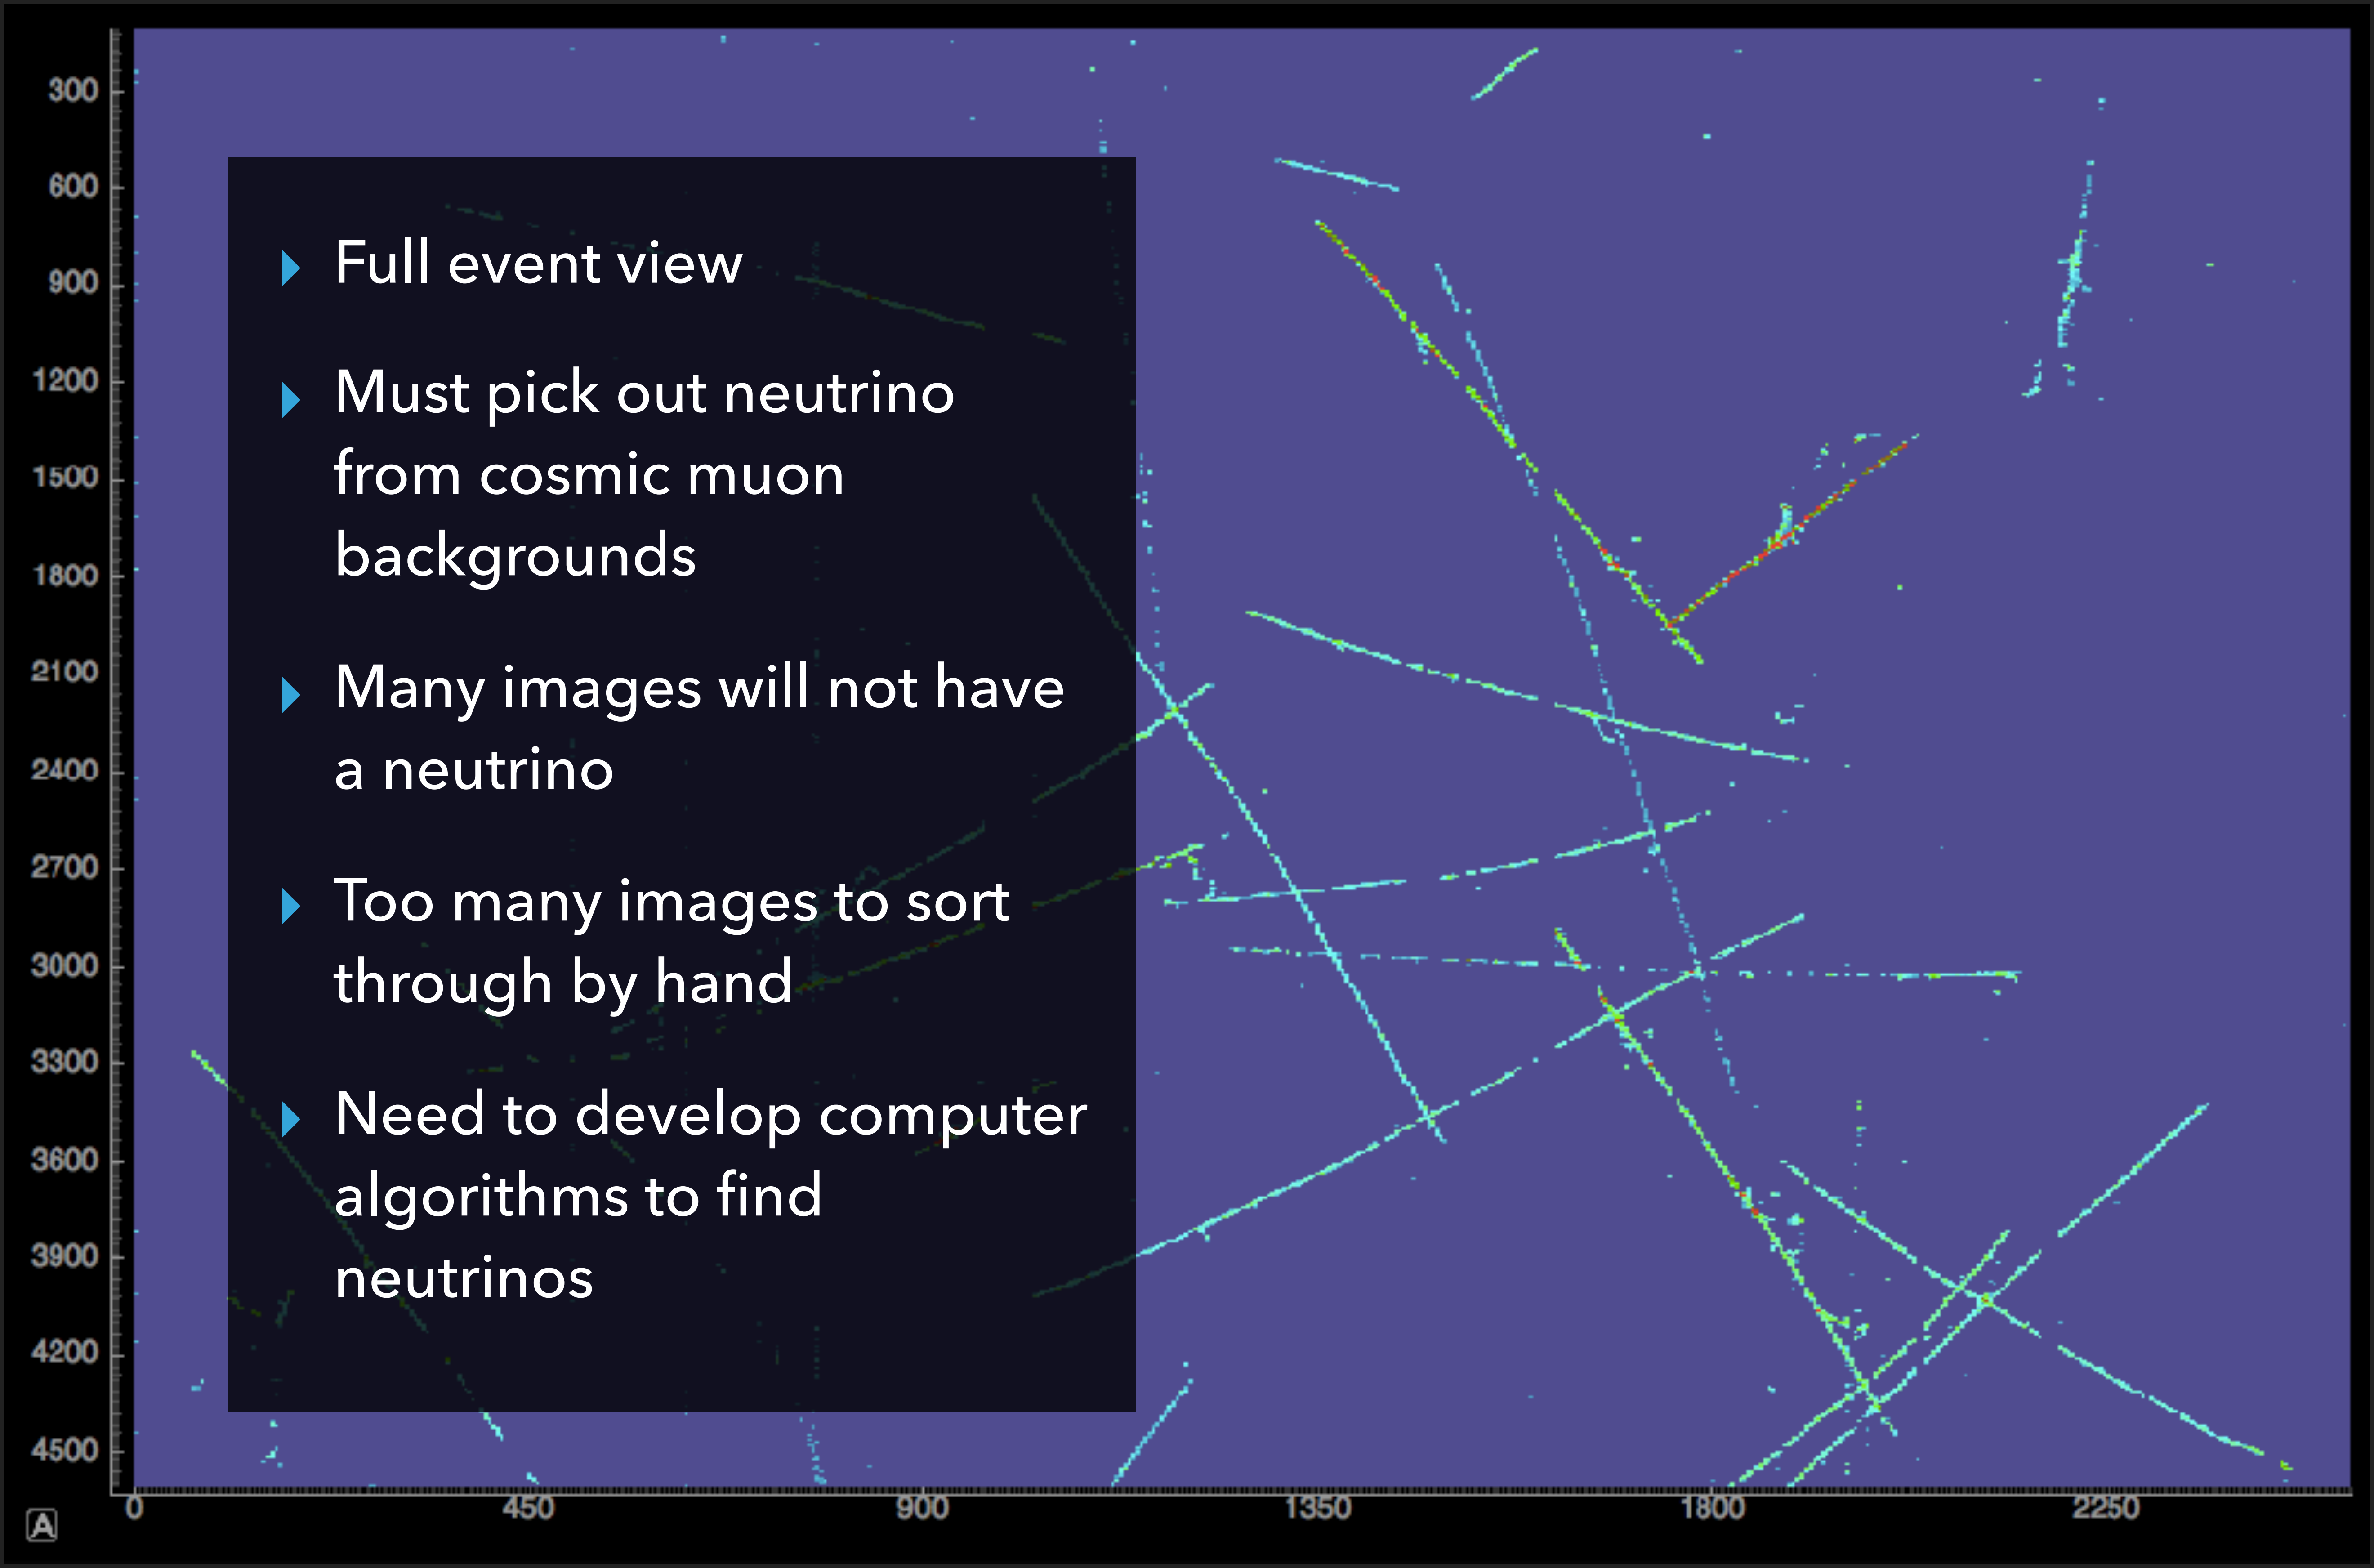

- Full event view
- Must pick out neutrino from cosmic muon backgrounds
- Many images will not have a neutrino
- Too many images to sort through by hand
- Need to develop computer algorithms to find neutrinos

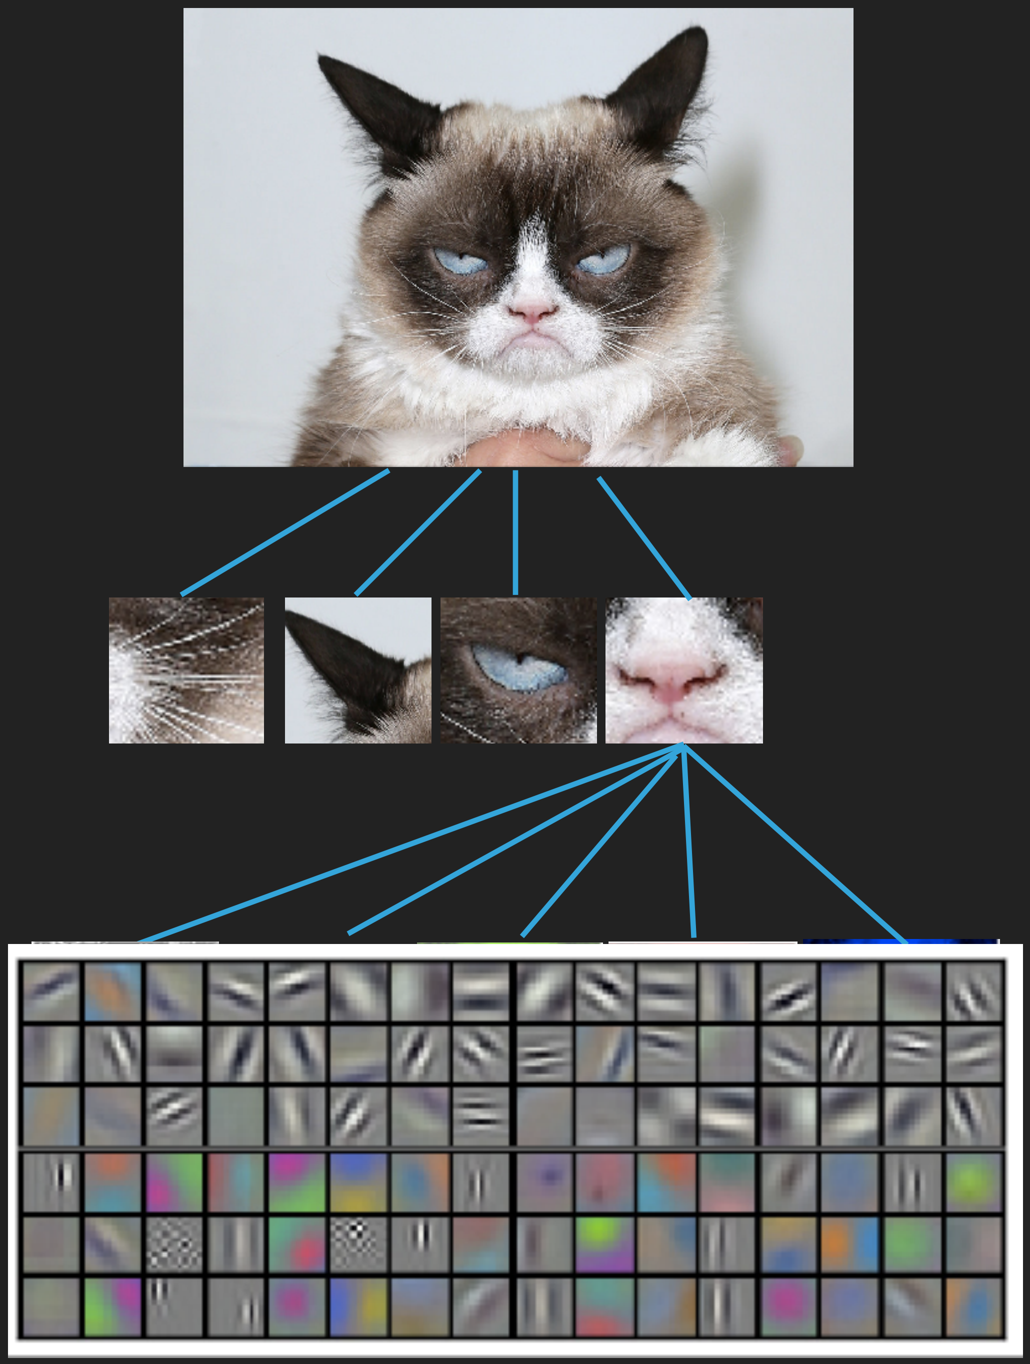

- To analyze an image, e.g. recognize as cat, decompose an object into a collection of small features
- Features composed of different patterns, lines and colors
- How to find the features and put them together?

CONVOLUTIONAL NEURAL NETWORKS

- Applying convolutional neural nets (CNN)
- Very adept at image analysis
- Primary advantages: scalable and generalizable technique
- Successfully applied to many different types of problems

Face detection

Video analysis for self-driving cars

Defeating humans at Go

CONVOLUTIONAL NEURAL NETWORKS

- CNNs differ from "traditional" neural nets in their structure
- CNN "neuron" looks for local, translation-invariant patterns among inputs

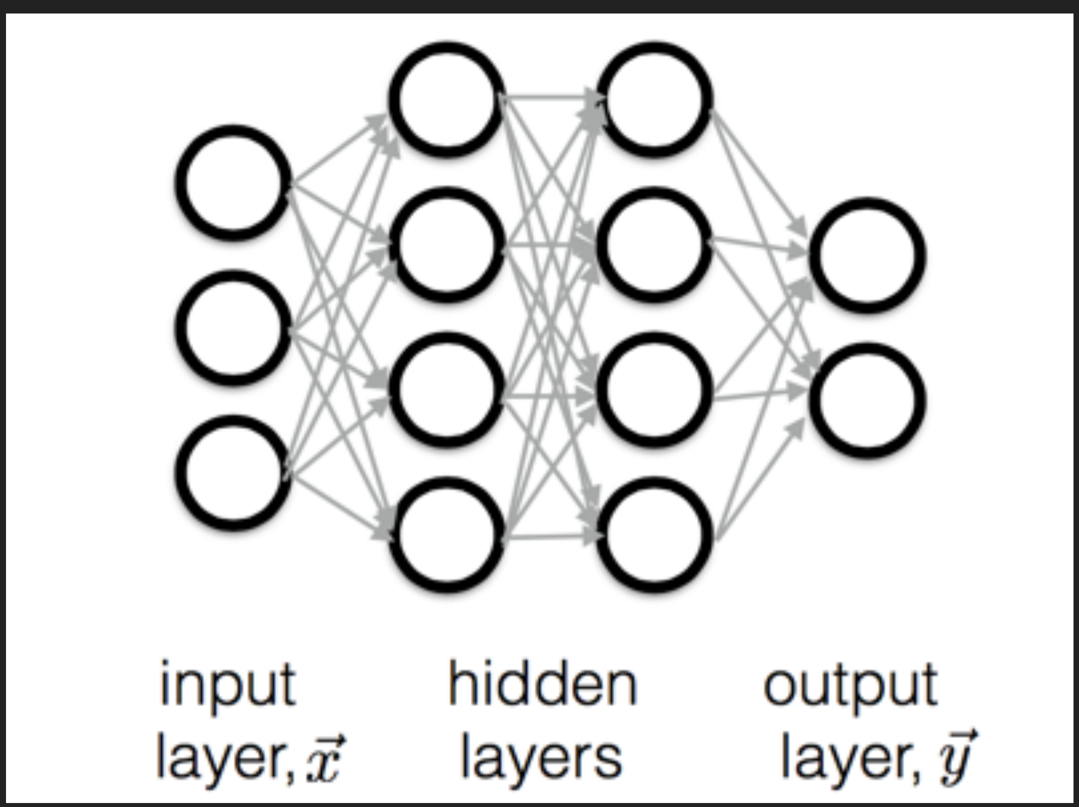

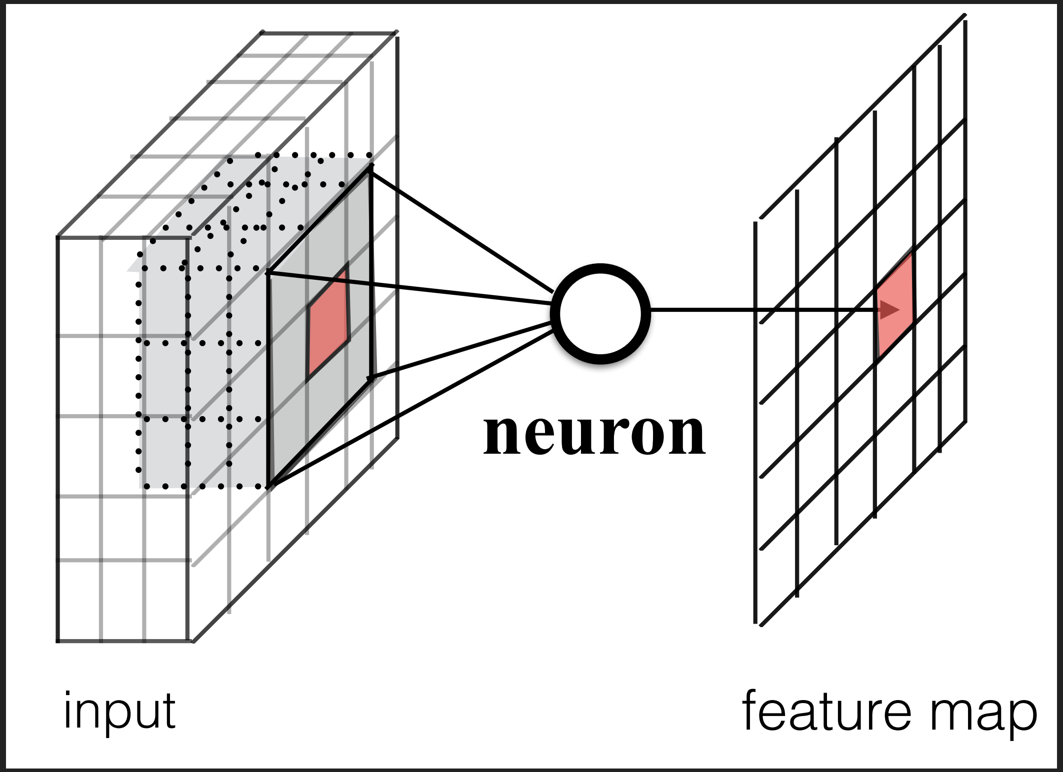

$$f_{i,j}(X) = \sigma\left(W_i \cdot X_j + b_i\right),$$

CONVOLUTIONAL FILTER

- Core operation in a CNN is the convolutional filter – identifies the location of patterns in an image
- Here regions of light and dark are where the pattern (or its inverse) matched well within the image

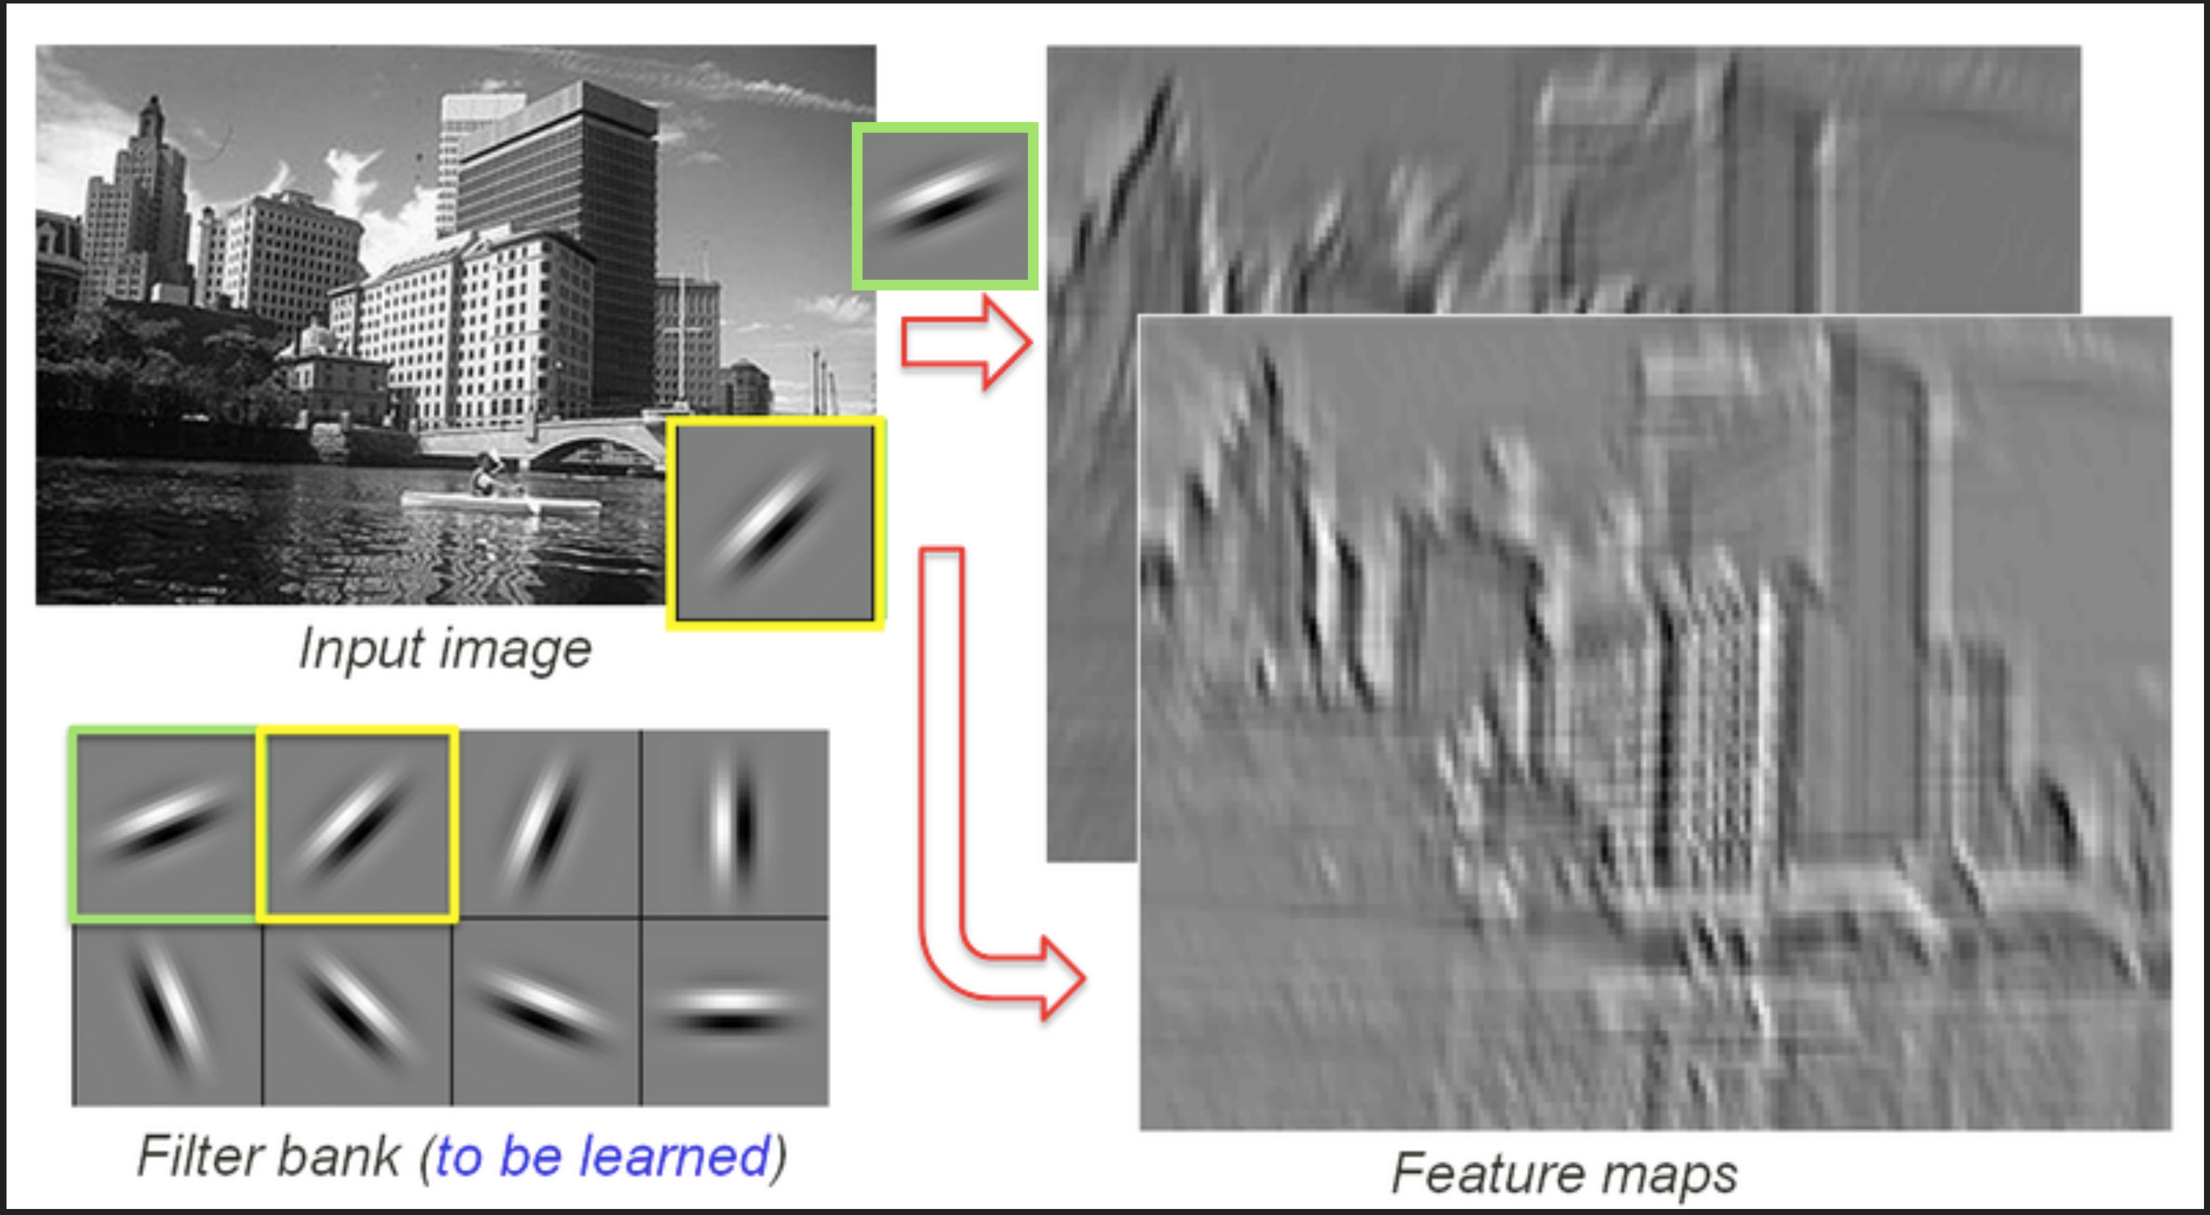

CONVOLUTIONAL FILTER

- one neuron produces one feature map
- operation takes as input an image and outputs an image

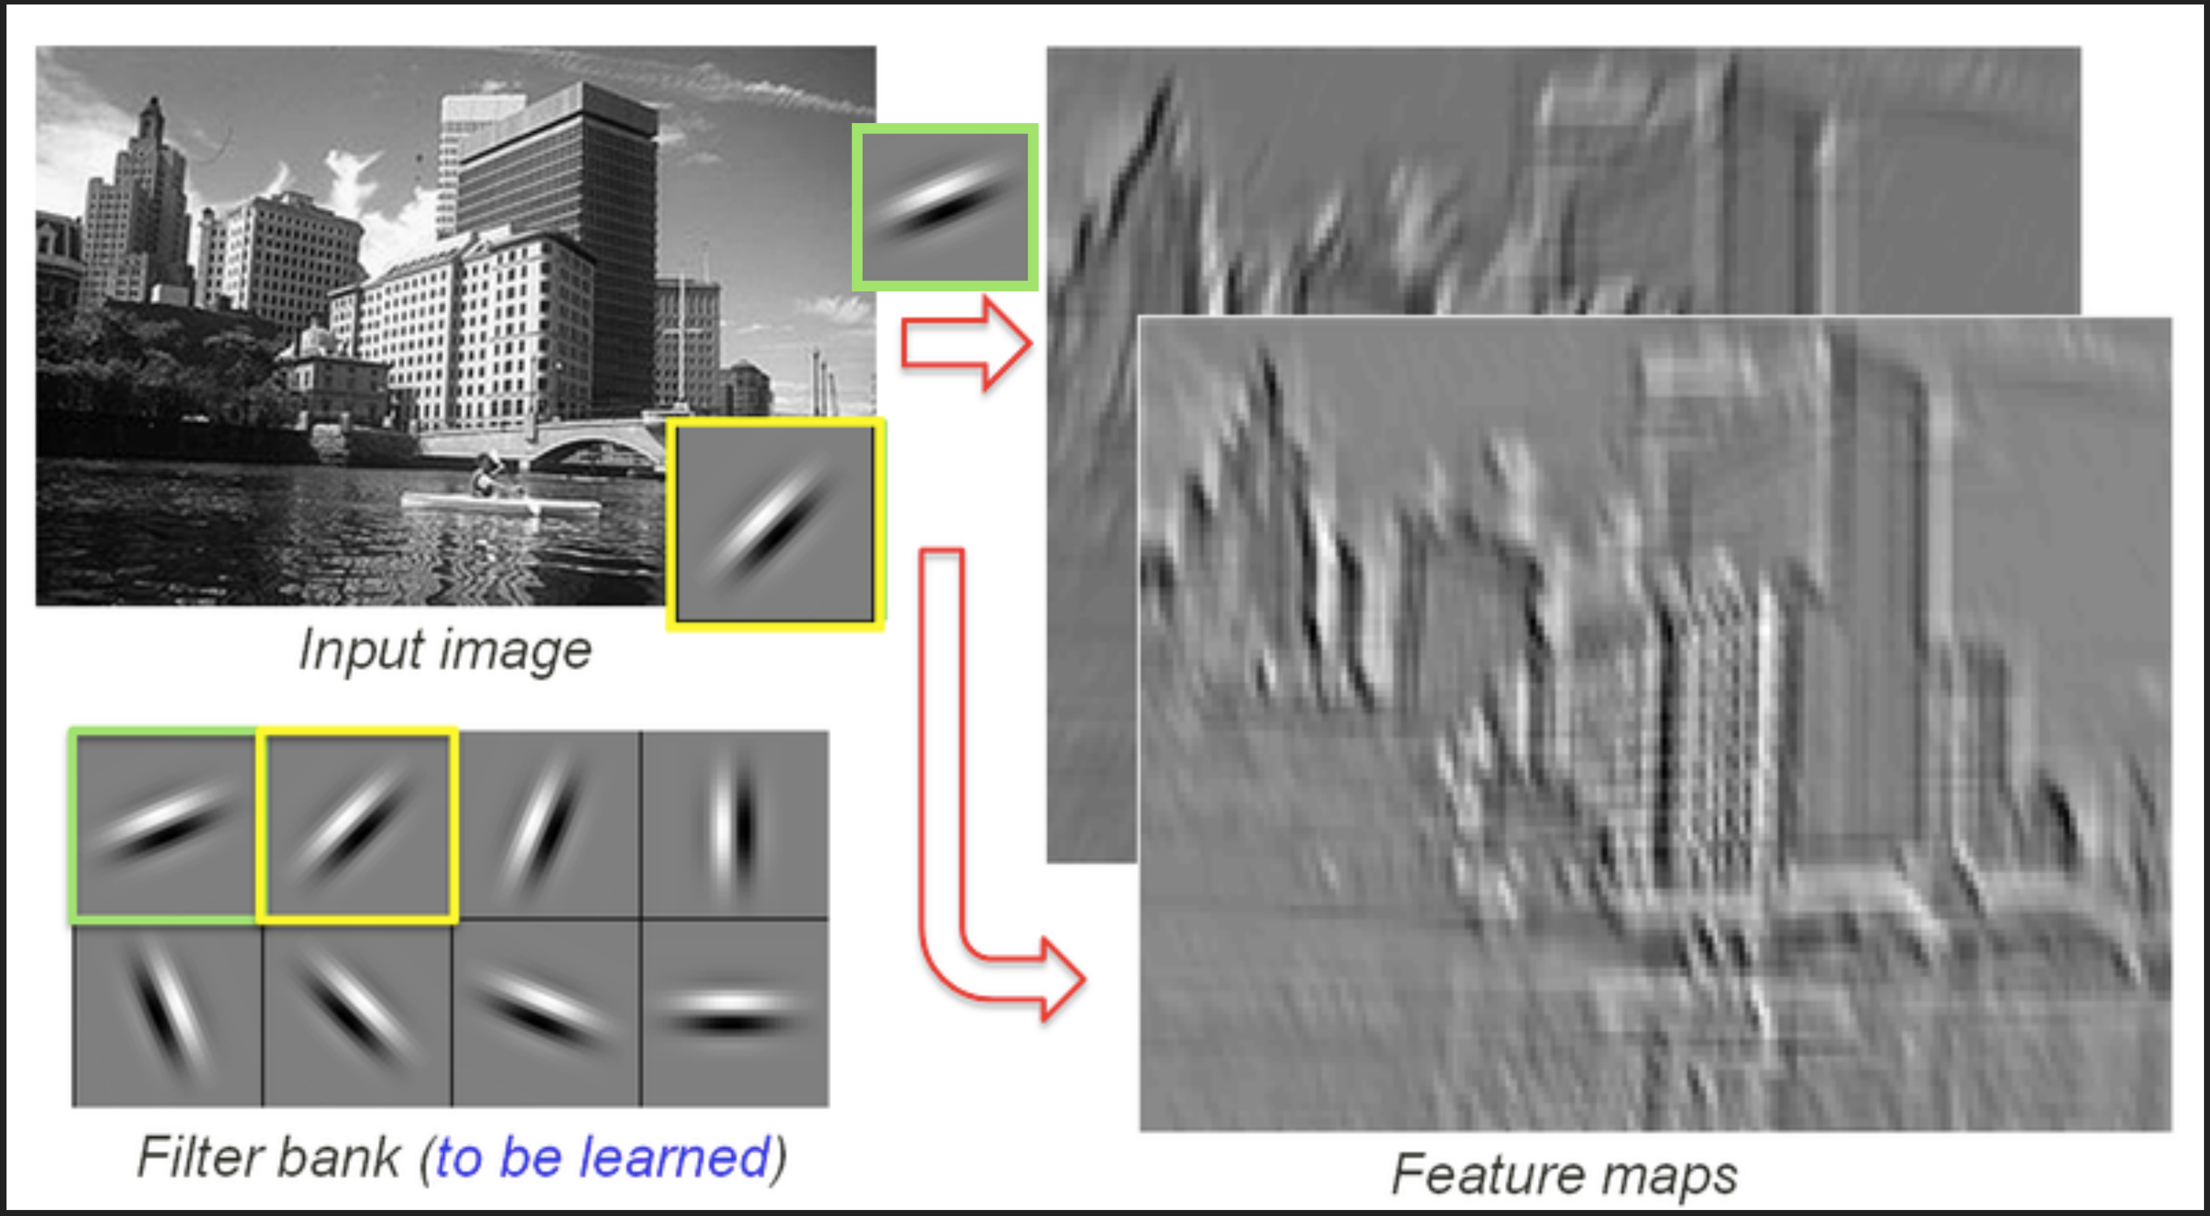

CNN NETWORKS

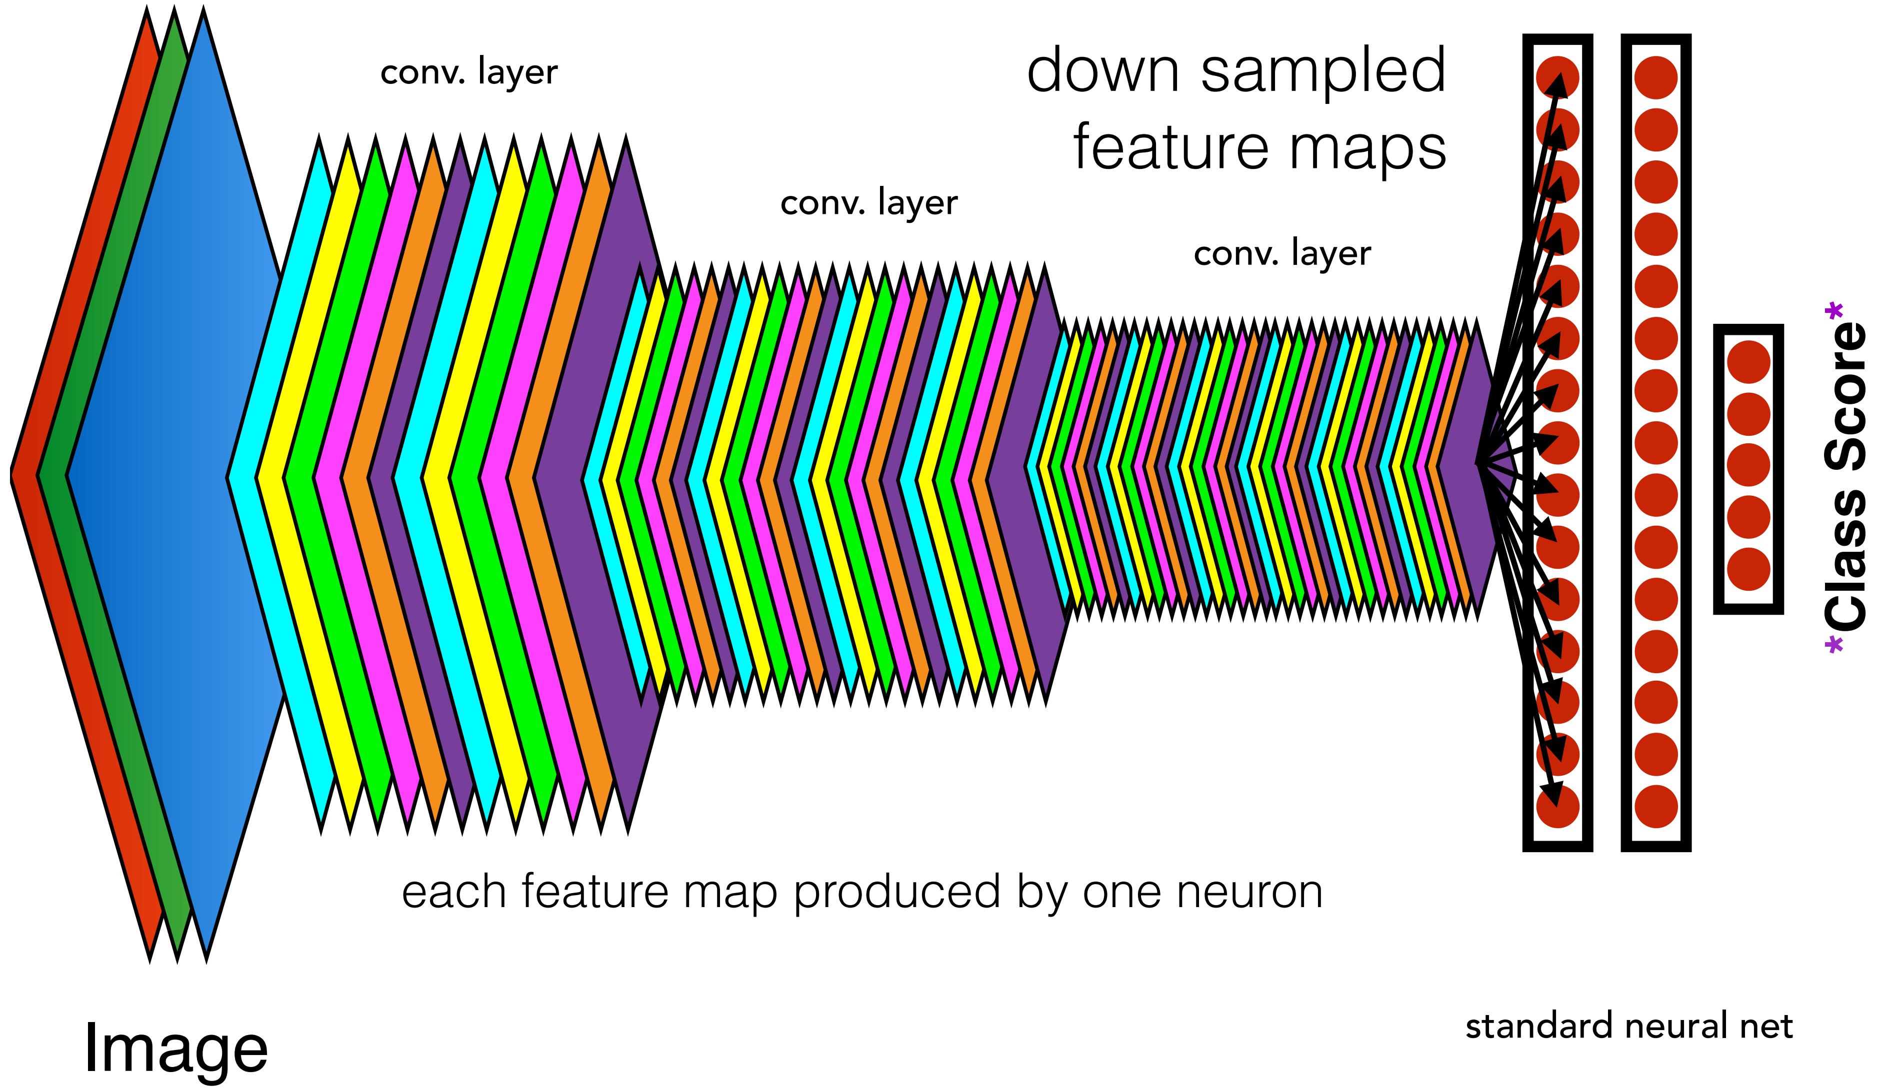

use many layers to assemble patterns into complex image features

CONVOLUTIONAL NETWORKS

- Consider the task of recognizing faces
- Begin with image pixels (layer 1)
- Start by applying convolutions of simple patterns (layer 2)
- Find groups of patterns by applying convolution on feature maps (layer 3)
- Repeat
- Eventually patterns of patterns can be identified as faces (layer 4)

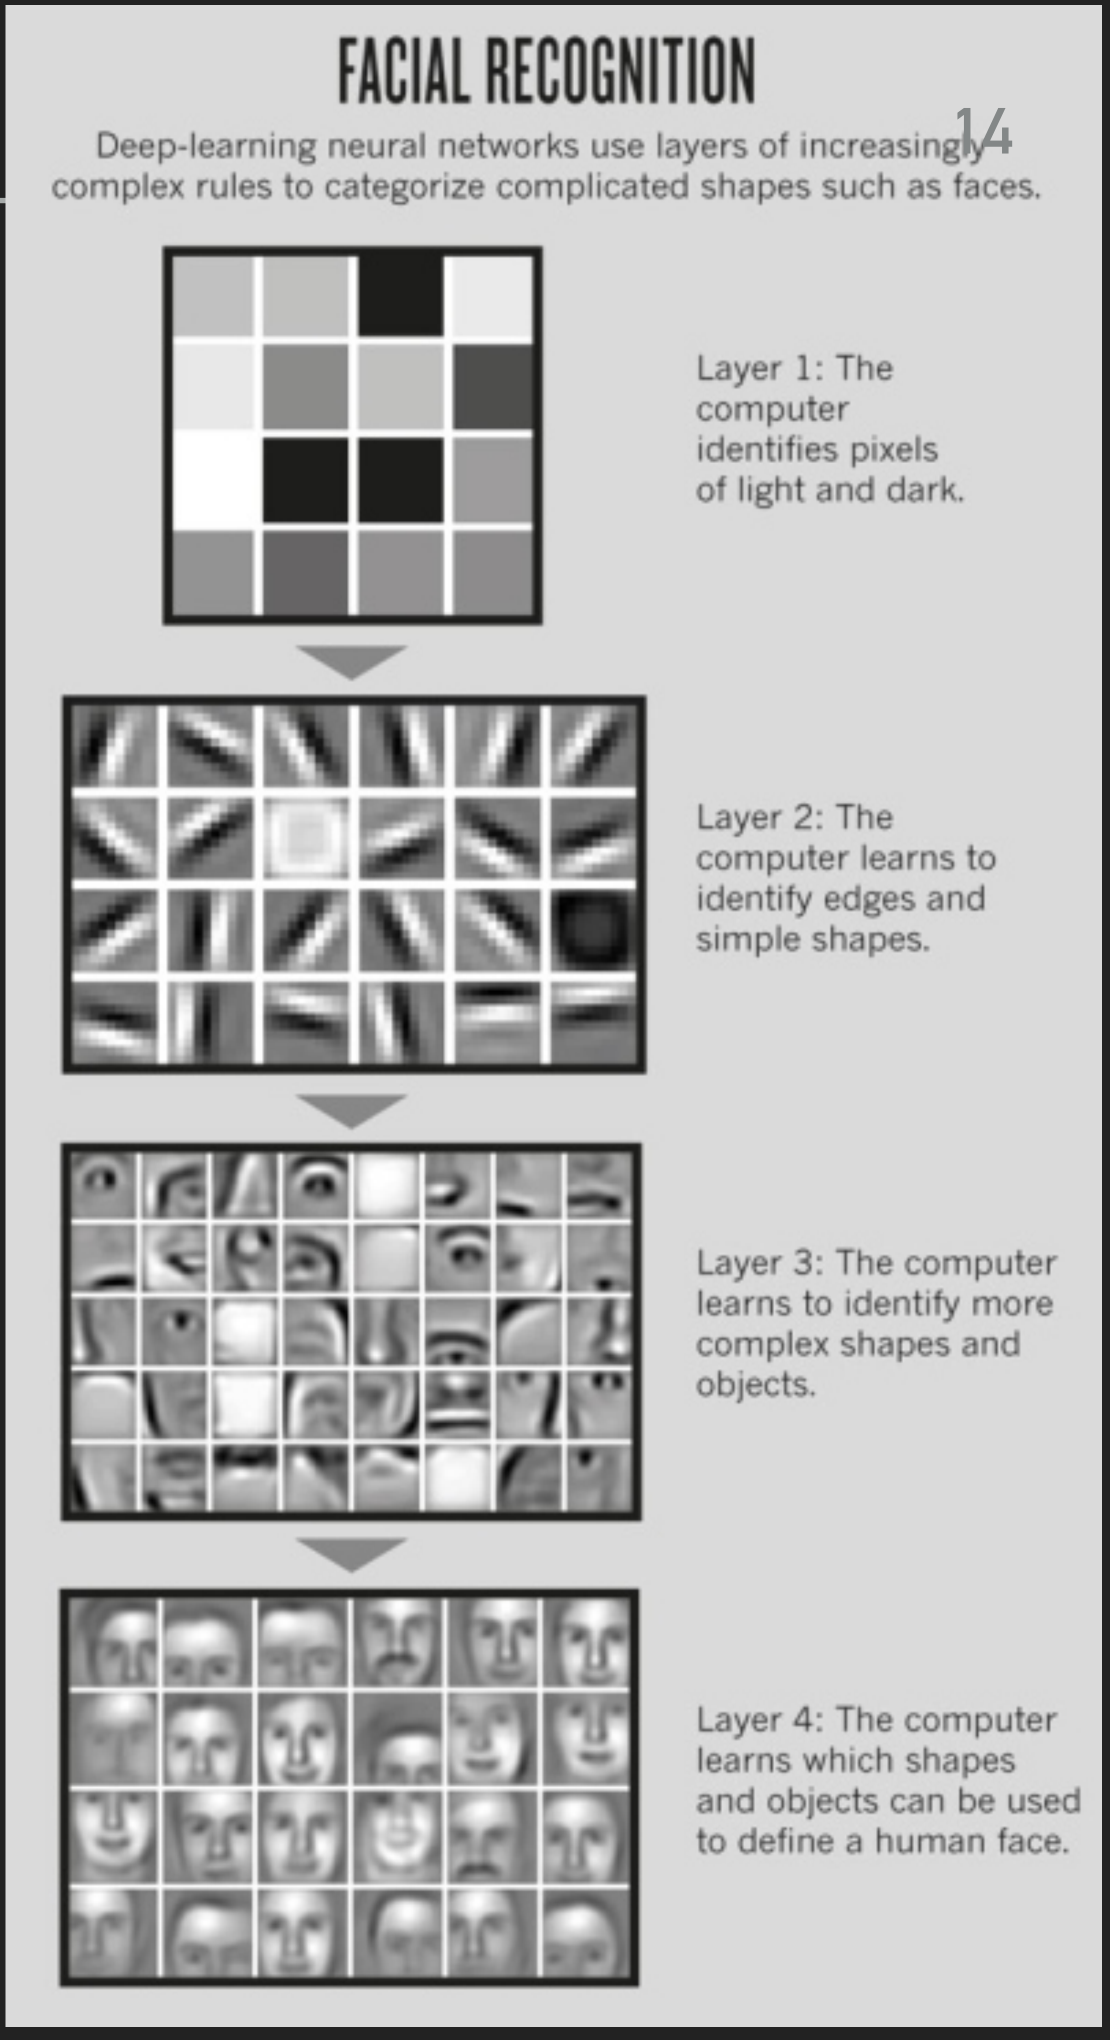

CONVOLUTIONAL NETWORKS

- CNNs learn these patterns (or convolutional filters) by themselves
- That's why CNNs are effective for many different tasks

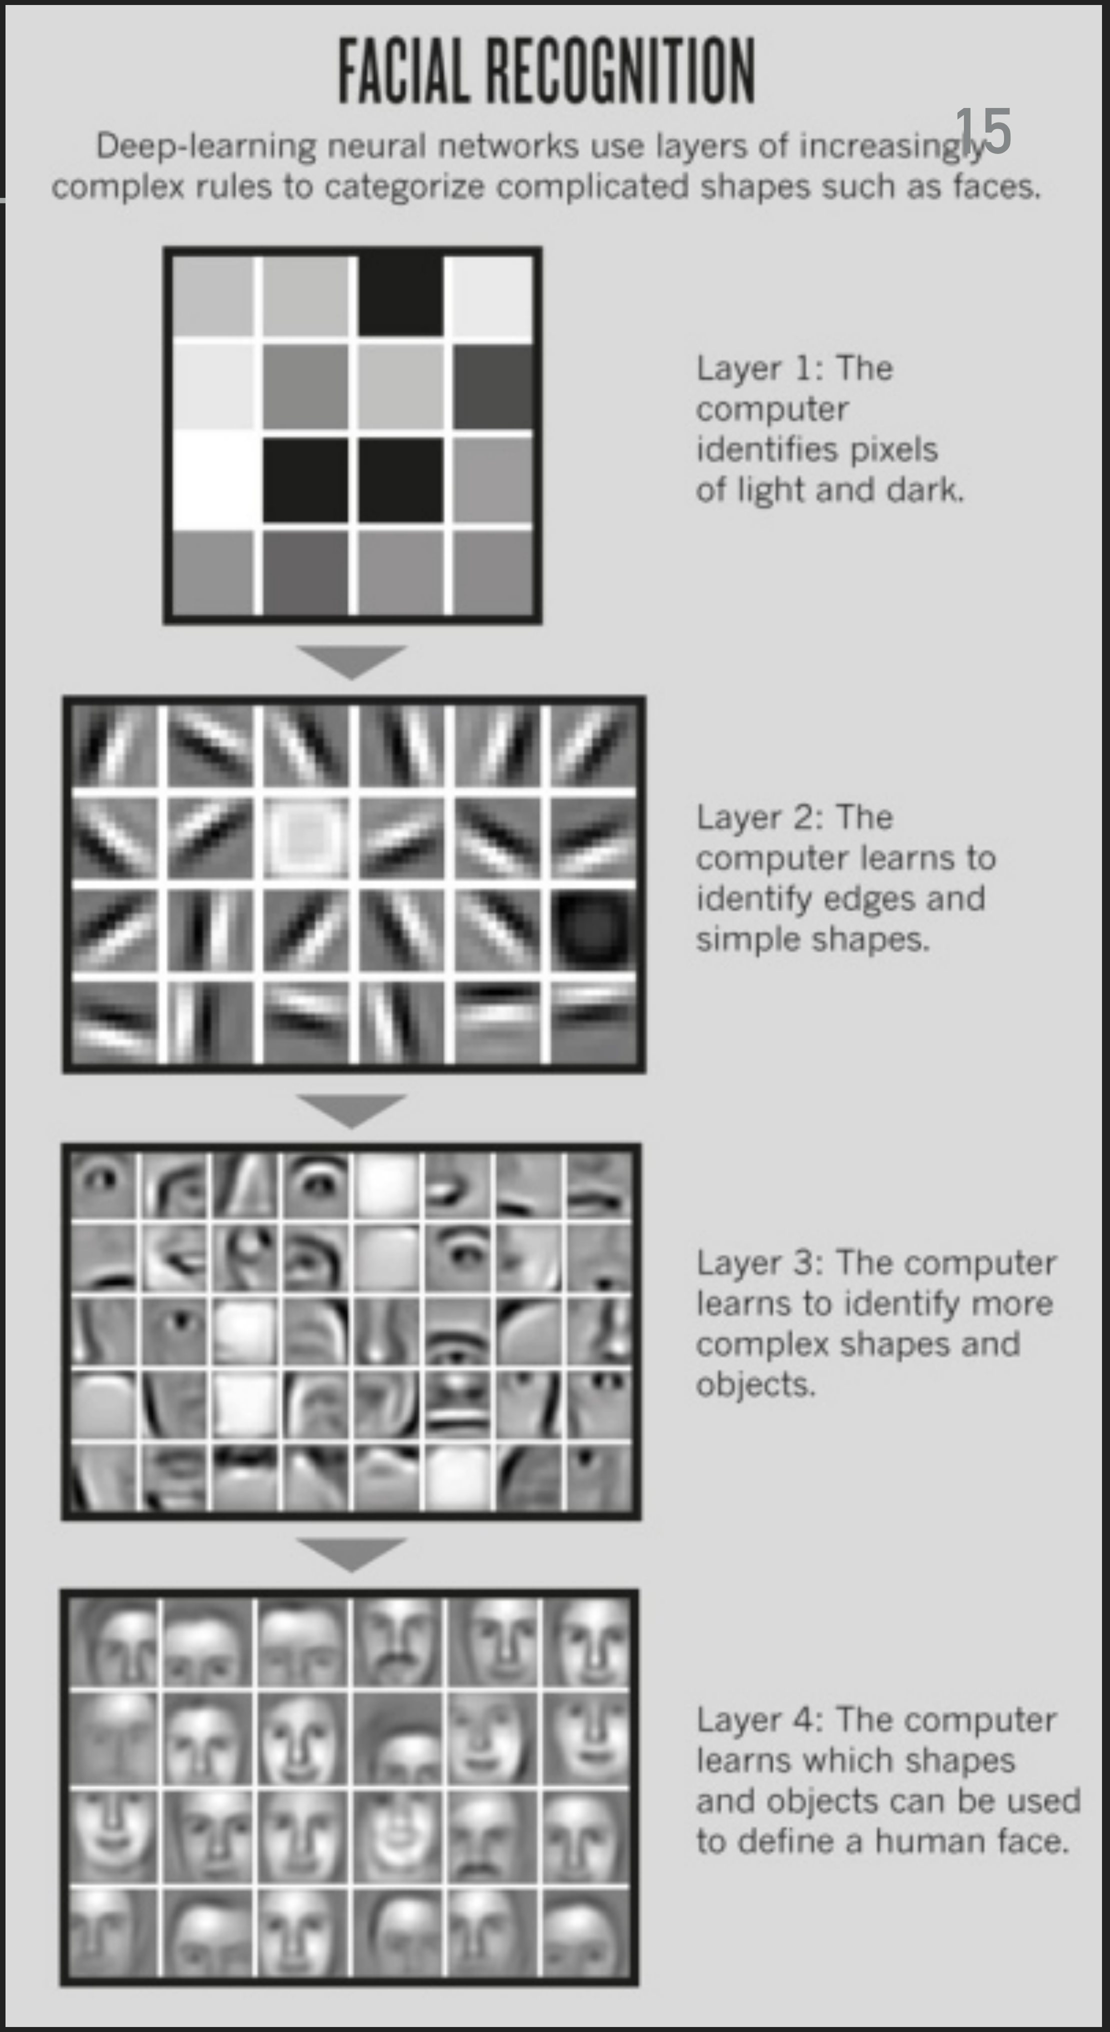

CNNS IN MICROBOONE (AND LARTPCS)

- Explored several CNN algorithms that perform tasks directly applicable to our problem
  - Image classification
  - Object detection
  - Pixel labeling

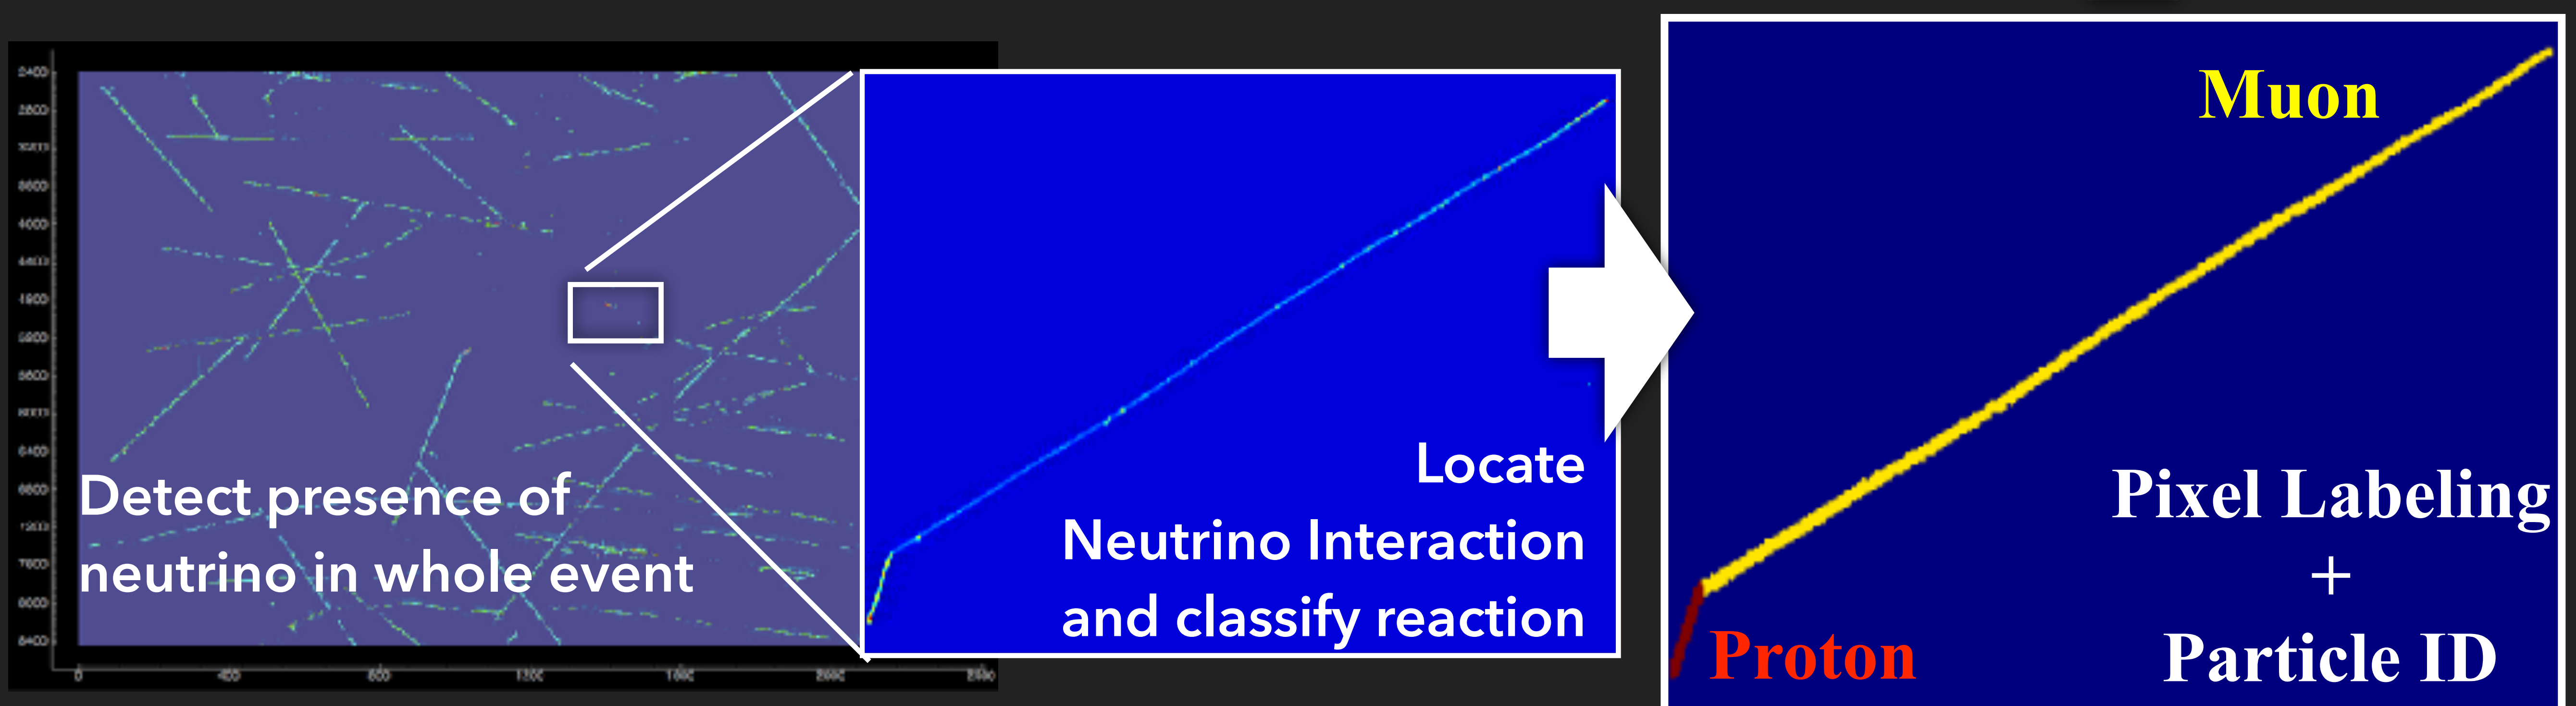

$$\nu_\mu + n \rightarrow \mu + p$$

PROOF OF PRINCIPLE STUDY

- Study with images from simulation
- To start: can network tell these four particles apart?
- Important particles in analyses

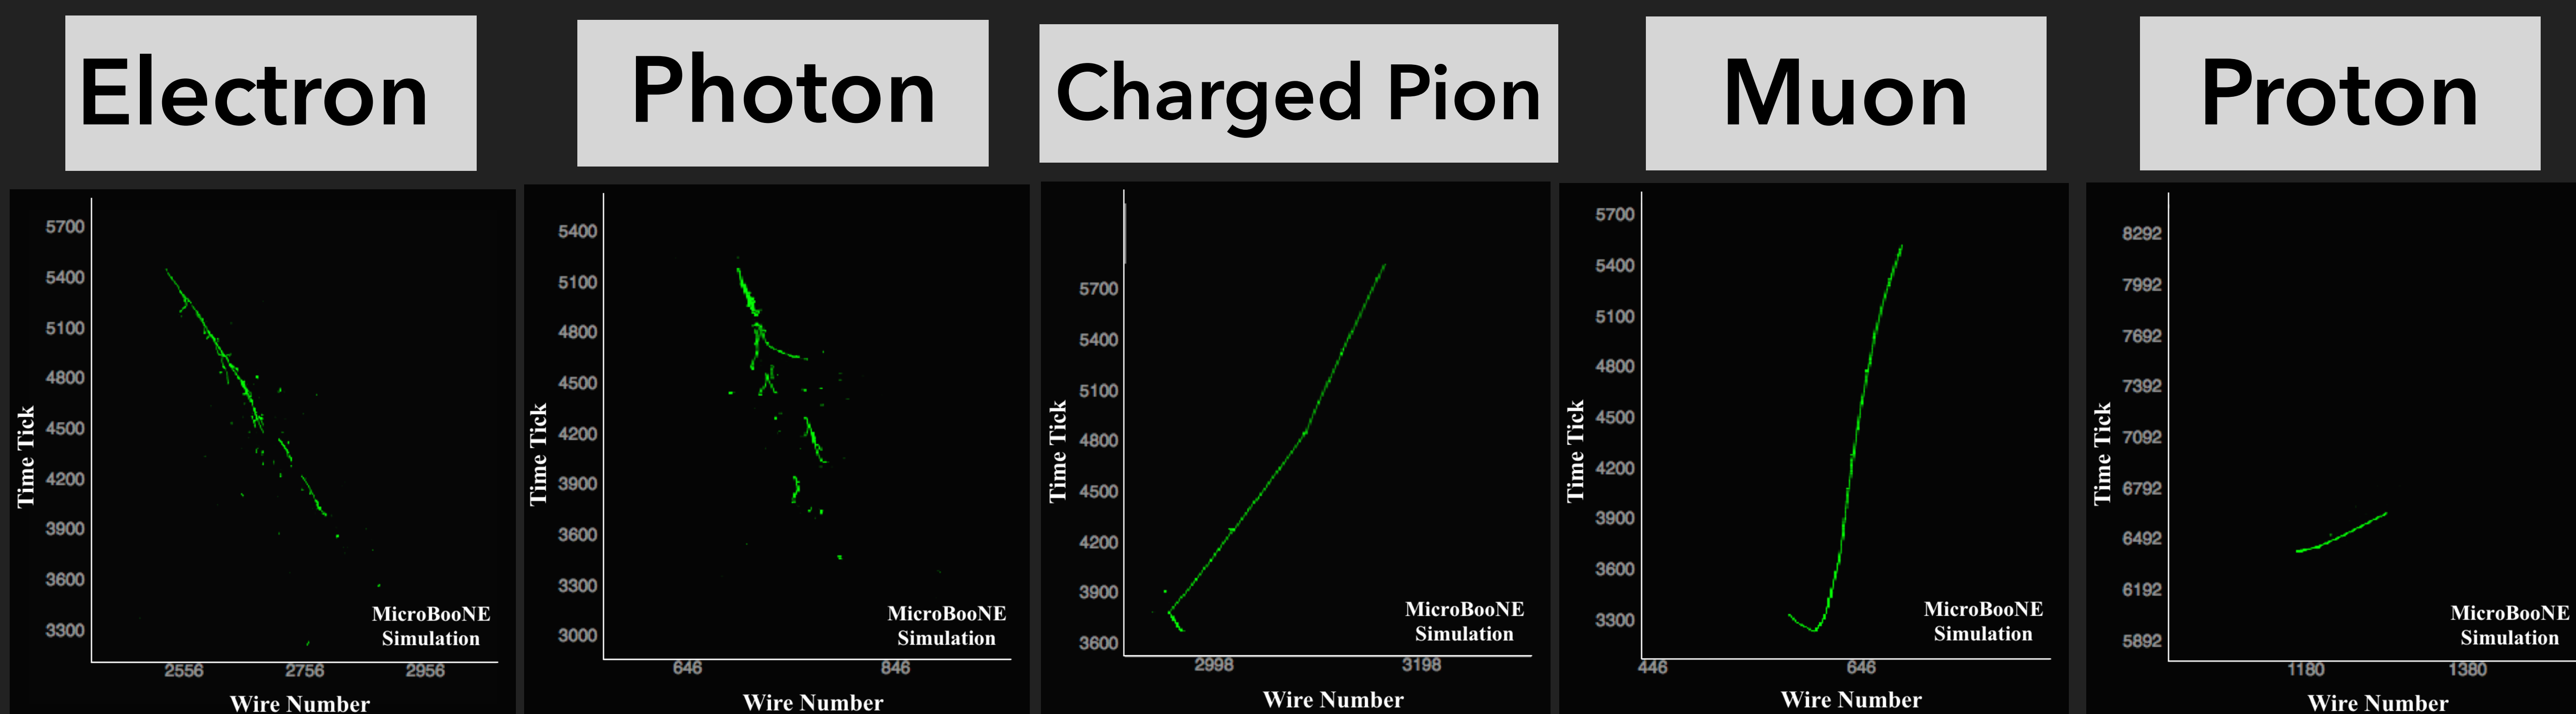

PROOF OF PRINCIPLE STUDY

- Study with images from simulation
- High-lighting electron ID: important for finding signal interactions in current/future LArTPCs $\nu_e + n \rightarrow e + p$

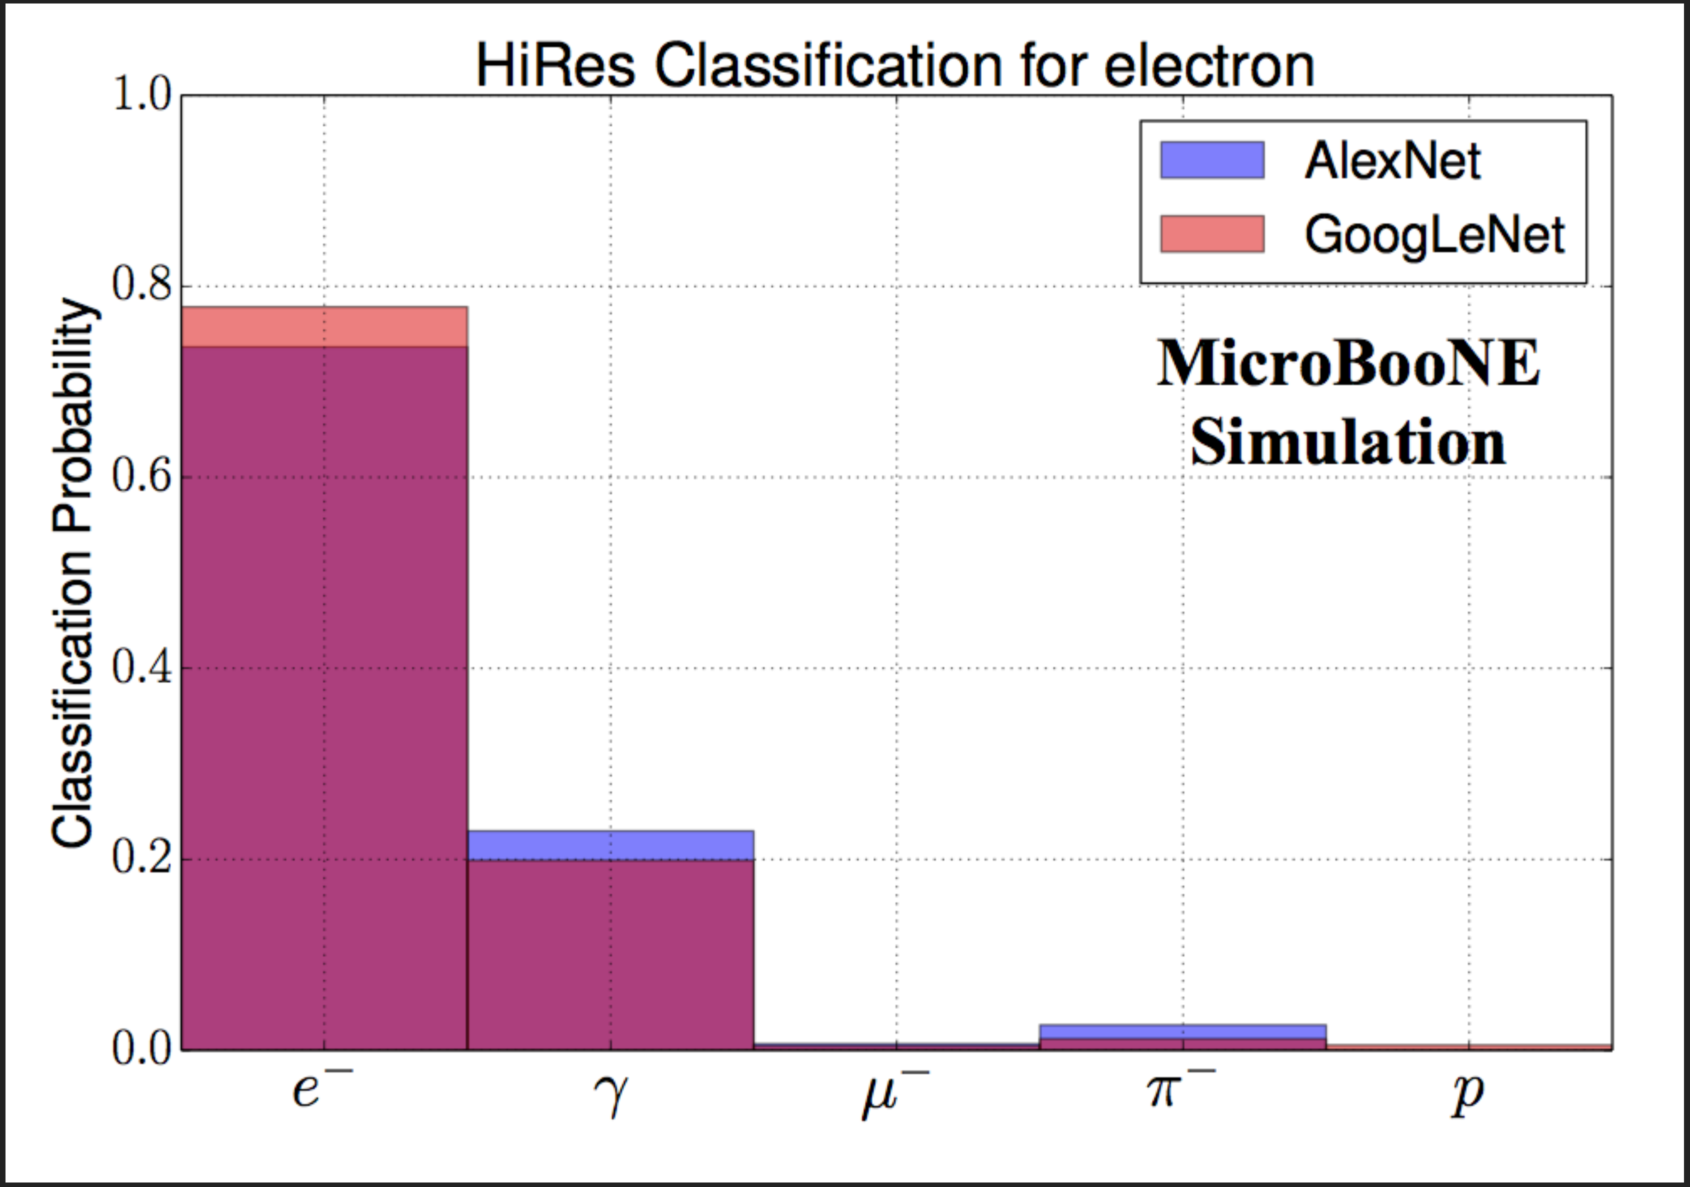

NEUTRINO INTERACTION DETECTION

- Explored class of problems known as objet detection for LArTPCs
- For surface near the detectors, could be used to locate regions of interest in the detector

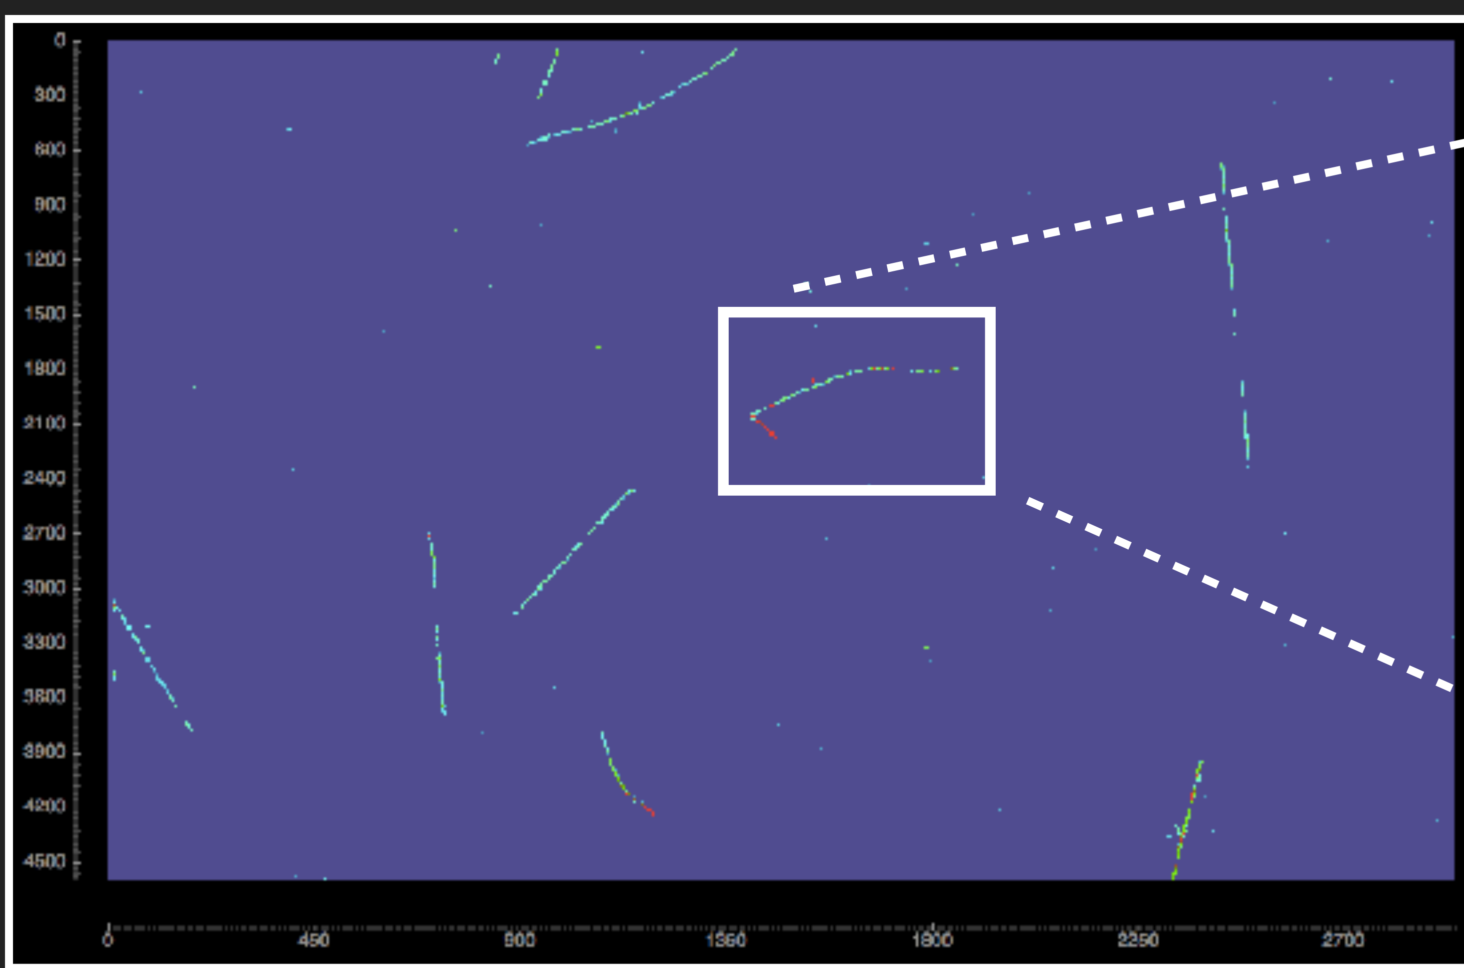

$$\nu_\mu + n \rightarrow \mu + p$$

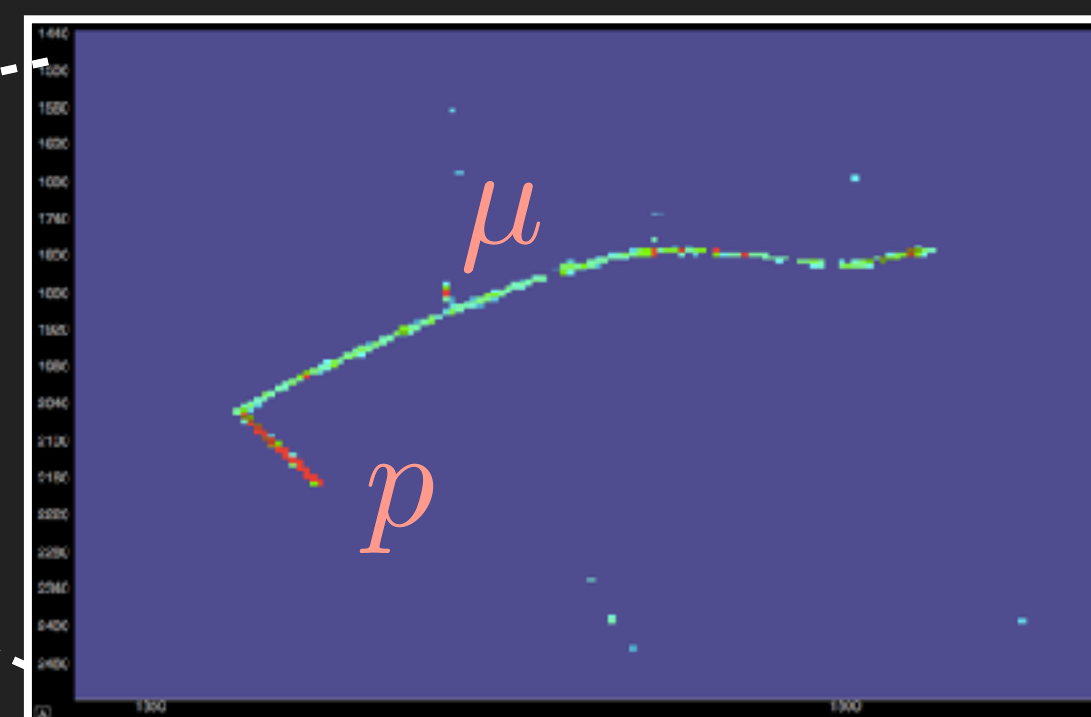

Note: had use reduce resolution image for network

RESULT: NEUTRINO DETECTION

- Key element in faster-RCNN is the Region Proposal Network
- Takes image features and determines if a given location contains an "object"
- Top regions with objects are passed to next stage, a typical classifier

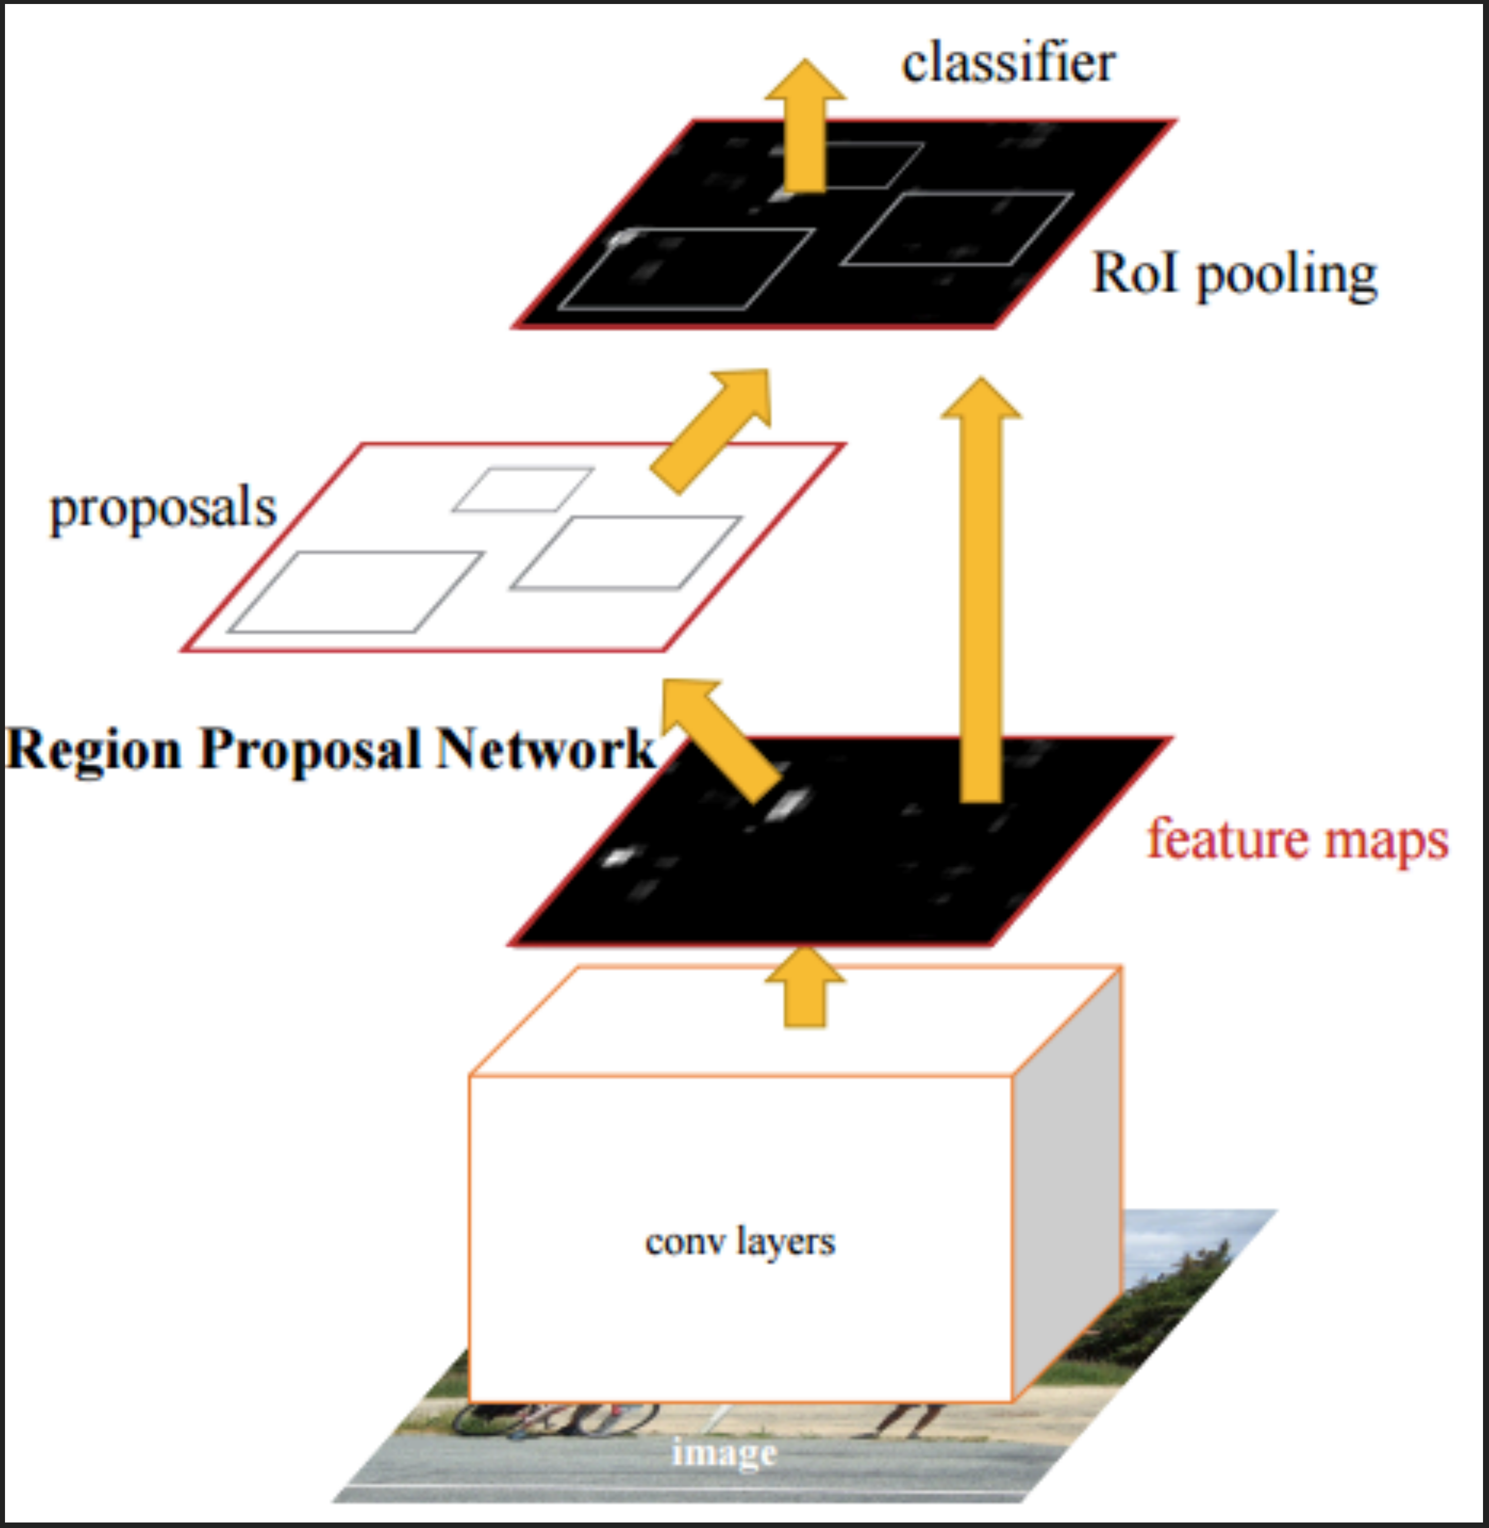

FASTER R-CNN

- Network output are classified regions of the image

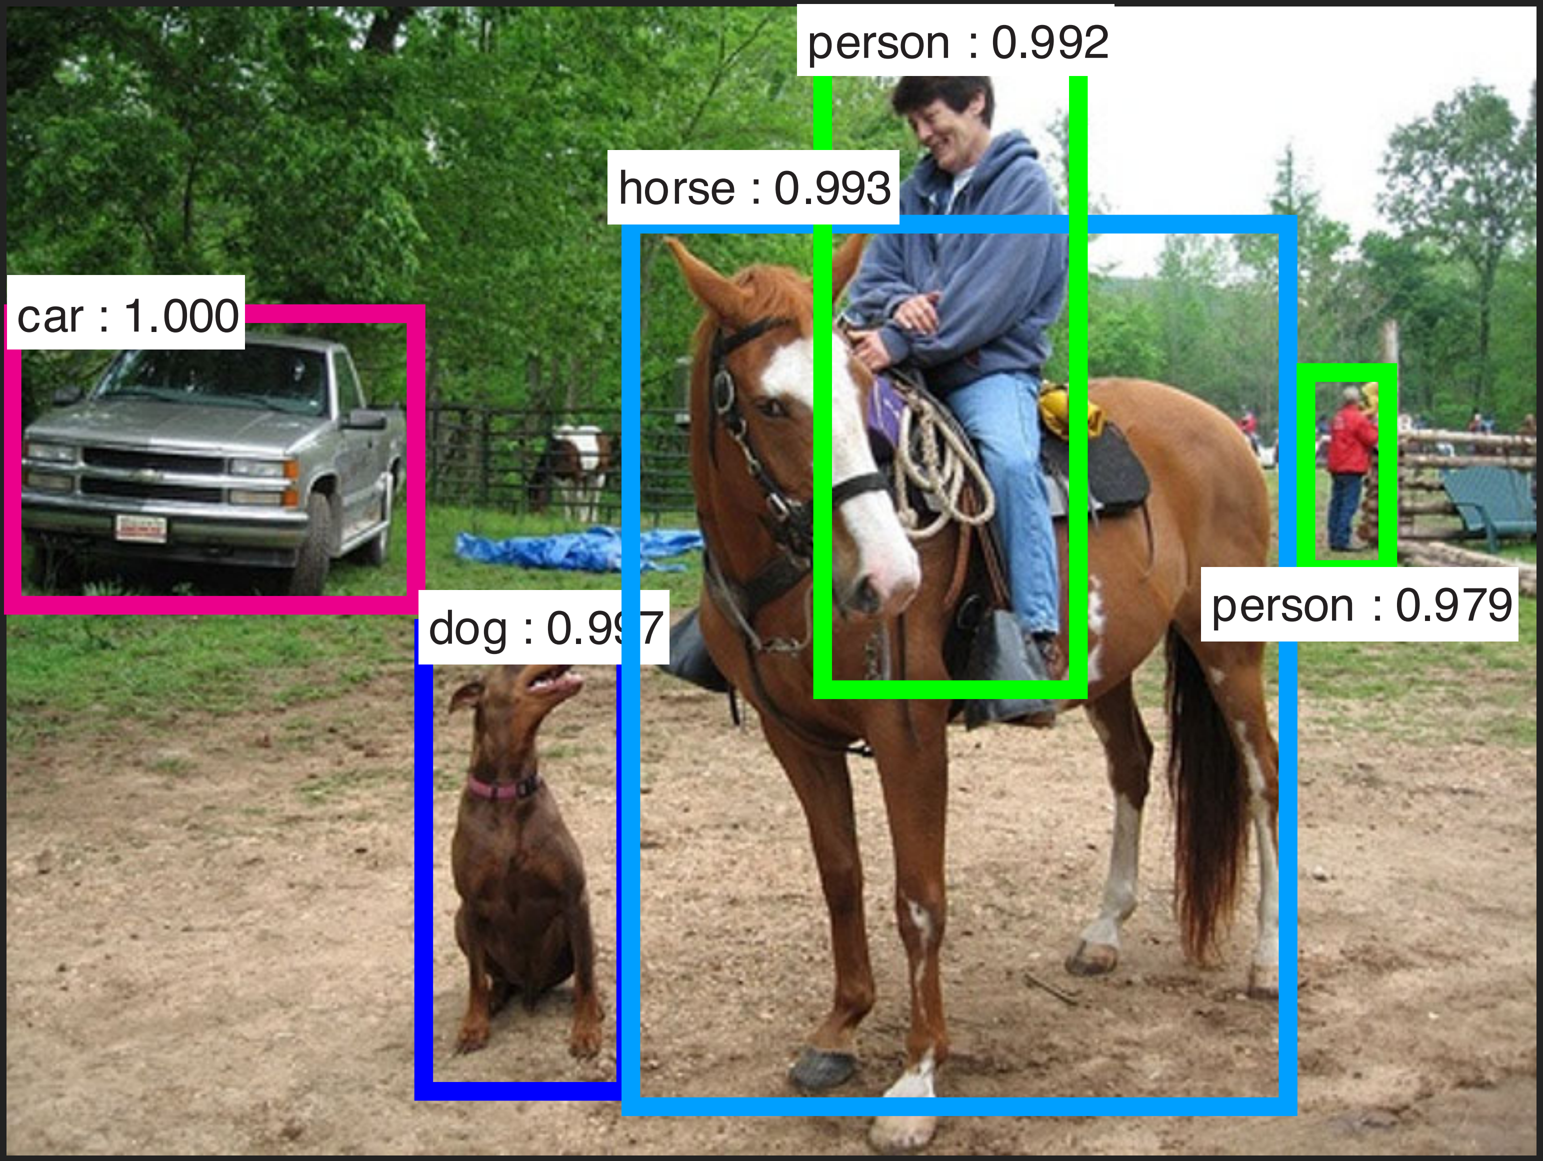

RESULT: NEUTRINO DETECTION

- Trained a network to place a bounding box around a neutrino interaction within a whole event view

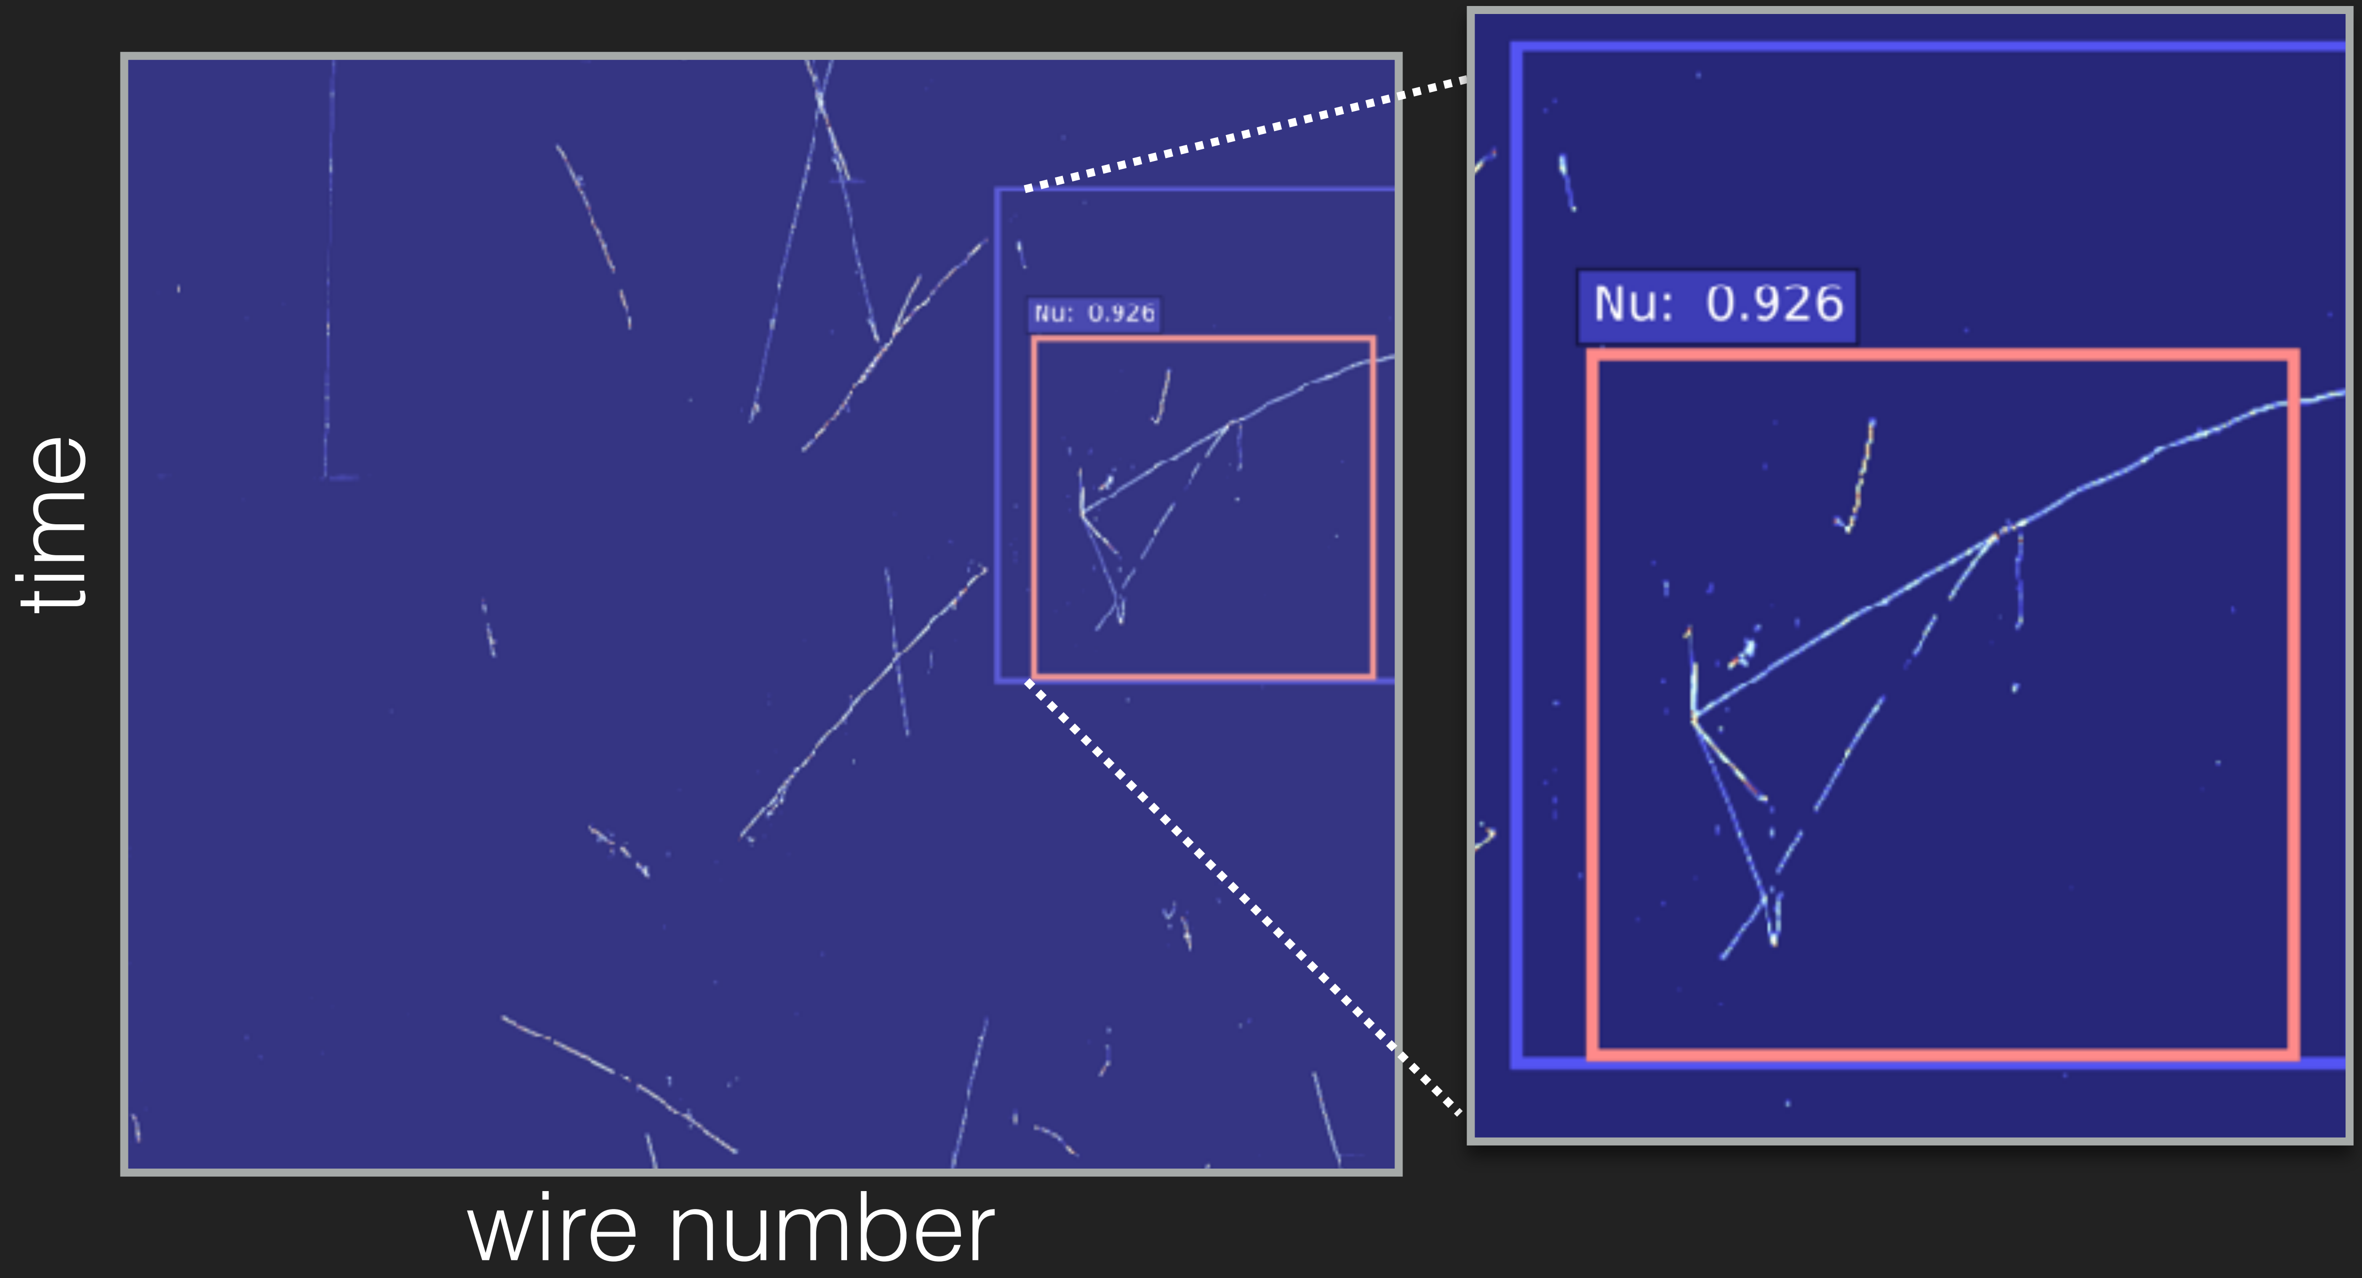

RESULT: NEUTRINO DETECTION

- Distribution of scores for regions overlapping with neutrinos (blue) versus background (red)

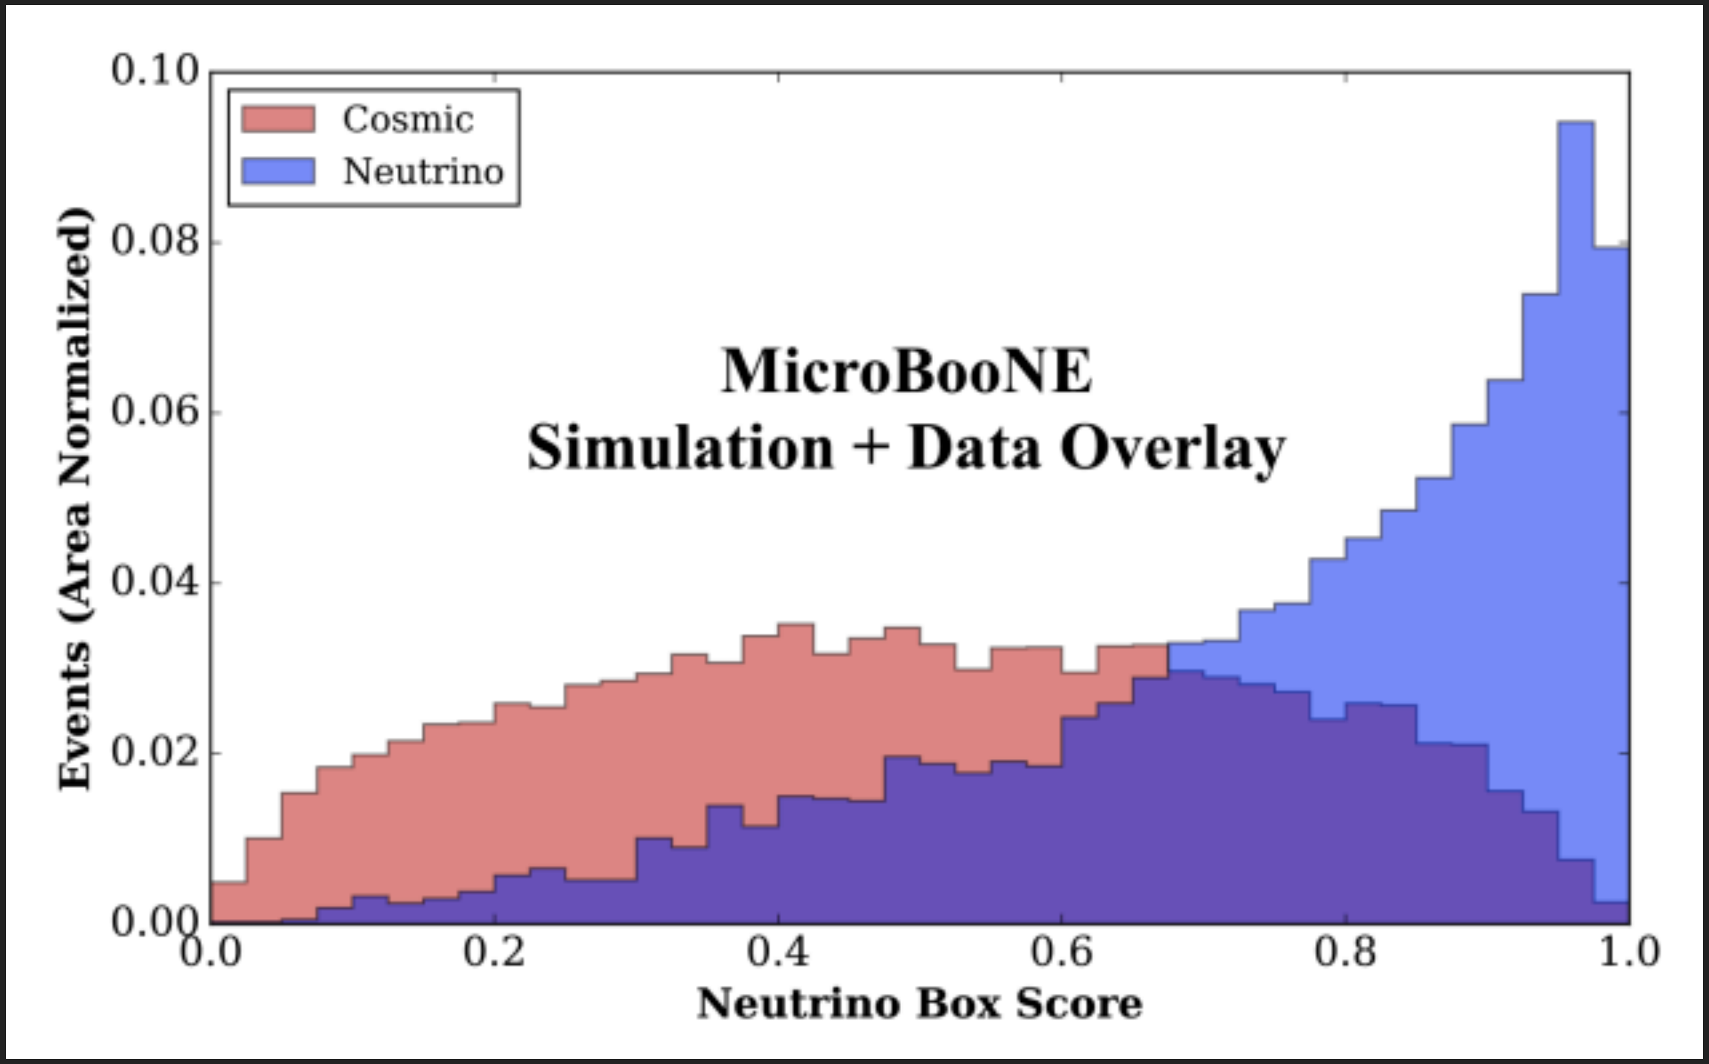

SEMANTIC SEGMENTATION

- This task asks the network to label the individual pixels as belong to some class

Image **FCN-8** **Label**

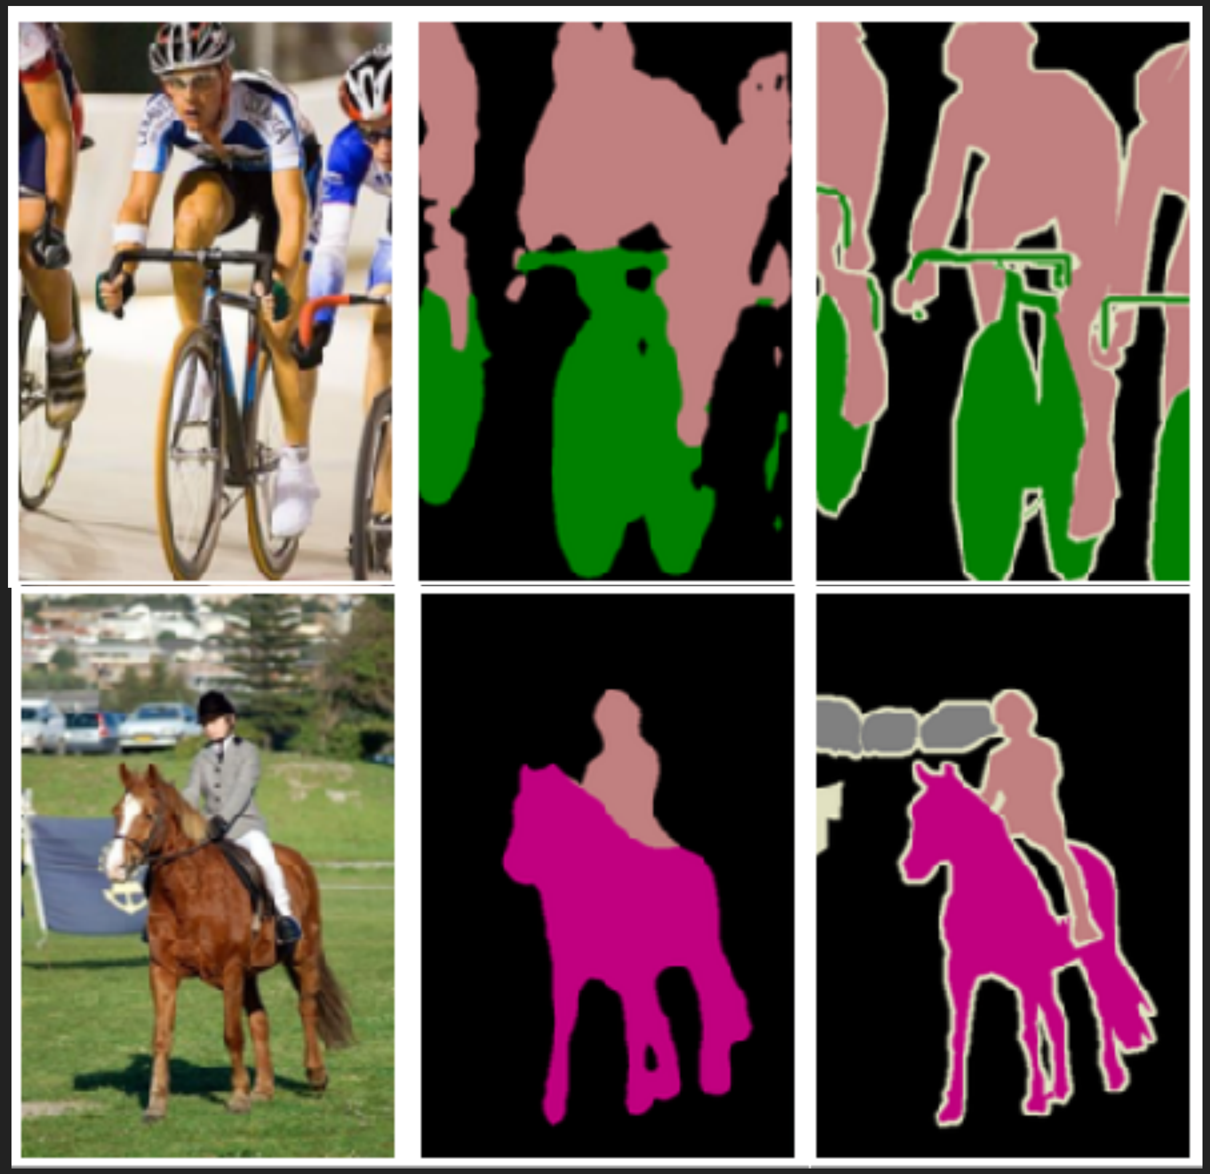

FCN-8: ***Fully-Convolutional-Network*** (FCN)

How is it different from ***Image Classification***?

Cartoon of Image Classification

Encode

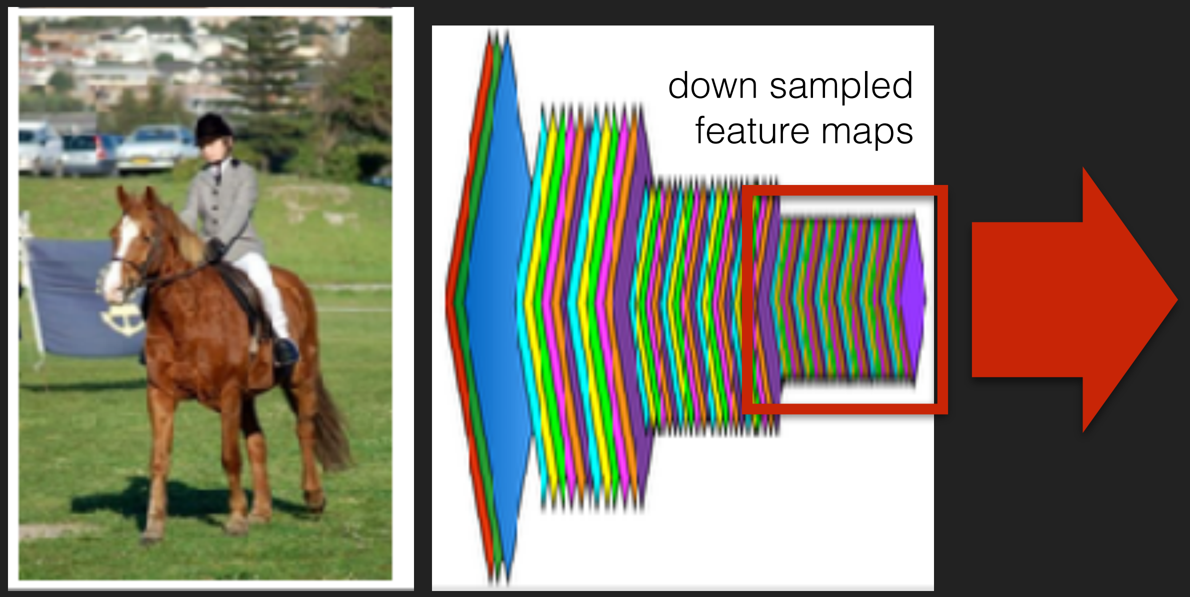

input image

class vector

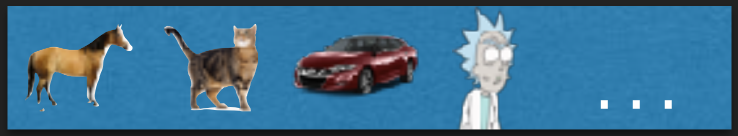

- Convolution layers find collection of complex features
- Features found combined to determine most likely objects in whole images

How is it different from *Image Classification*?

Cartoon of Image Classification

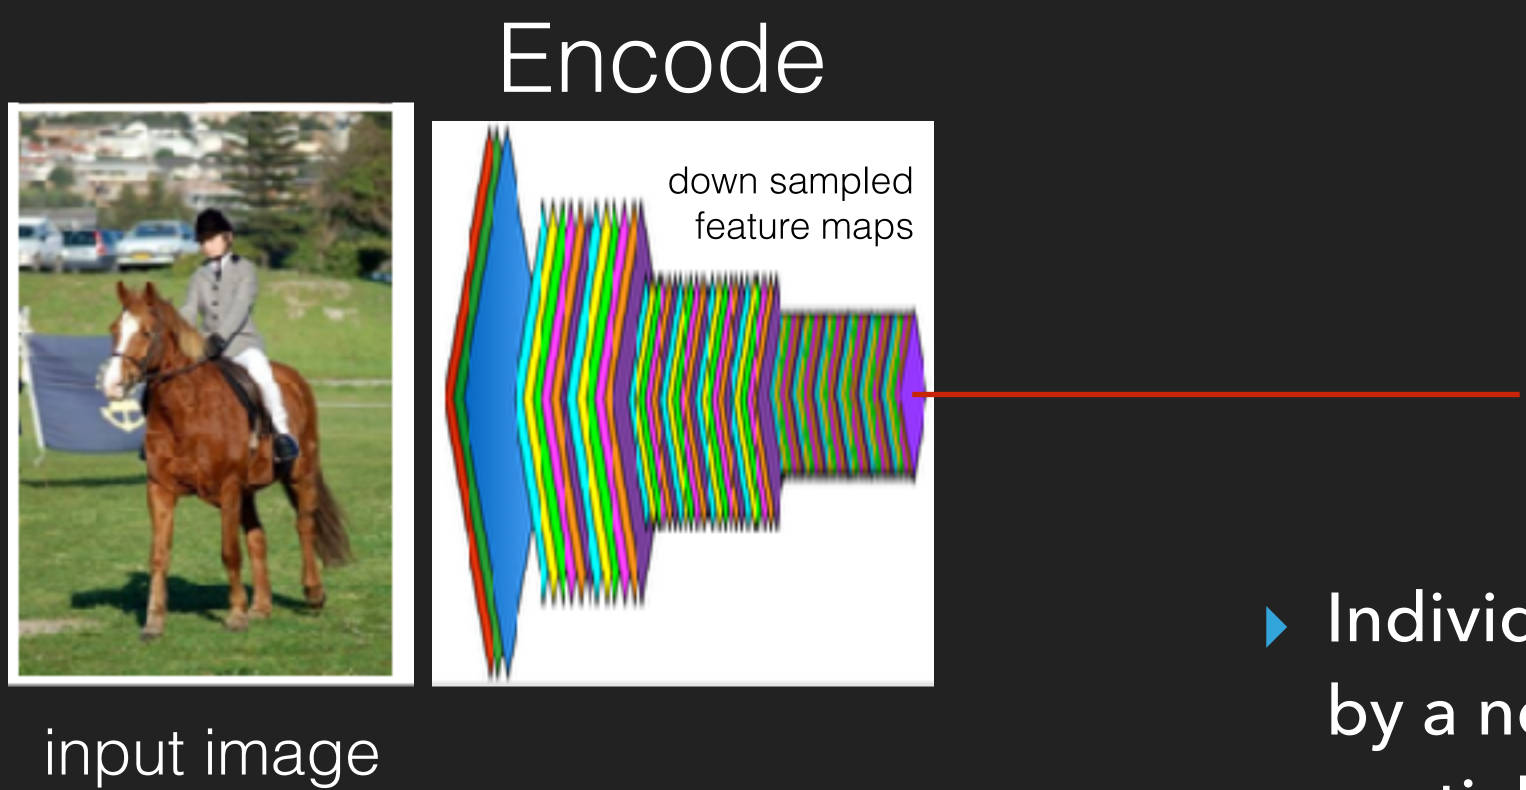

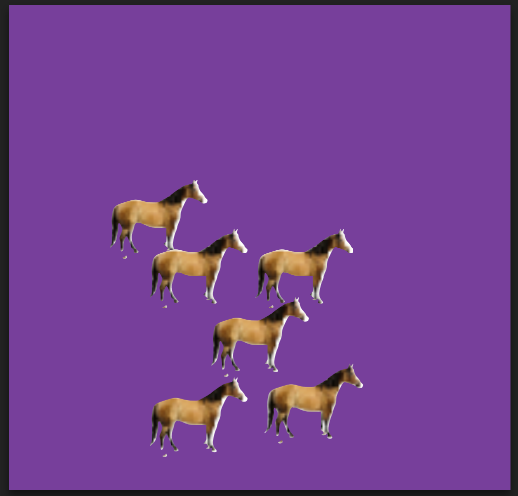

- Individual feature maps (produced by a neuron in a layer) contain spatial information
- However, down-sampled
- For semantic segmentation, we want to use this information

How is it different from ***Image Classification***?

<u>Cartoon of Fully-Convolutional SS Network</u>

Encode

Decode

input image

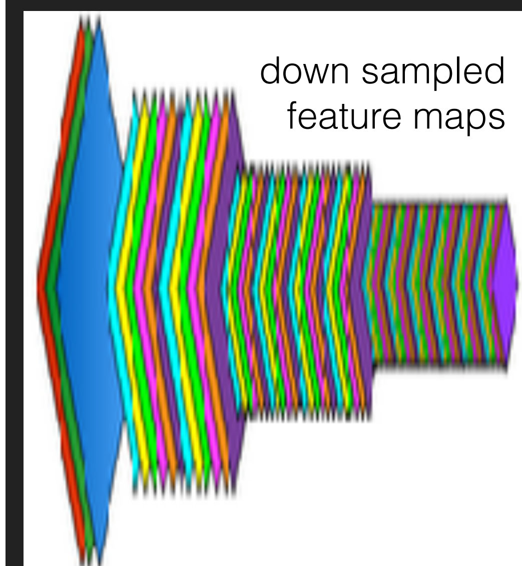

feature
vector

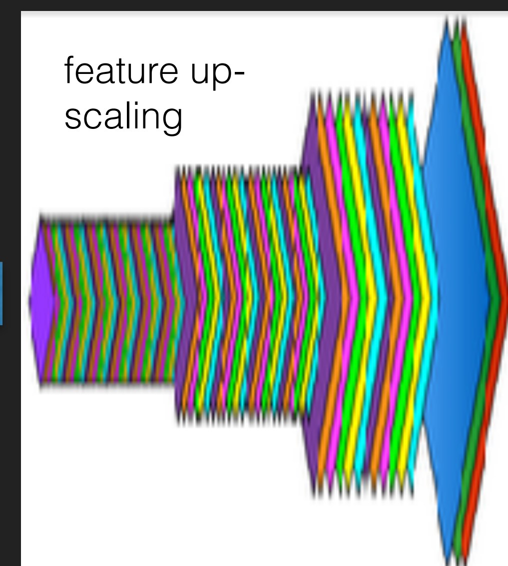

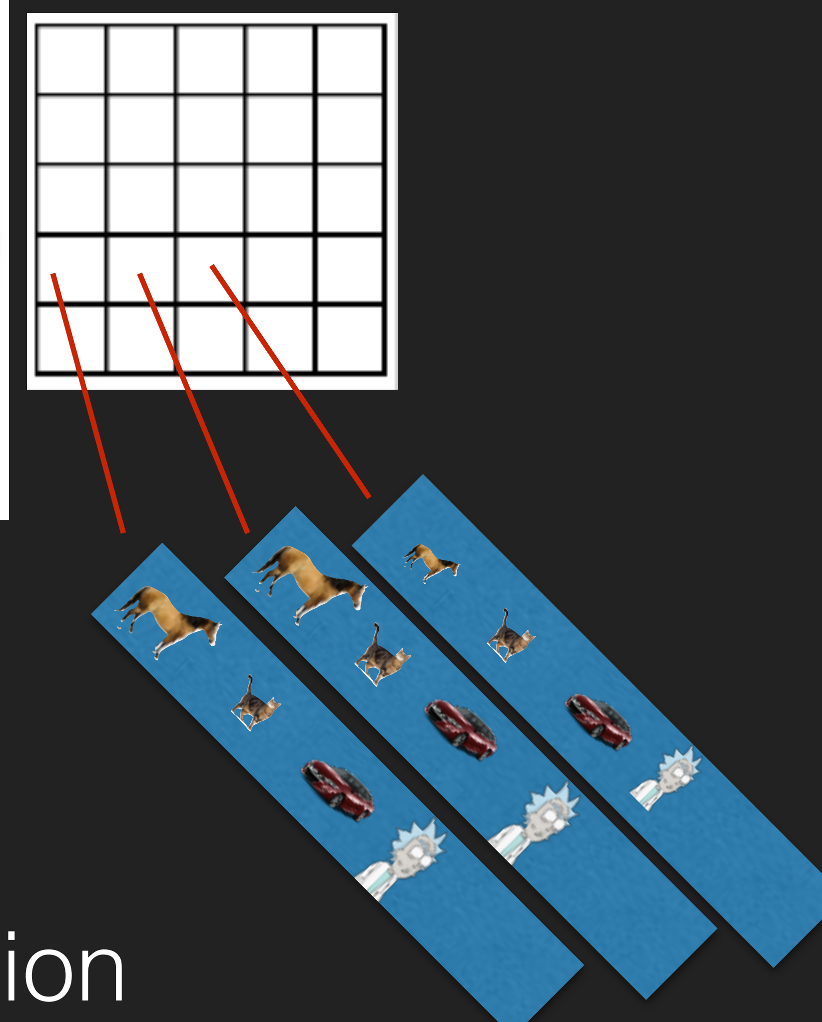

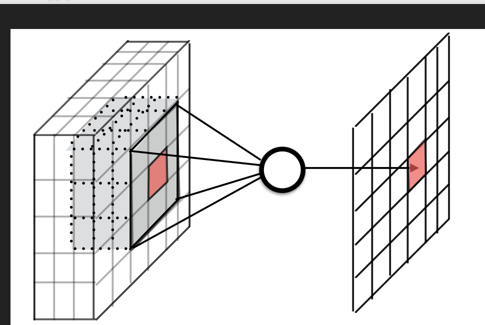

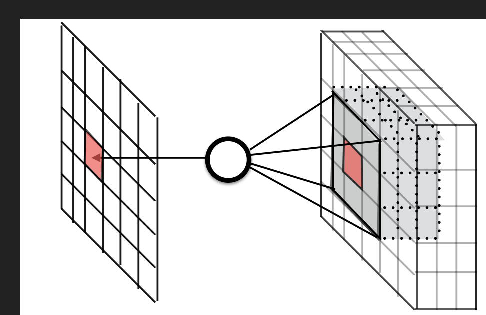

convolutions

learned projection

How is it different from ***Image Classification***?

Cartoon of Fully-Convolutional SS Network

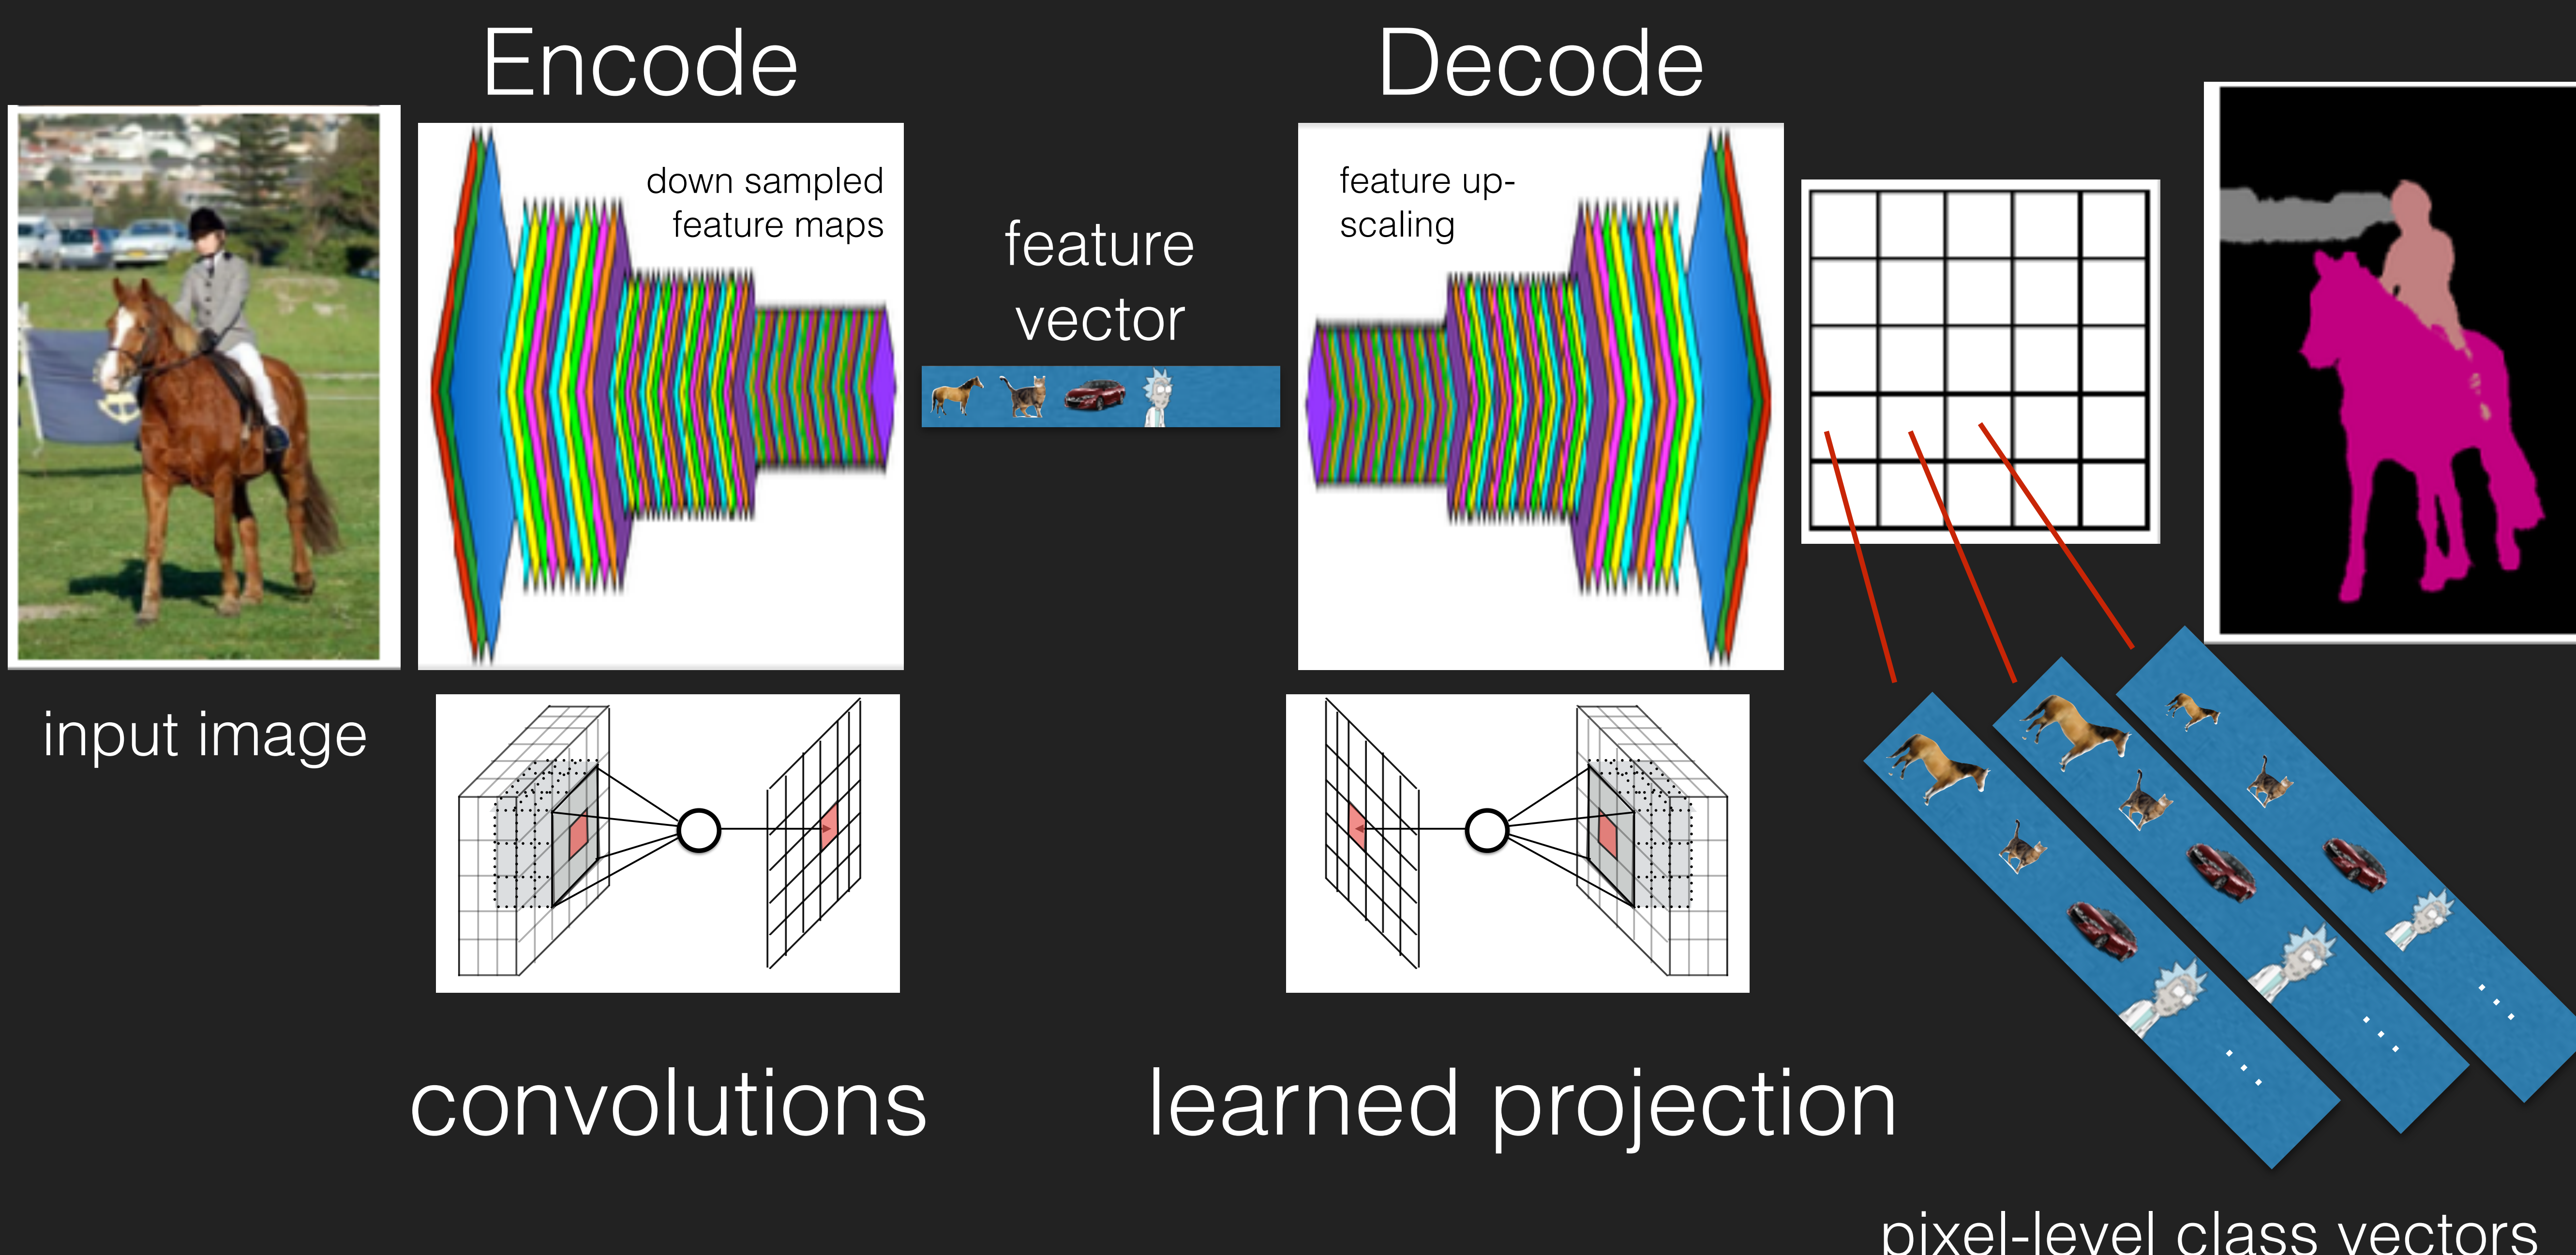

SEMANTIC SEGMENTATION IN LARTPC

Supervised Training (UB)

- Assign pixel-wise "weight" to penalize mistakes
- Weights inversely proportional to each "category" of pixel count
- Useful for LArTPC images (low information density)
- U-Net (arXiv:1505.04597)

"Weight" Image
(for training)

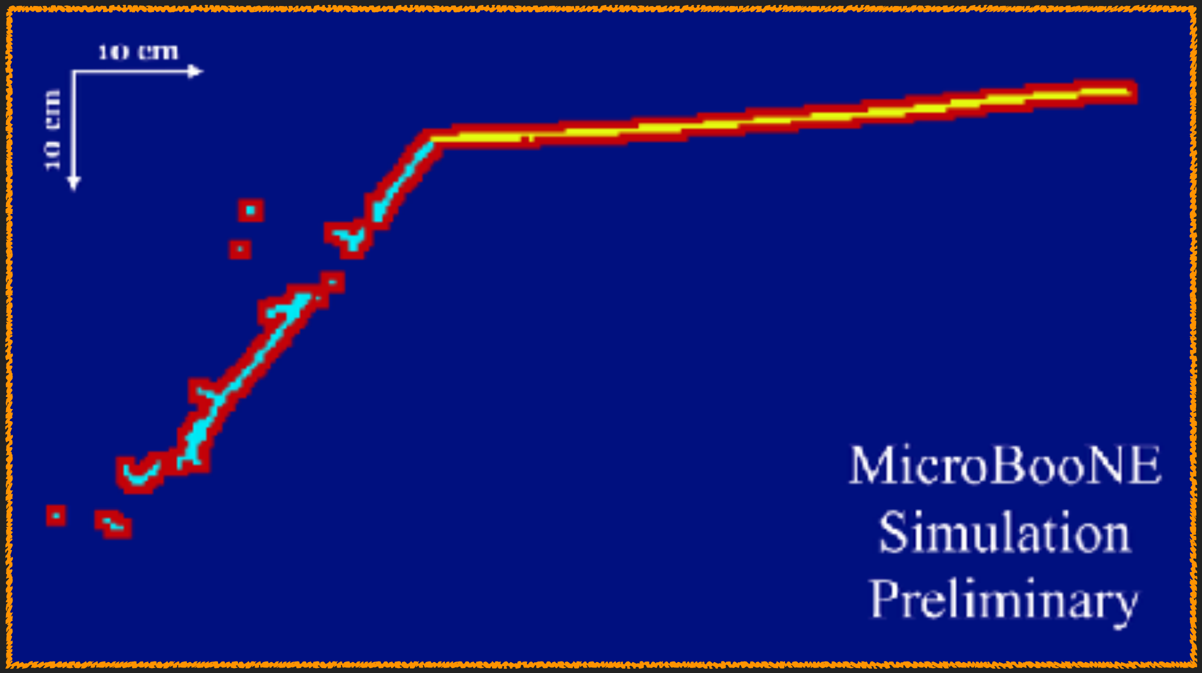

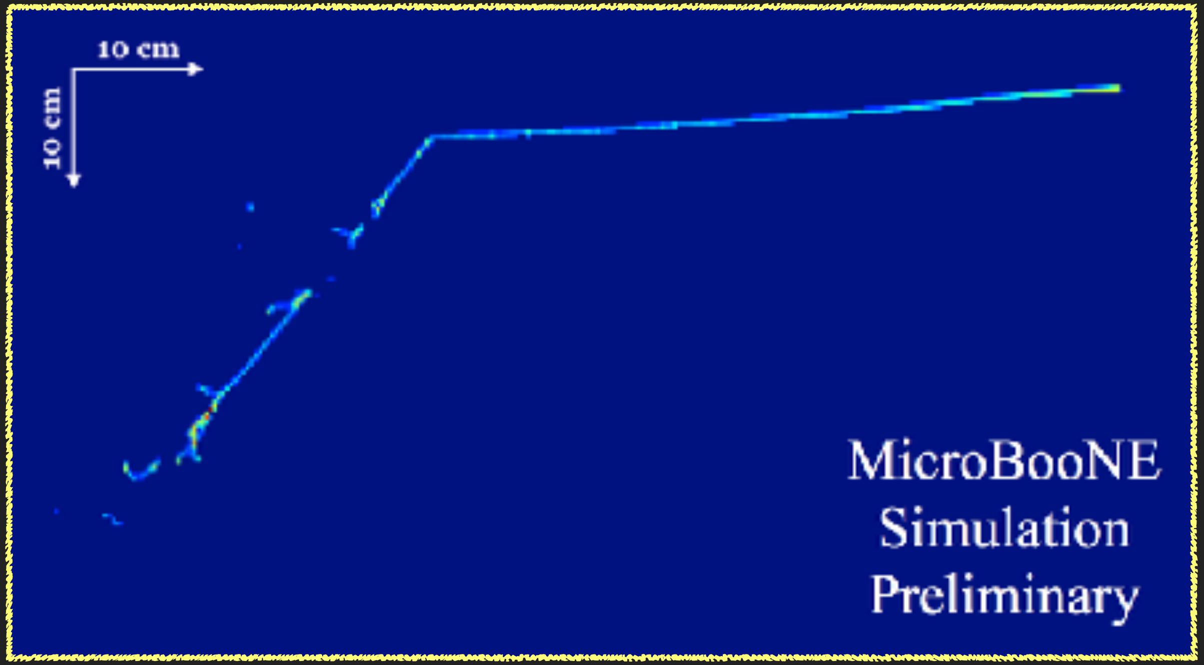

Input Image

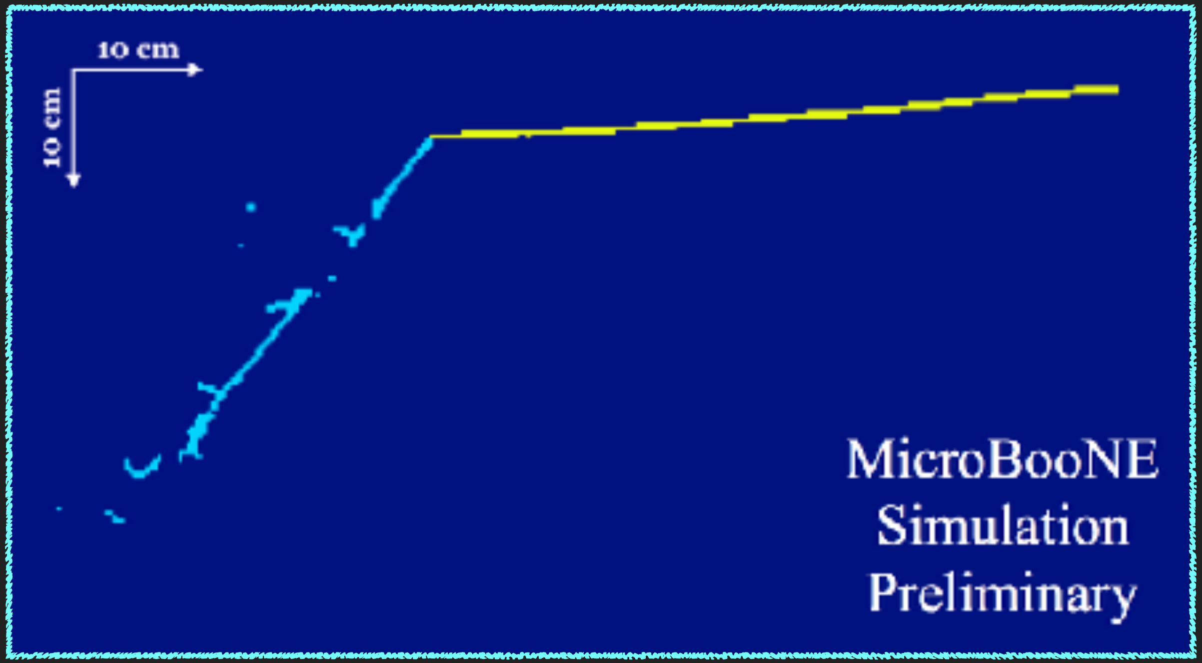

"Label" Image
(for training)

SEMANTIC SEGMENTATION

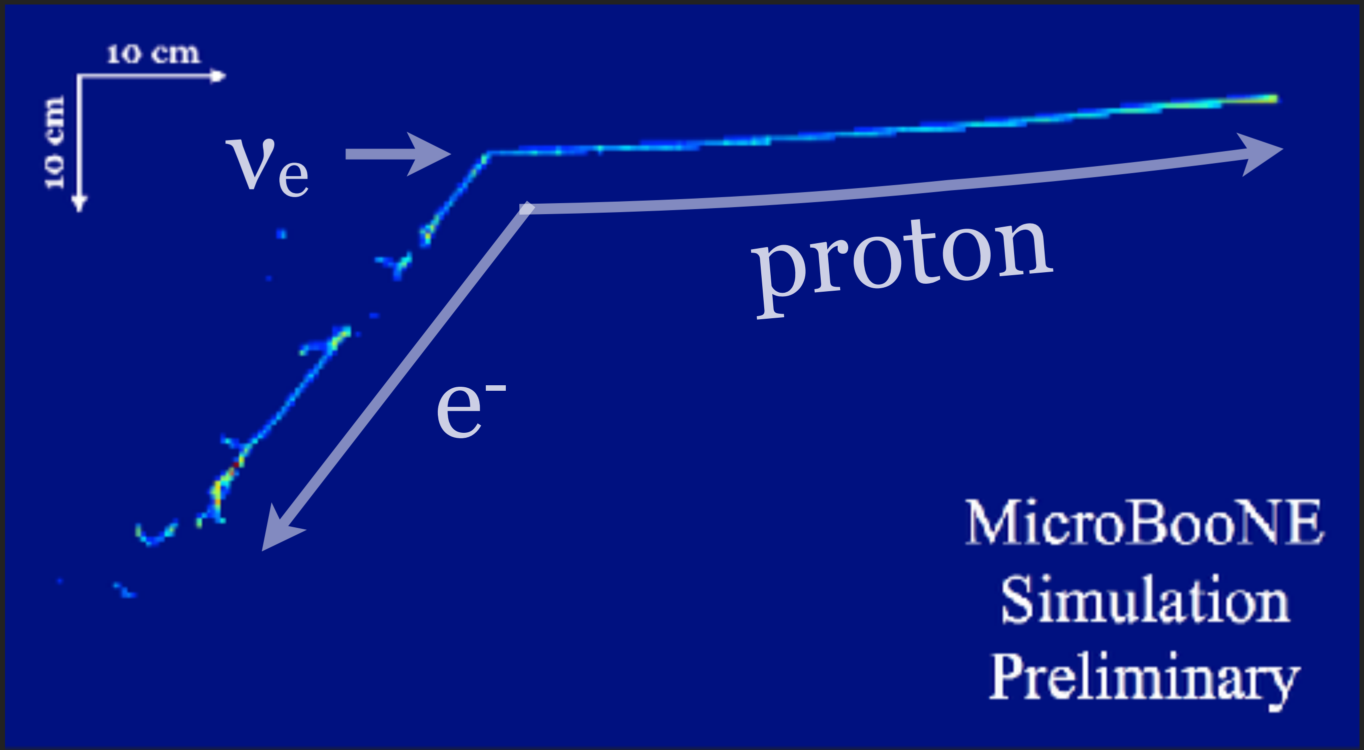

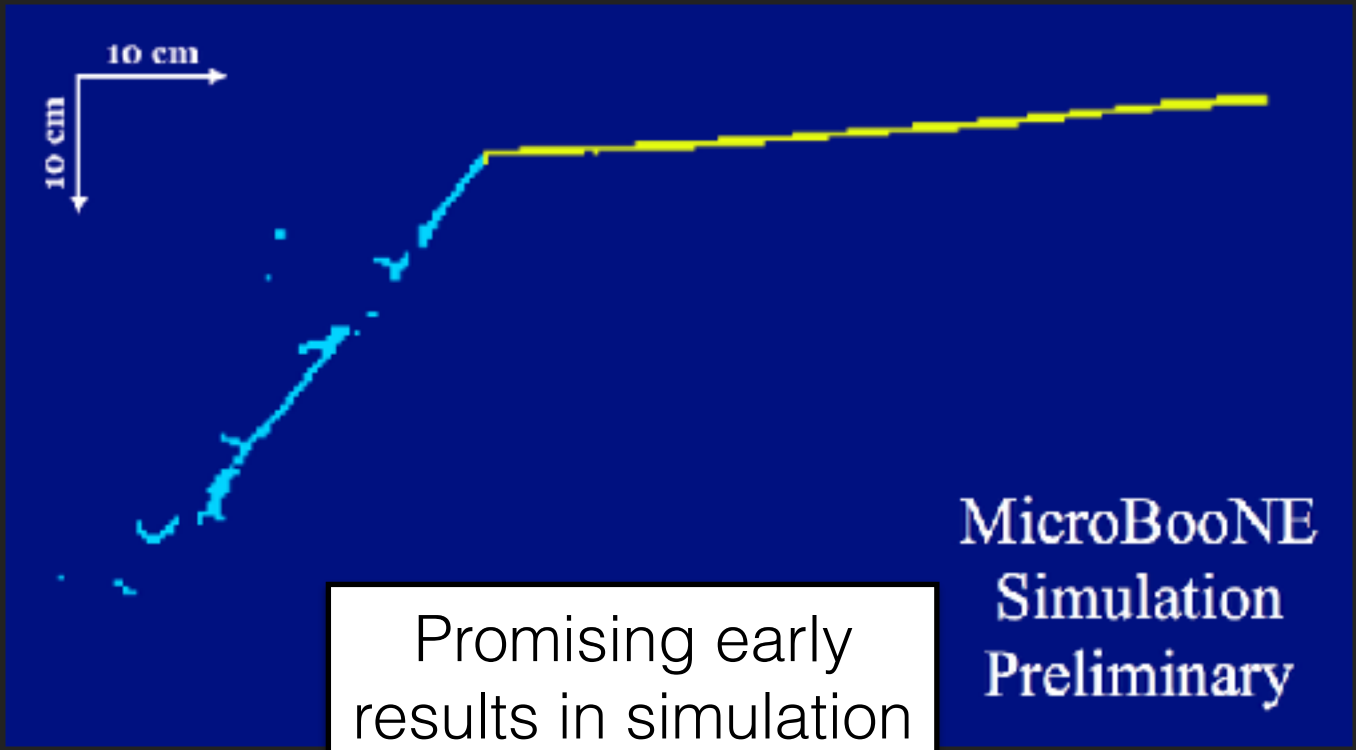

Promising early results in simulation and data samples

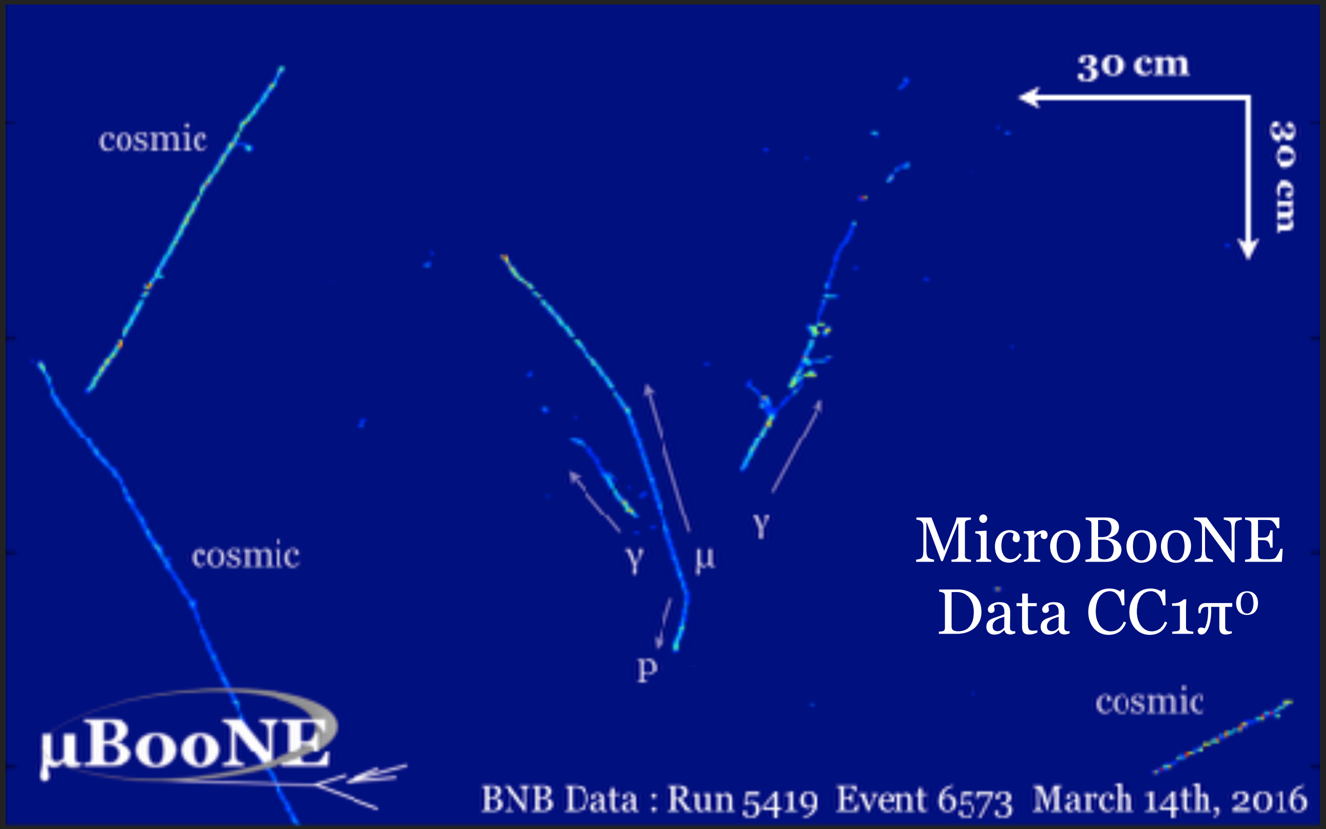

ADC Image

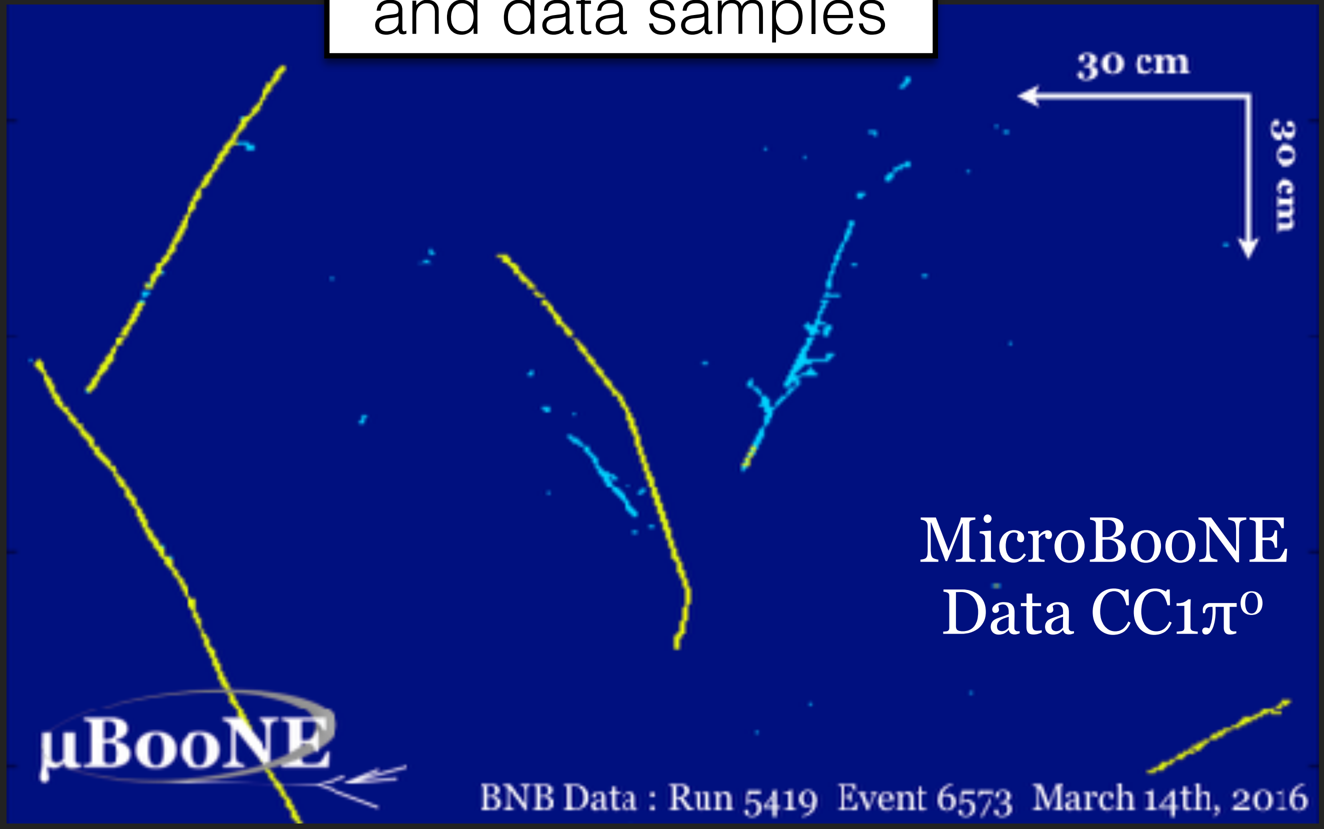

Network Output

NEXT STEPS

- We have incorporated some of the techniques we've developed into an analysis looking the low energy excess
  - See L. Yates talk on Thursday
  - Incorporates PID and Semantic Segmentation
- On-going effort to mitigate systematics from training on MC events
  - Testing on cosmic ray samples
  - Semantic aware-training
  - Feature-constrained training (to avoid leaning MC-specific features)

SUMMARY

- MicroBooNE is helping to pioneer the use of CNNs for LArTPC data
  - Classification, object detection, semantic segmentation
  - Details in paper: JINST 12 (02) P02017
- Also, working to understand how to bridge the MC-data divide
- Incorporating techniques into physics analyses
  - *See L. Yates Talk Thursday (Neutrino II afternoon, Comitium)*
- HEP-Friendly (i.e. ROOT) interfaces to Caffe and Tensorflow
  - LArCV: https://github.com/LArbys/LArCV
  - Caffe 1-fork: https://github.com/LArbys/caffe
  - Starting to think about LArSoft integration

THANK YOU

- Thanks for your attention
- And thank you to the funding agencies for making this work possible

BACK-UPS

Image, Network	Classified Particle Type				
	e^- [%]	γ [%]	μ^- [%]	π^- [%]	proton [%]
HiRes, AlexNet	73.6 ± 0.7	81.3 ± 0.6	84.8 ± 0.6	73.1 ± 0.7	87.2 ± 0.5
LoRes, AlexNet	64.1 ± 0.8	77.3 ± 0.7	75.2 ± 0.7	74.2 ± 0.7	85.8 ± 0.6
HiRes, GoogLeNet	77.8 ± 0.7	83.4 ± 0.6	89.7 ± 0.5	71.0 ± 0.7	91.2 ± 0.5
LoRes, GoogLeNet	74.0 ± 0.7	74.0 ± 0.7	84.1 ± 0.6	75.2 ± 0.7	84.6 ± 0.6

LONG-TERM VISION FOR DL IN LARTPCS

- Current:
  - replace/augment traditional algorithm tasks: PID, clustering, 2D->3D reconstruction
  - limit to tasks one can check with some kind of cosmic ray sample on DATA: MicroBooNE, protodune will have data
  - Systematics aware-training
  - employ in analyses
- Near-term:
  - SBND will have lots of neutrino interaction data
  - Train for tasks targeting neutrino interactions
  - Unsupervised techniques where Networks cluster data itself
- End-goal:
  - Recurrent Neural Network systems that perform interaction hypothesis search
  - Fast Hypothesis generation through Generative networks (e.g. GAN)
  - Reinforcement learning to teach network to solve interaction using self-taught decision tree for calling reco. algorithms
  - Output components of decision process to humans

How is it different from ***Image Classification***?

Example CNN for Image Classification

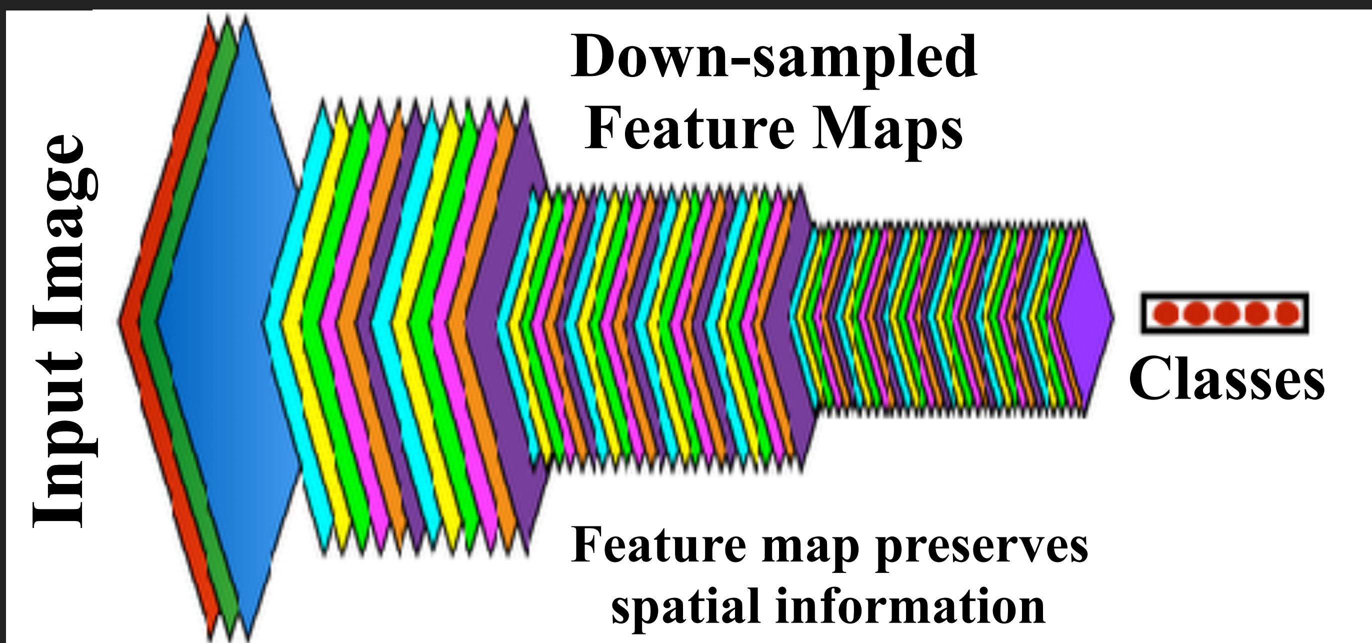

- Classification network reduces the whole image into final “class” 1D aray
- SSNet, after extracting class feature tensor, interpolates back into original image size

Example CNN for Semantic Segmentation

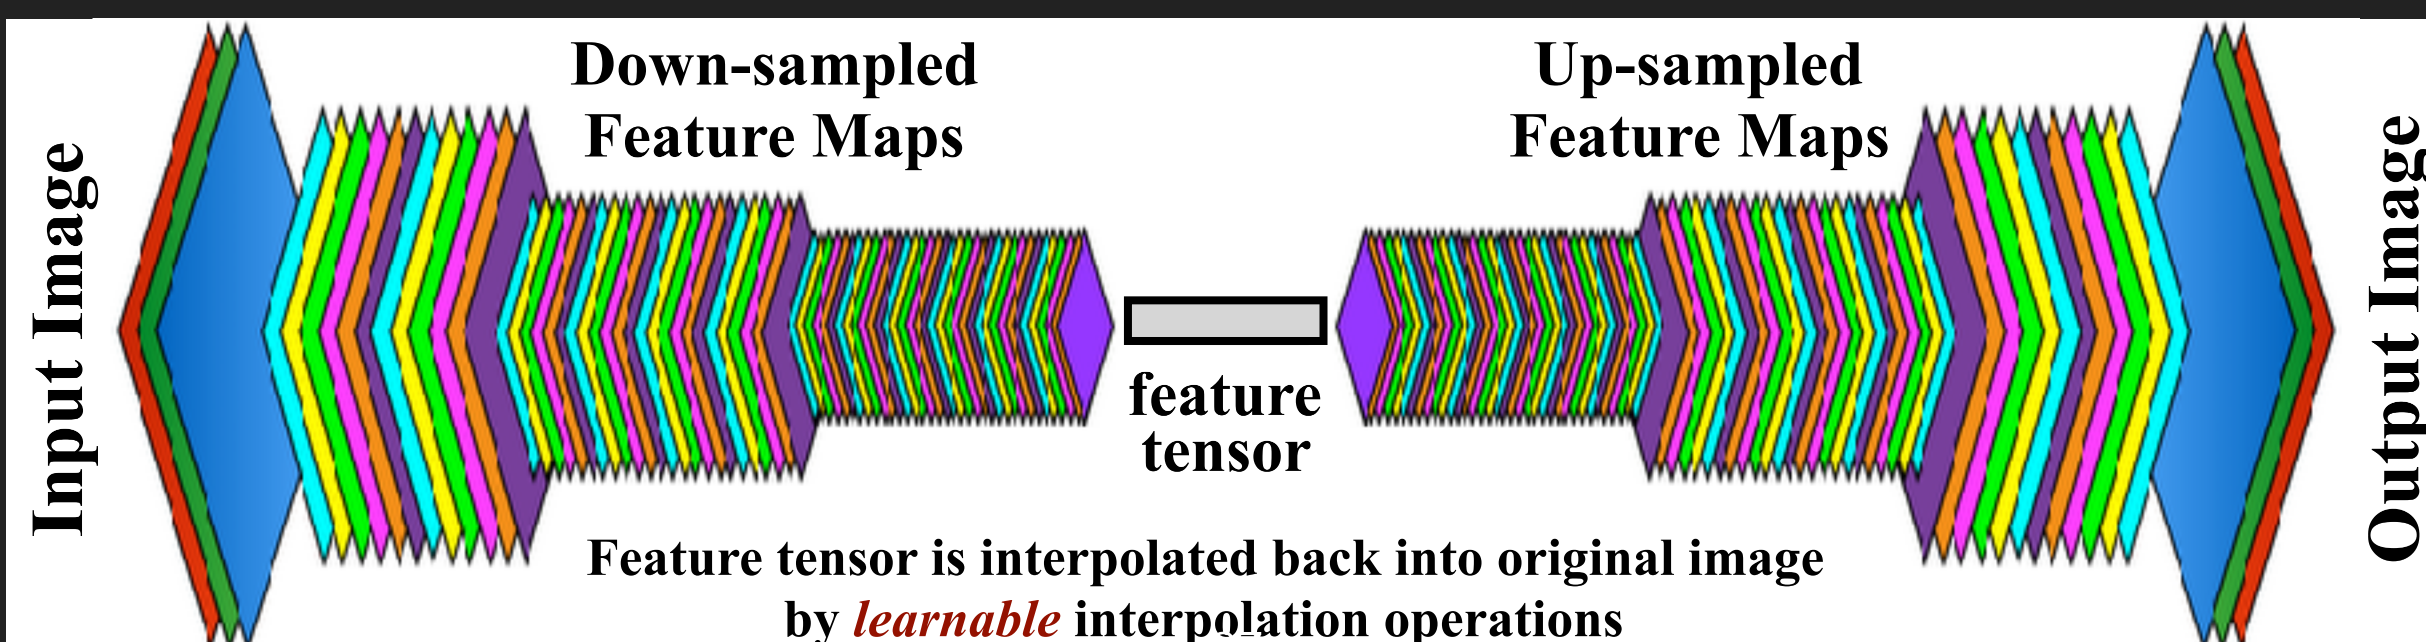

uBoone U-ResNet (or UBURN) Architecture

- U-Net gets it name from its graph diagram: network composed of a collapsing and expanding half, plus connections between low level and high-level feature maps

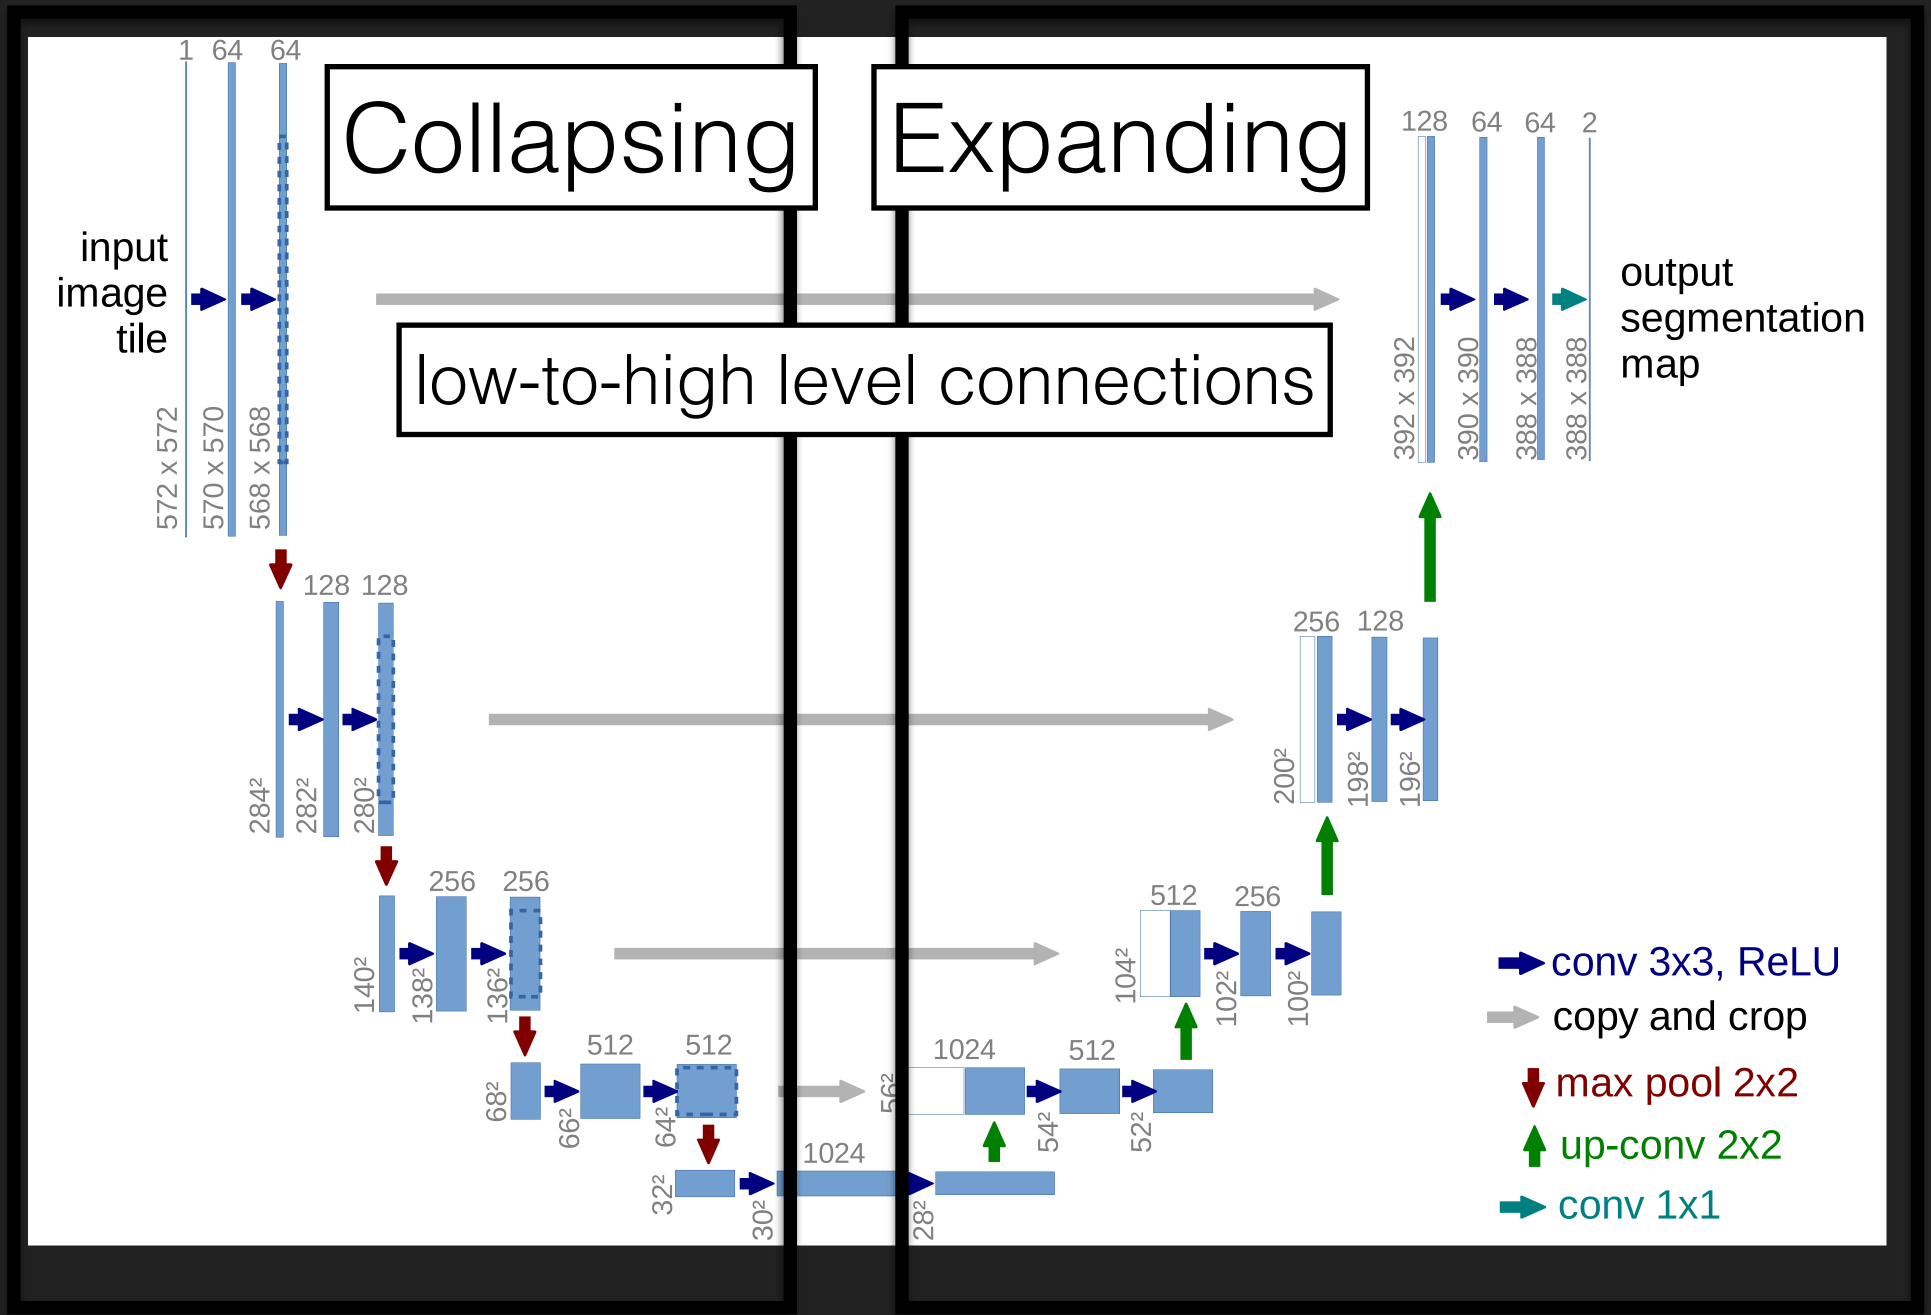

CLASSIFICATION

- Network used in paper
- Uses ResNet modules
- BatchNorm
- DropOut
- Convolution "stem" (purple and gold) where weights shared across application of 3 views

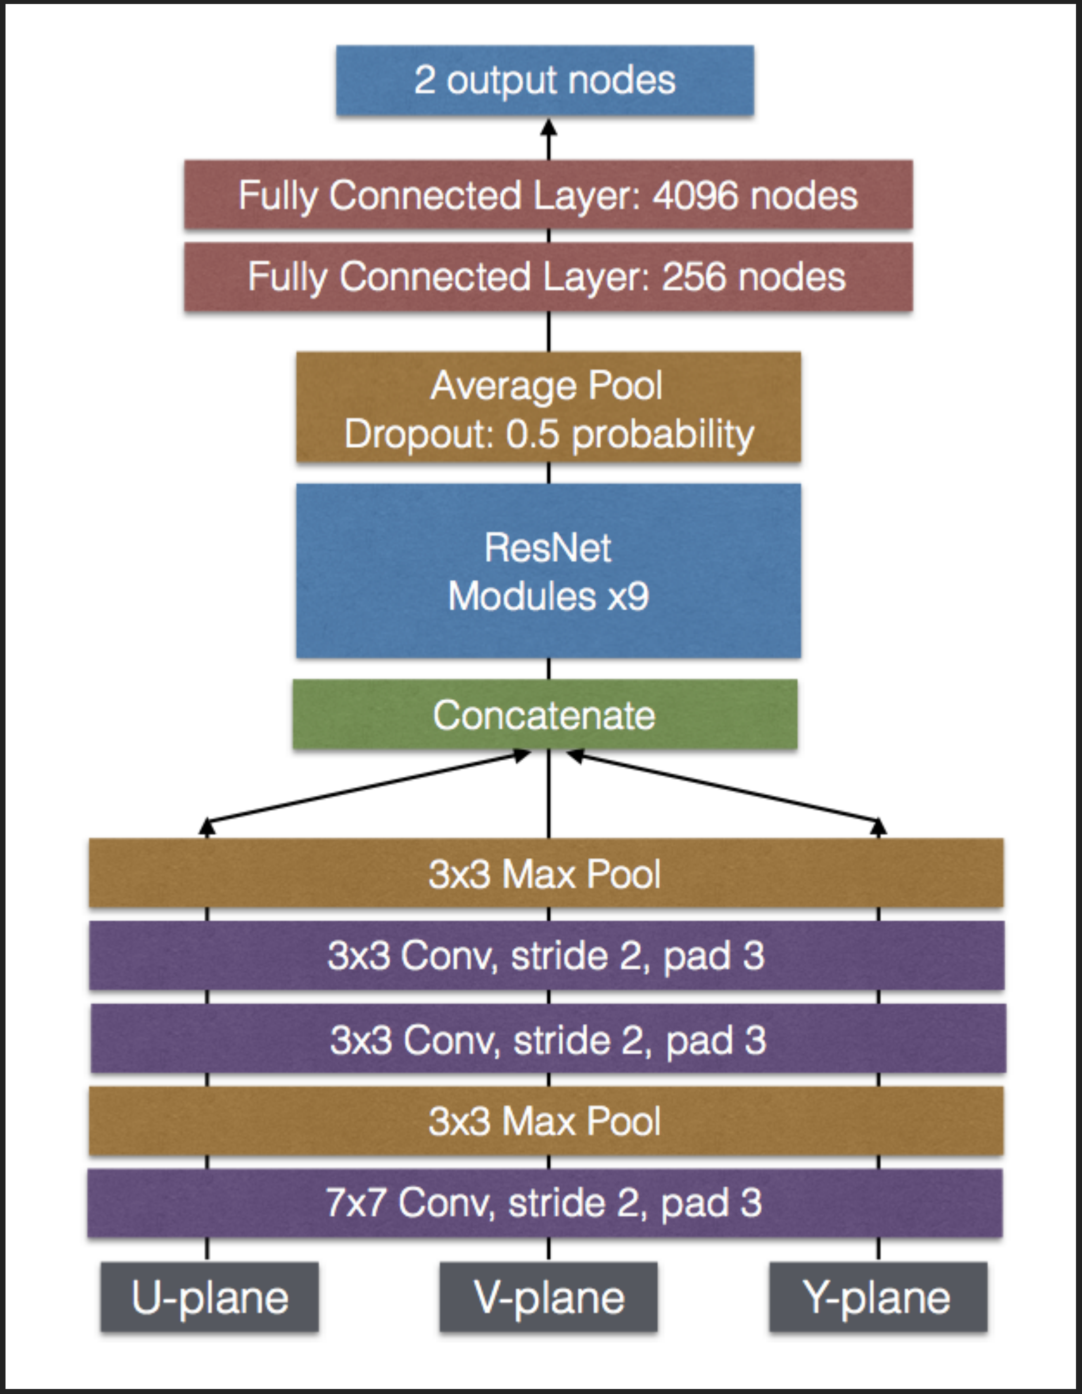

Generative Adversarial Networks (GANs)

A GAN is a CNN that takes in a random vector and transforms it into an image. The image produced is then fed through a classifier CNN, which classifies the image as either real or fake.

The goal of a GAN is to produce images that the classifier thinks are real.

A GAN that uses feature mapping has a modified goal: to produce images that, when fed through the classifier, cause the neurons in the classifier network to activate in the same way as they would when viewing real images.

http://arxiv.org/abs/1511.06434

arXiv:1606.03498

Feature Matching in GANs

Standard GAN: GAN is rewarded when classifier network classifies the image as real. **arXiv:1511.06434**

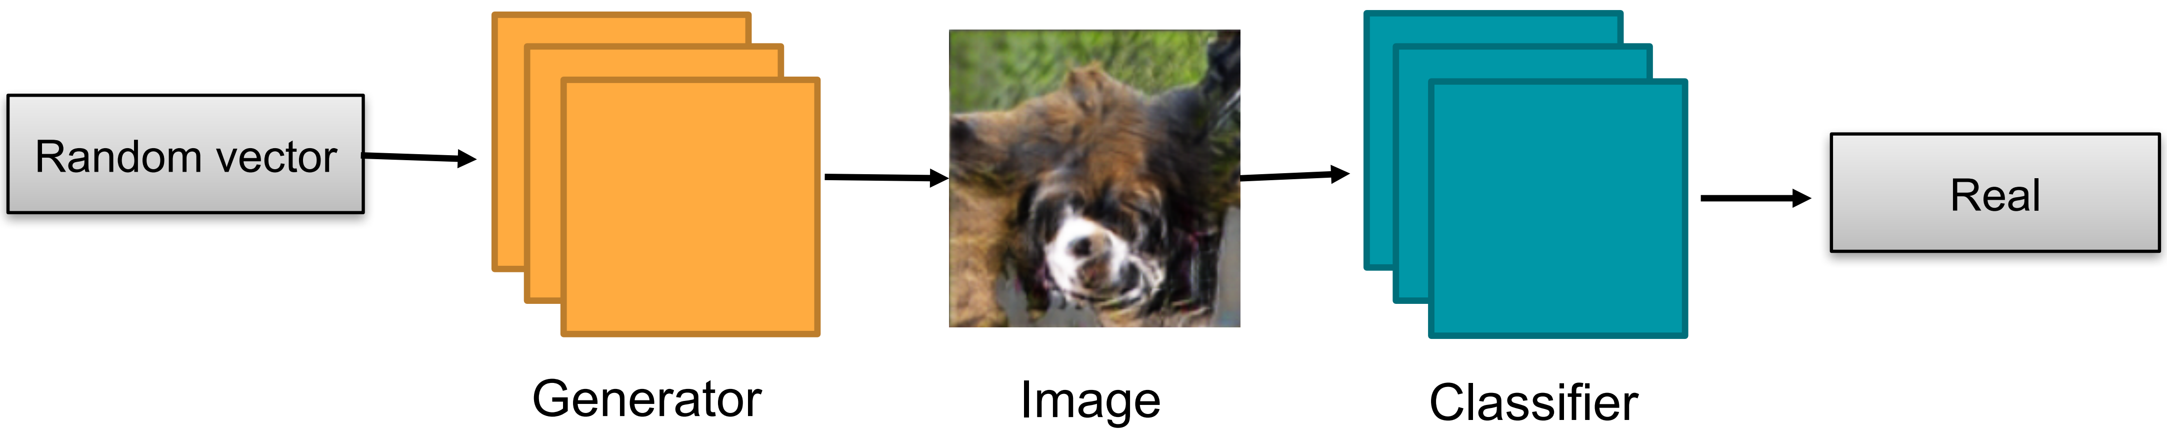

Feature-matching GAN: GAN is rewarded when neurons in an intermediate layer of the classifier network activate in the same way as when viewing a real image. **arXiv:1606.03498**

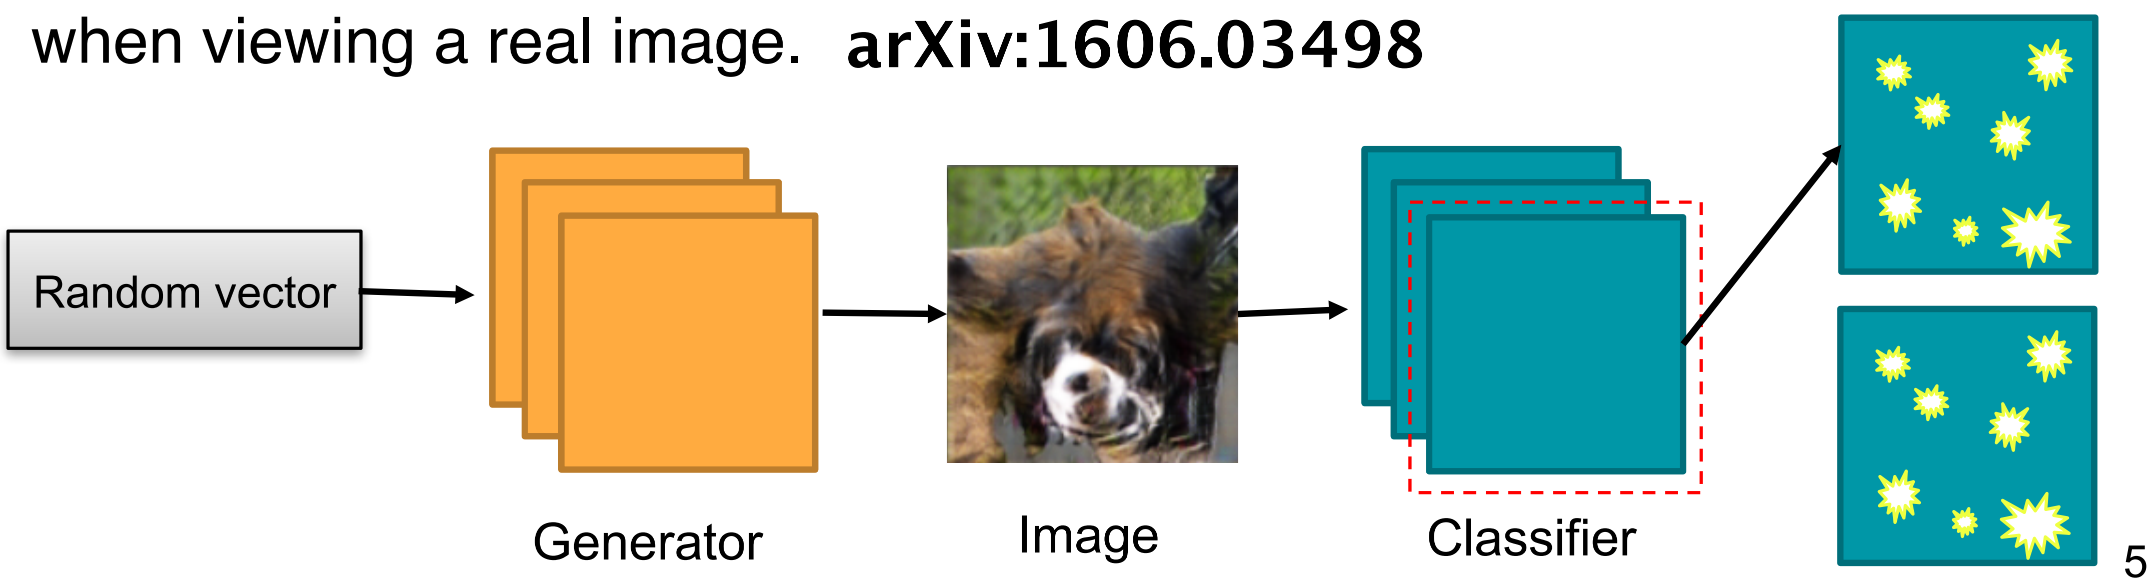

Real image activation

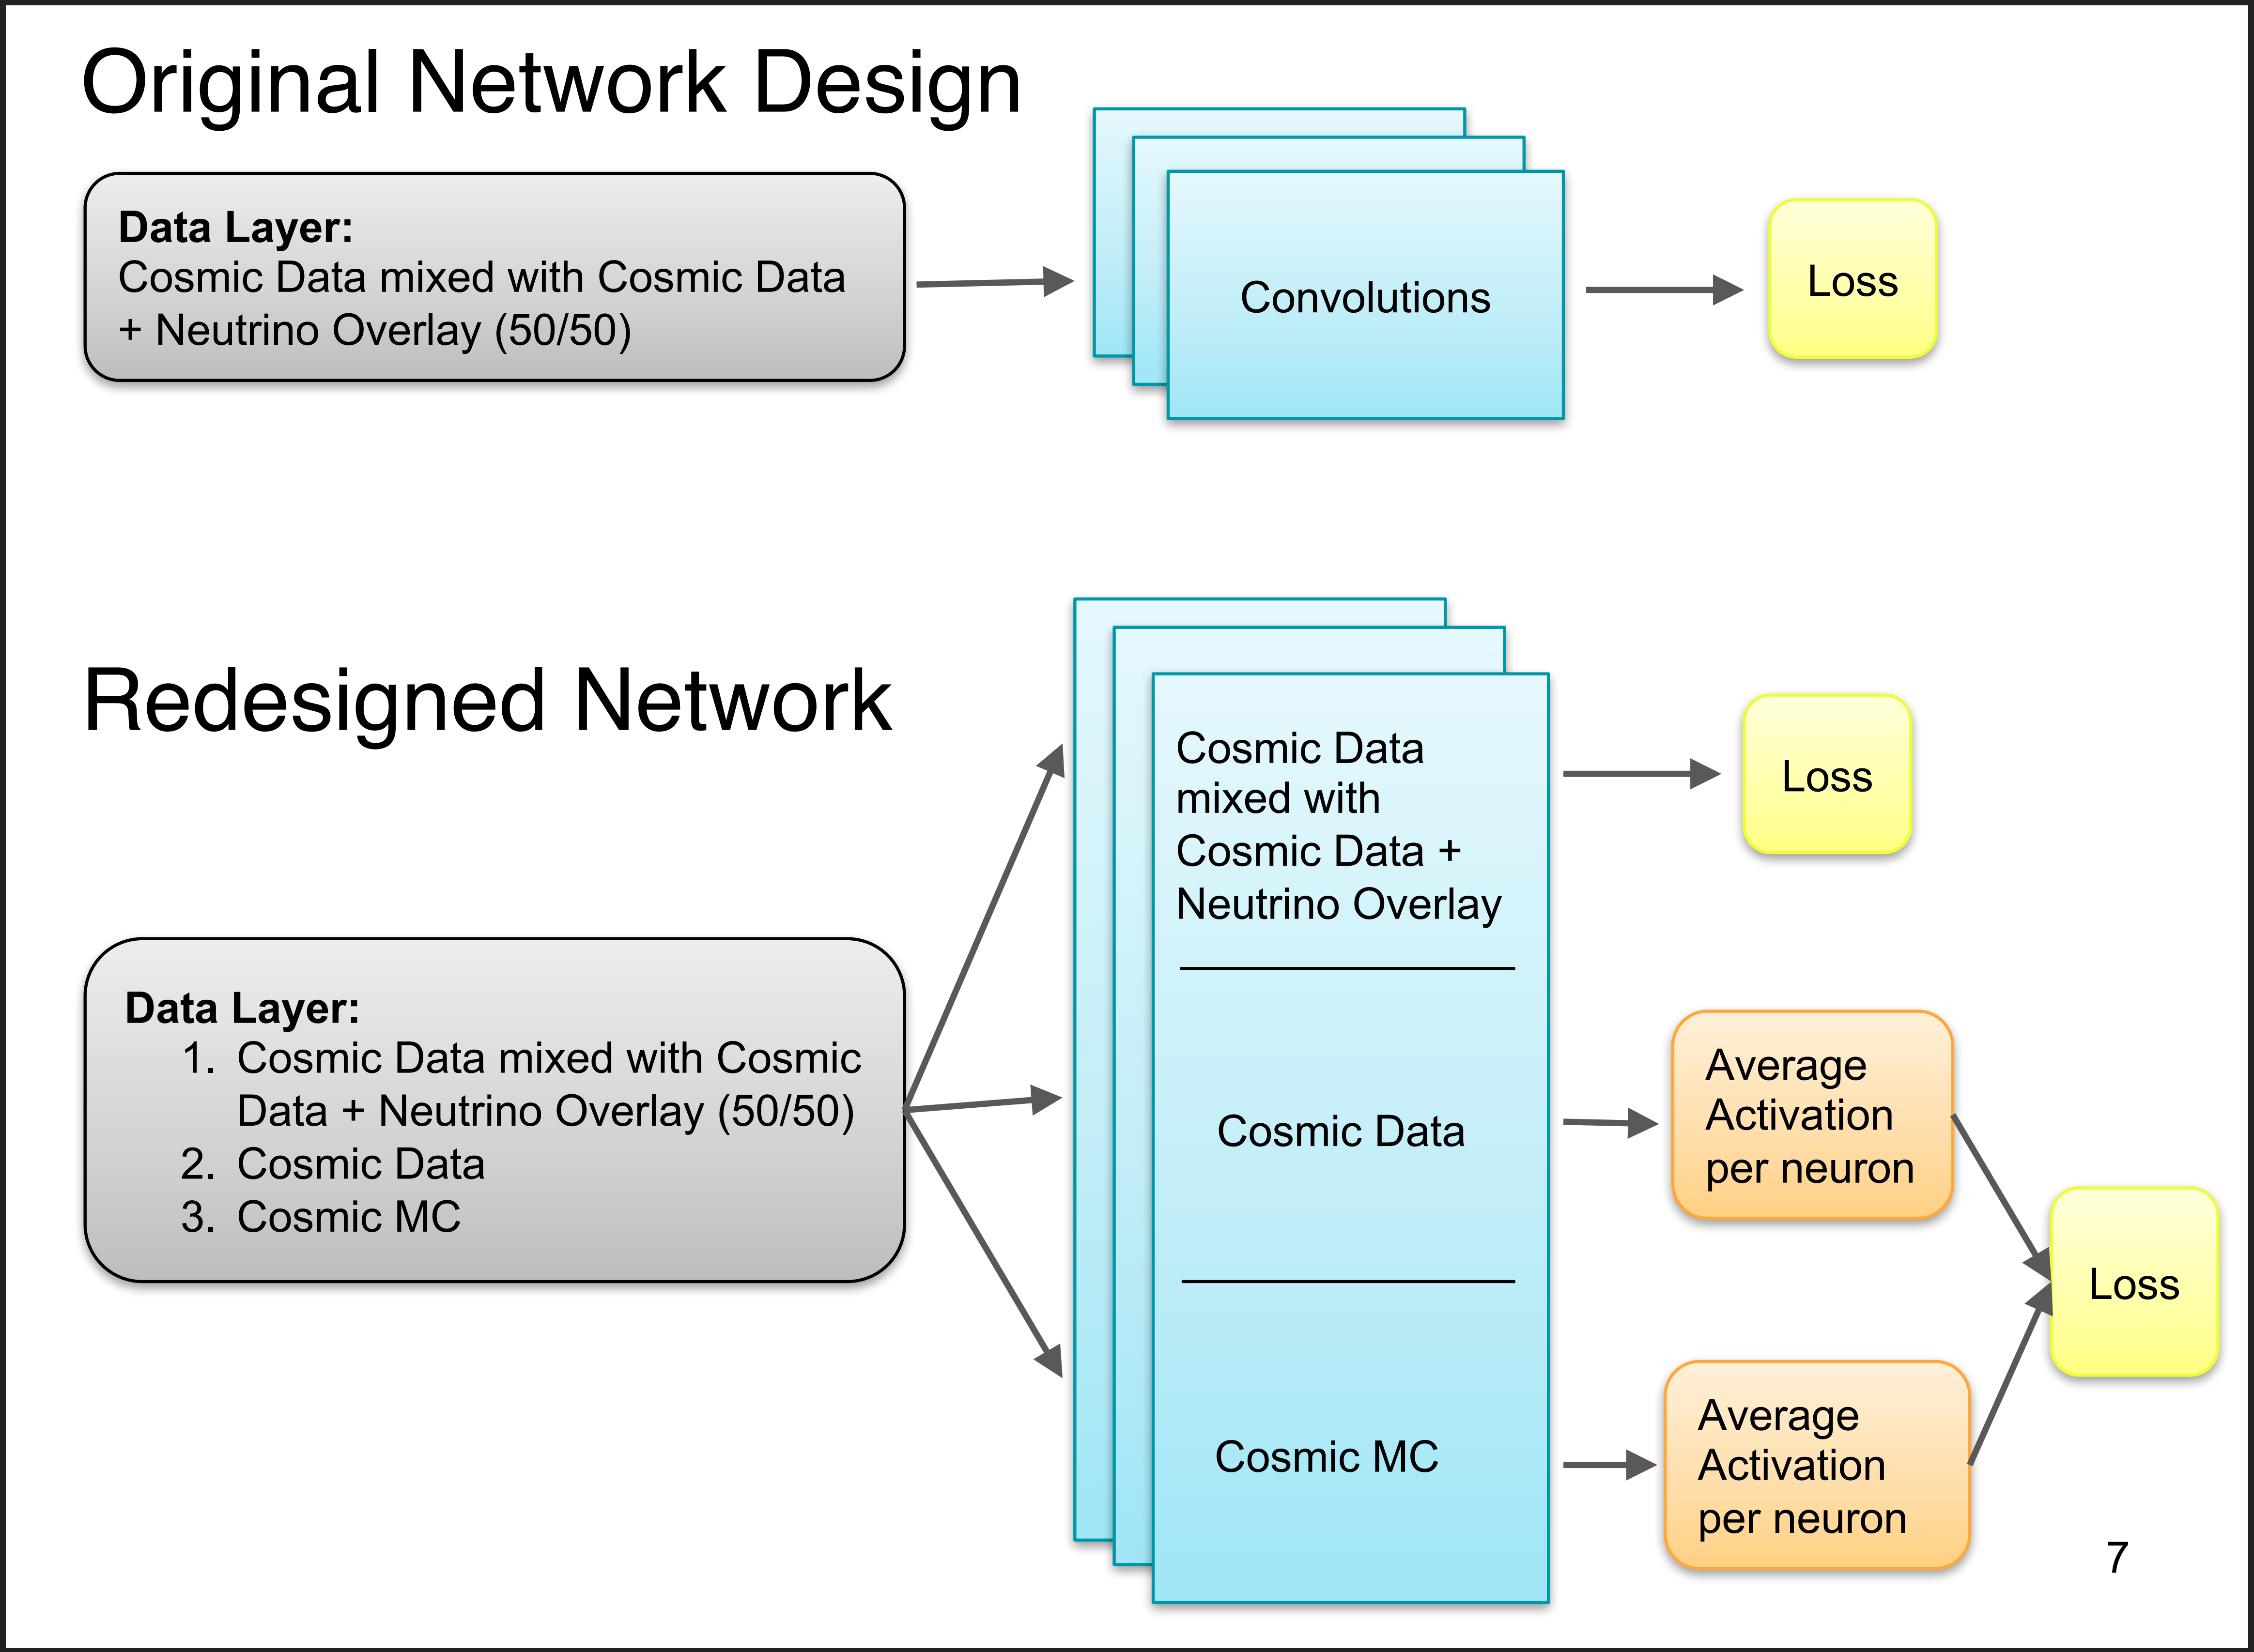
Original Network Design
Data Layer:
Cosmic Data mixed with Cosmic Data + Neutrino Overlay (50/50)
Convolutions
Loss
Redesigned Network
Data Layer:
1. Cosmic Data mixed with Cosmic Data + Neutrino Overlay (50/50)
2. Cosmic Data
3. Cosmic MC
Cosmic Data mixed with Cosmic Data + Neutrino Overlay
Cosmic Data
Cosmic MC
Loss
Average Activation per neuron
Average Activation per neuron
Loss
7

WHAT IS STABILITY TRAINING?

https://arxiv.org/pdf/1604.04326.pdf

- Small perturbations in images can cause large shifts in classification scores
- We modify our loss function with a "Stability Term"
- Run "original image" and "original image plus gaussian noise" and minimize difference in score

$$L_{total} = L_0 + L_{stability}$$

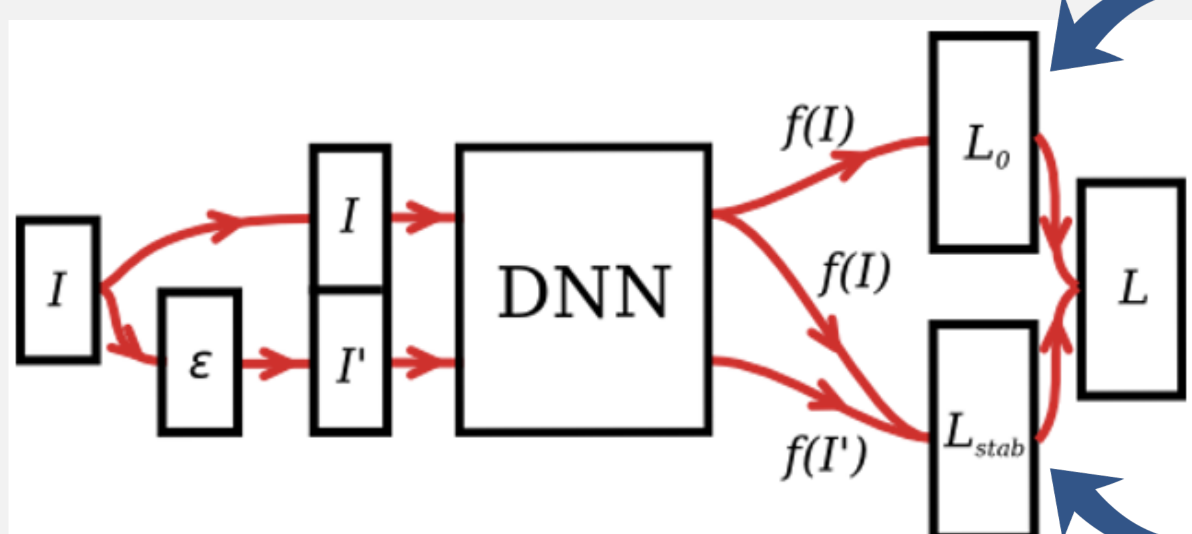

"Was classification correct?"

"Did perturbation change the score?"