MACHINE LEARNING AT FERMILAB

A summary of recent applications in HEP and ML community efforts

Fernanda Psihas

Indiana University

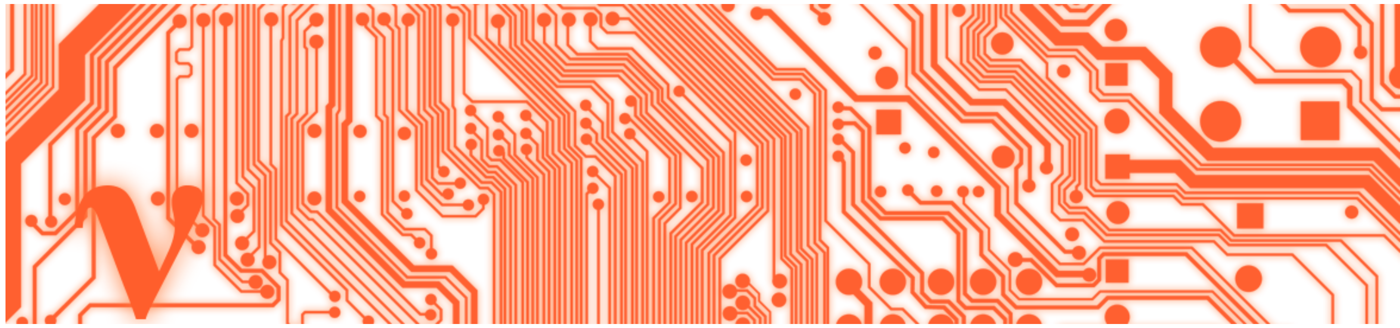

What Is Machine Learning?

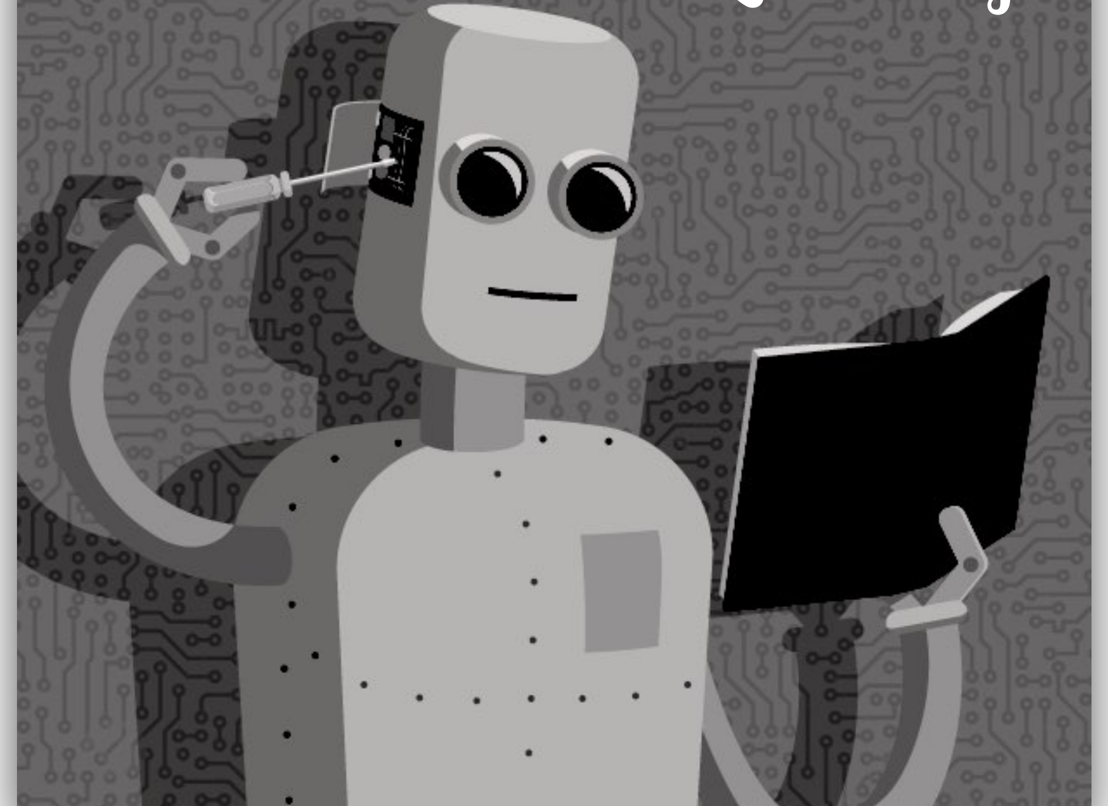

Applications in Fermilab Experiments

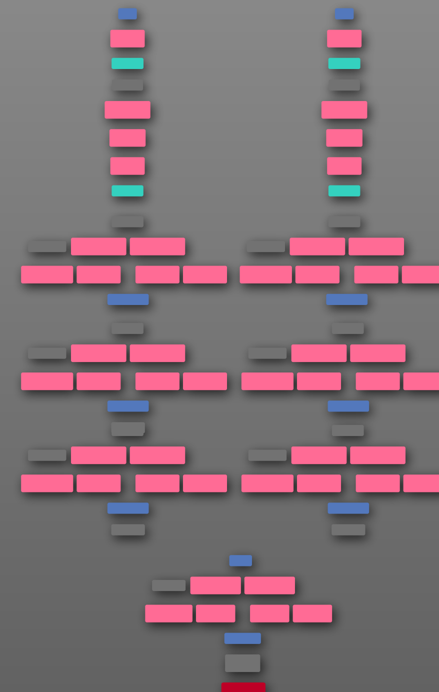

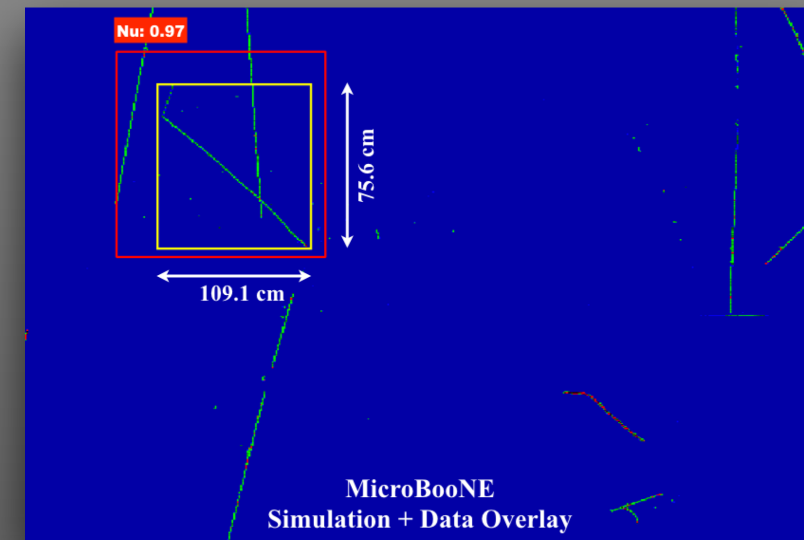

Challenges in ML for HEP

SYSTEMATICS

MODEL COMPLEXITY

ROBUSTNESS

PHYSICS INTERPRETATION

Resources and ML Community

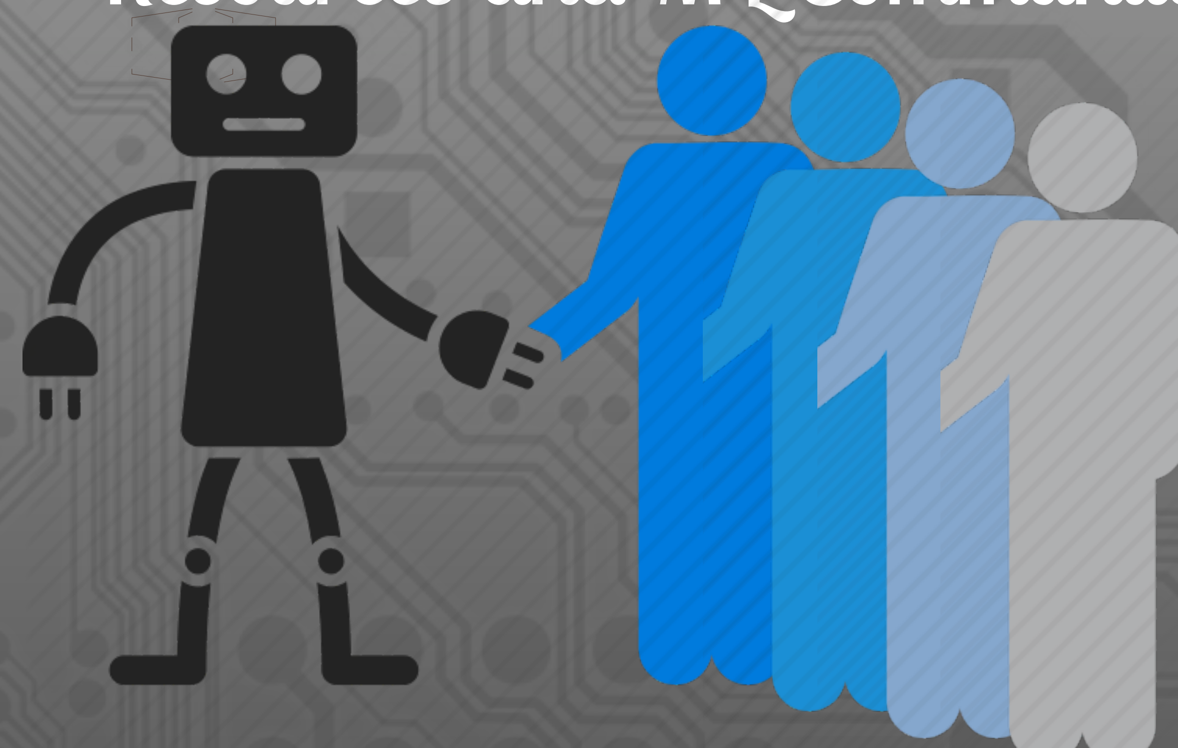

Machine Learning Algorithms

Algorithms whose PERFORMANCE for a given TASK improves with EXPERIENCE

TASKS IN HEP

Identification, Reconstruction, MC Generation, Pattern Recognition, Estimation of Physics Quantities

PERFORMANCE METRICS

Accuracy, Running time, Sensitivity, Bias reduction

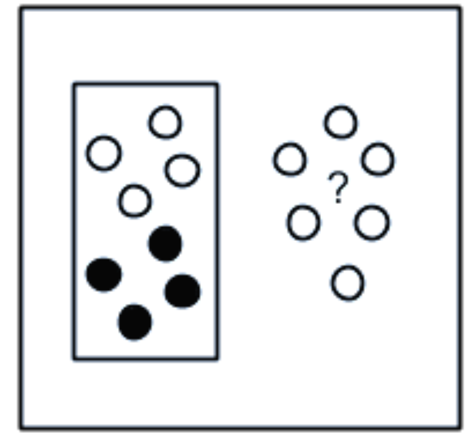

Supervised Learning Algorithms

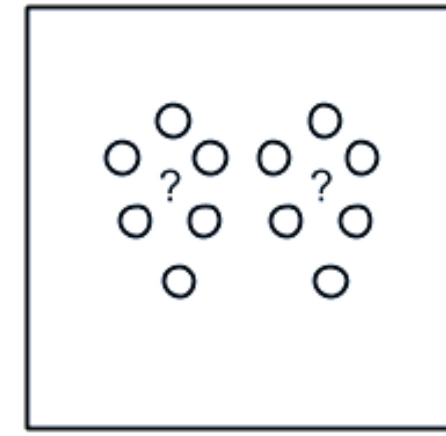

Unsupervised Learning Algorithms

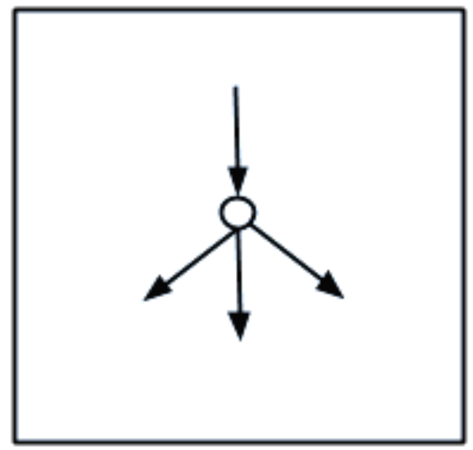

Artificial Neural Network Algorithms

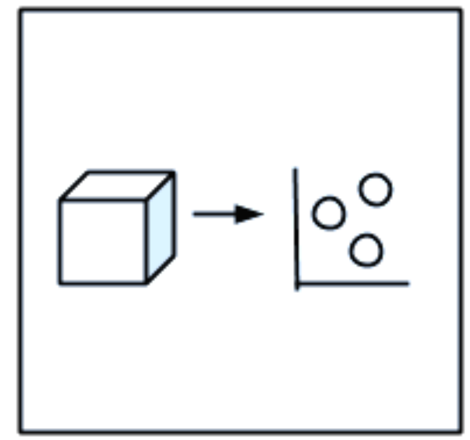

Dimensional Reduction Algorithms

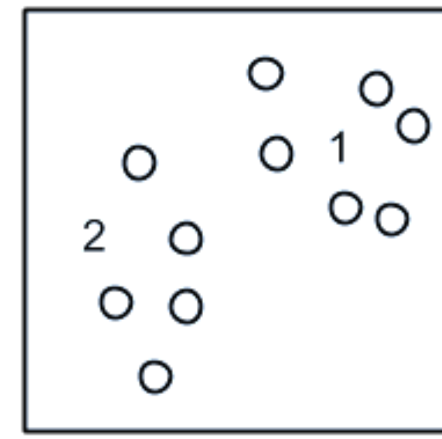

Clustering Algorithms

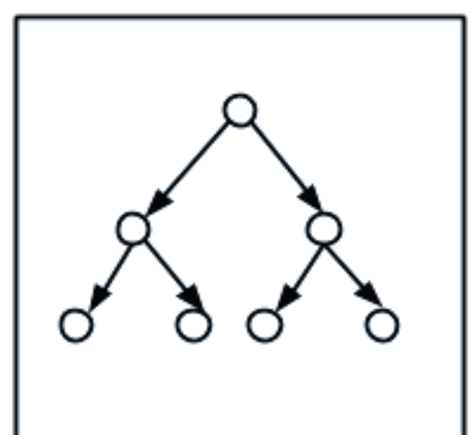

Decision Tree Algorithms

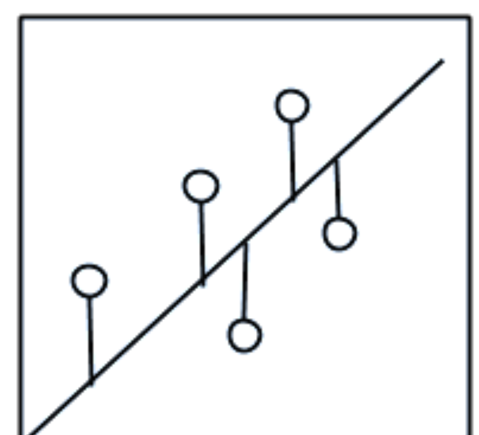

Regression Algorithms

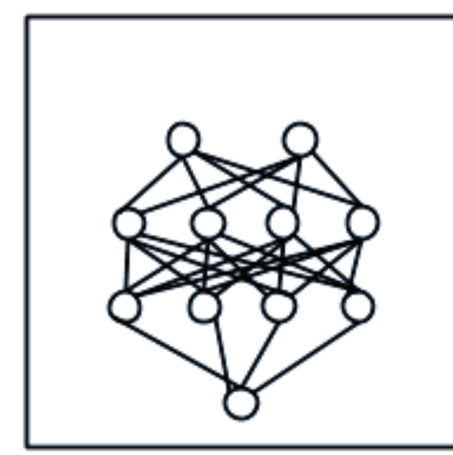

Deep Learning Algorithms

APPLICATIONS IN HEP

ML APPLICATIONS IN HEP

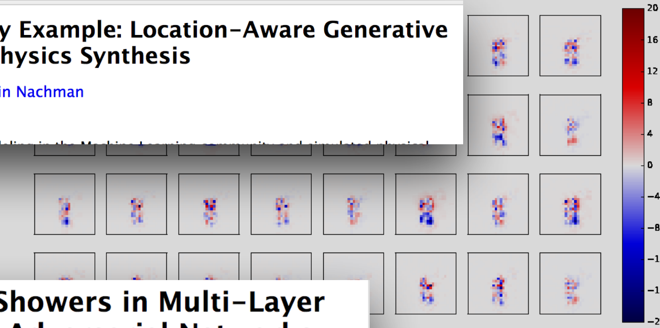

Learning Particle Physics by Example: Location-Aware Generative Adversarial Networks for Physics Synthesis

Luke de Oliveira, Michela Paganini, Benjamin Nachman

(Submitted on 20 Jan 2017)

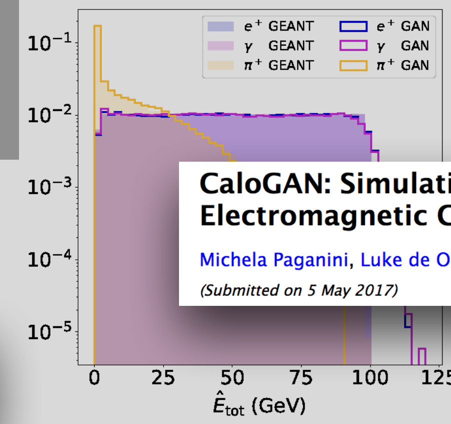

CaloGAN: Simulating 3D High Energy Particle Showers in Multi-Layer Electromagnetic Calorimeters with Generative Adversarial Networks

Michela Paganini, Luke de Oliveira, Benjamin Nachman

(Submitted on 5 May 2017)

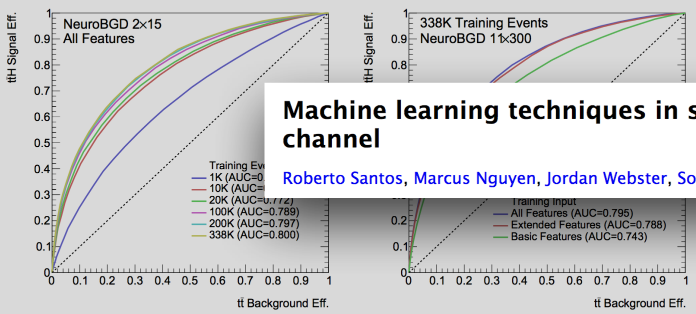

Machine learning techniques in searches for $t\bar{t}h$ in the $h \to b\bar{b}$ decay channel

Roberto Santos, Marcus Nguyen, Jordan Webster, Soo Ryu, Jahred Adelman, Sergei Chekanov, Jie Zhou

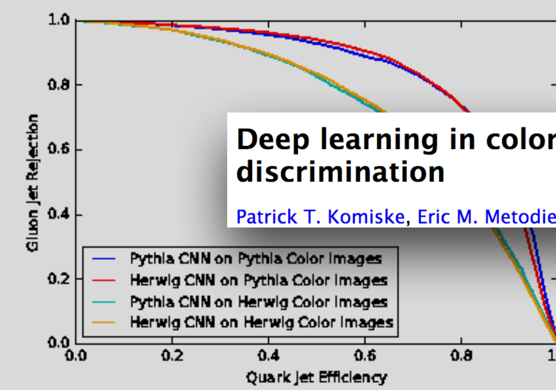

Deep learning in color: towards automated quark/gluon jet discrimination

Patrick T. Komiske, Eric M. Metodiev, Matthew D. Schwartz

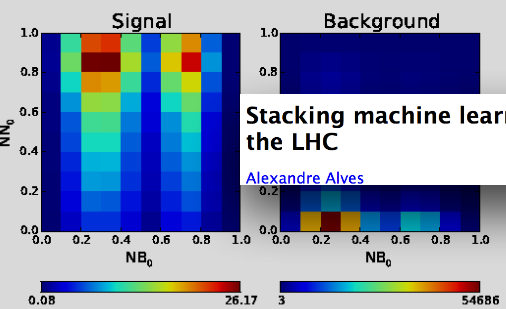

Stacking machine learning classifiers to identify Higgs bosons at the LHC

Alexandre Alves

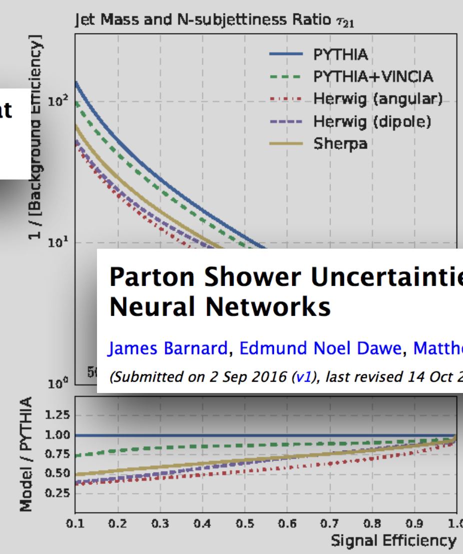

Parton Shower Uncertainties in Jet Substructure Analyses with Deep Neural Networks

James Barnard, Edmund Noel Dawe, Matthew J. Dolan, Nina Rajcic

(Submitted on 2 Sep 2016 (v1), last revised 14 Oct 2016 (this version, v2))

Deep Neural Networks to Enable Real-time Multimessenger Astrophysics

Daniel George, E. A. Huerta

(Submitted on 30 Dec 2016 (v1), last revised 4 Jan 2017 (this version, v2))

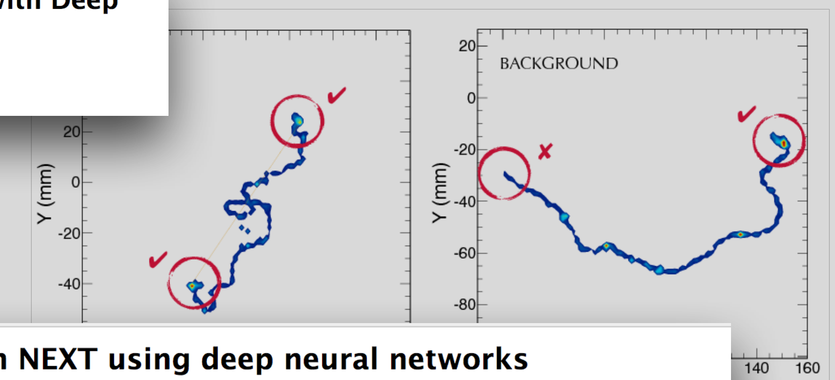

Background rejection in NEXT using deep neural networks

NEXT Collaboration: J. Renner, A. Farbin, J. Muñoz Vidal, J.M. Benlloch-Rodríguez, A. Botas, P. Ferrario, J.J.

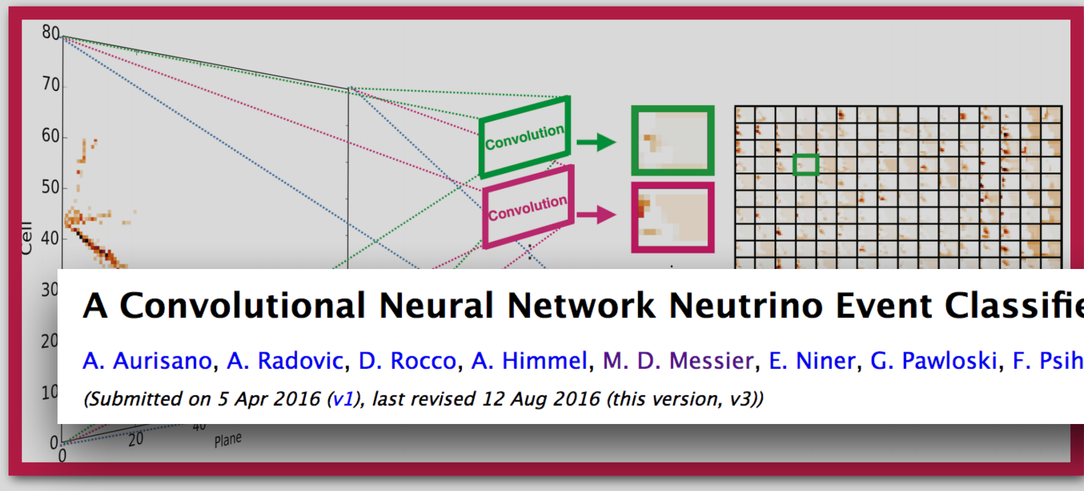

A Convolutional Neural Network Neutrino Event Classifier

A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. D. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, P. Vahle

(Submitted on 5 Apr 2016 (v1), last revised 12 Aug 2016 (this version, v3))

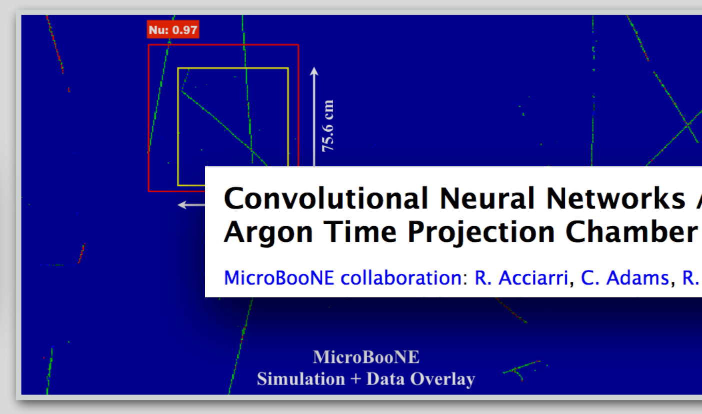

Convolutional Neural Networks Applied to Neutrino Events in a Liquid Argon Time Projection Chamber

MicroBooNE collaboration: R. Acciarri, C. Adams, R. An, J. Asaadi, M. Auger, L. Bagby, B. Baller, G. Barr, M. Bass,

MACHINE LEARNING

ML techniques employed in many CMS analyses, including the Higgs discovery.

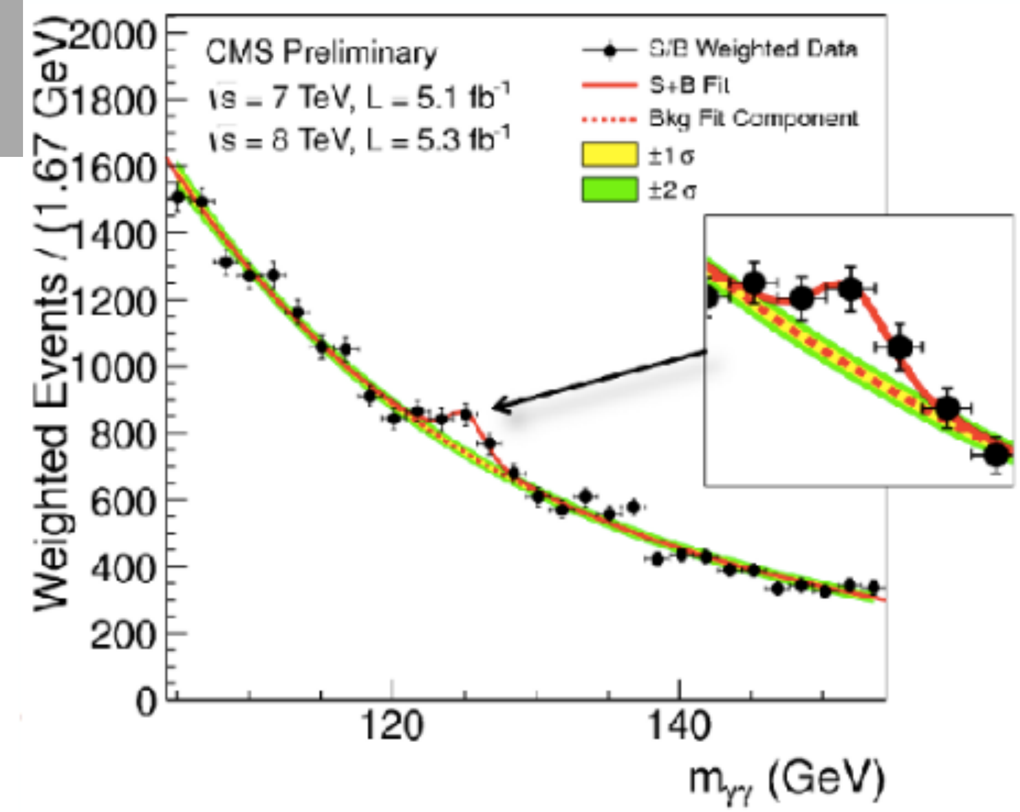

CMS has a broad program of machine learning applications for multiple tasks, including:

- ★ Reconstruction
- ★ PIDs and Tagging
- ★ MC Generation
- ★ Imaging Calorimetry
- ★ Tracking
- ★ Data Quality Monitoring
- ★ Trigger

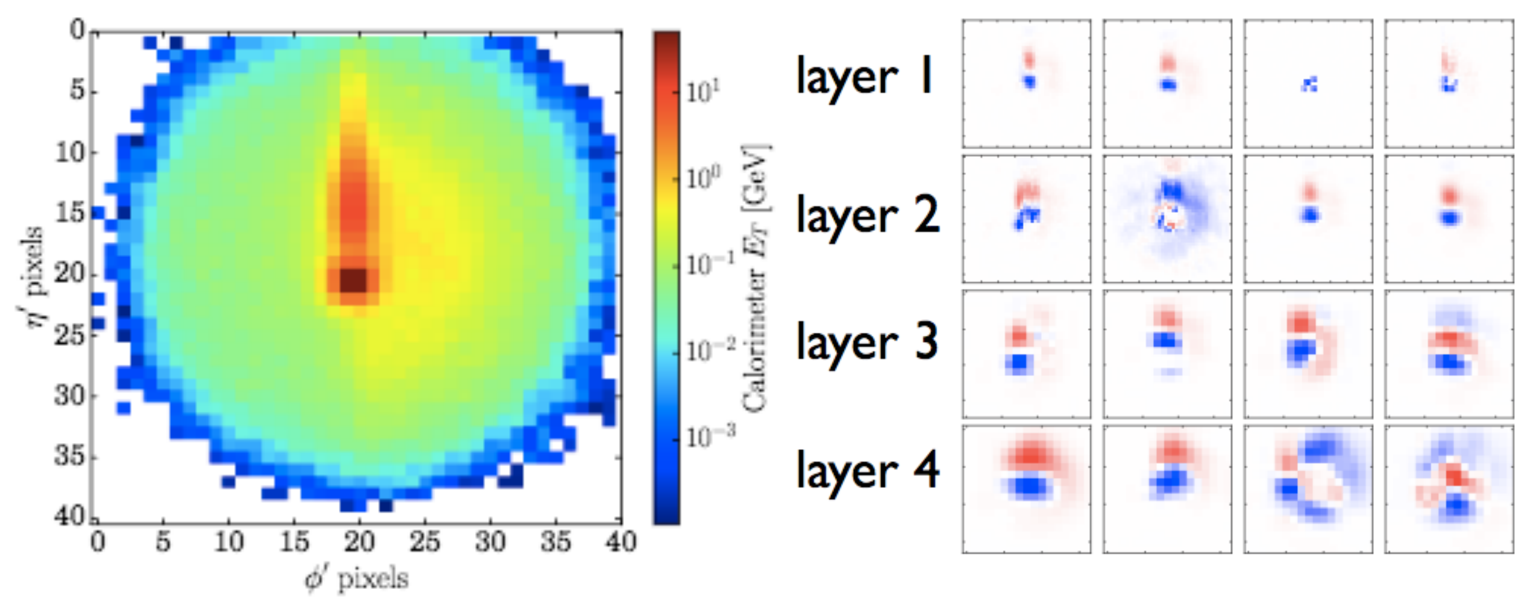

Deep-learning Top Taggers or The End of QCD?

Gregor Kasieczka, Tilman Plehn, Michael Russell, Torben Schell

(Submitted on 30 Jan 2017 (v1), last revised 16 May 2017 (this version, v2))

Convolutional Visual Network

NOvA has the first implementation of Convolutional Neural Networks on a HEP result.

- ★ *Advantage from extracting features to learn from, rather than learn from traditional reconstruction*
- ★ *CVN PID represented an equivalent increase of 30% exposure*

Ongoing program to incorporate deep learning for end-to-end reconstruction.

A Convolutional Neural Network Neutrino Event Classifier

A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. D. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, P. Vahle

(Submitted on 5 Apr 2016 (v1), last revised 12 Aug 2016 (this version, v3))

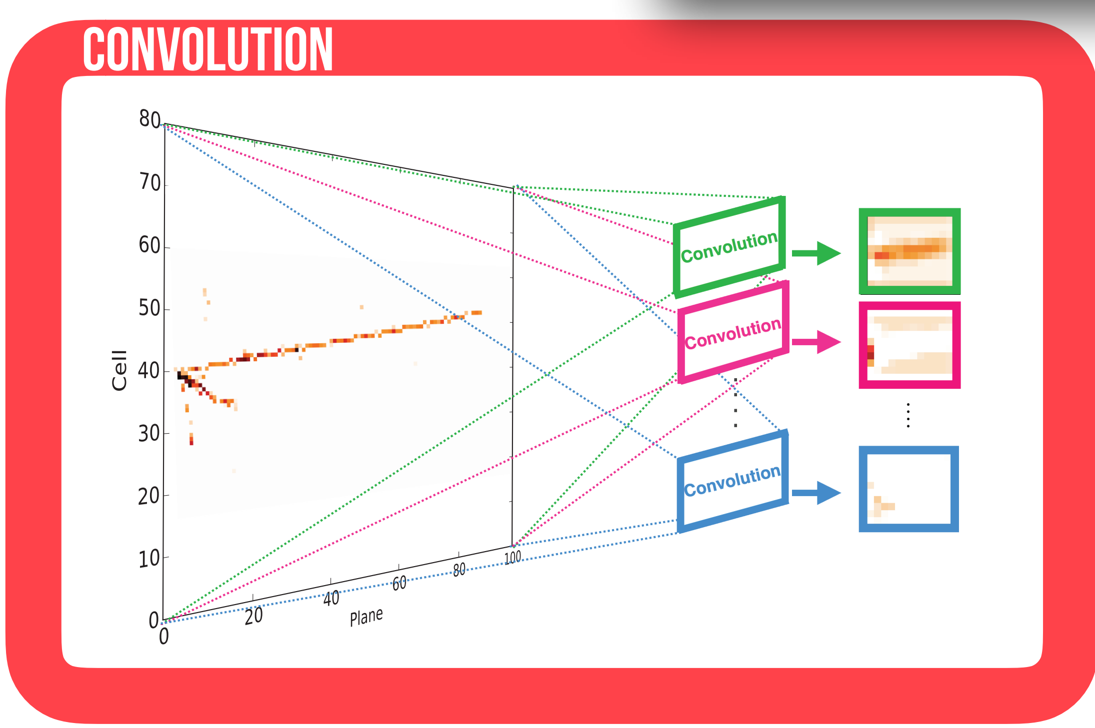

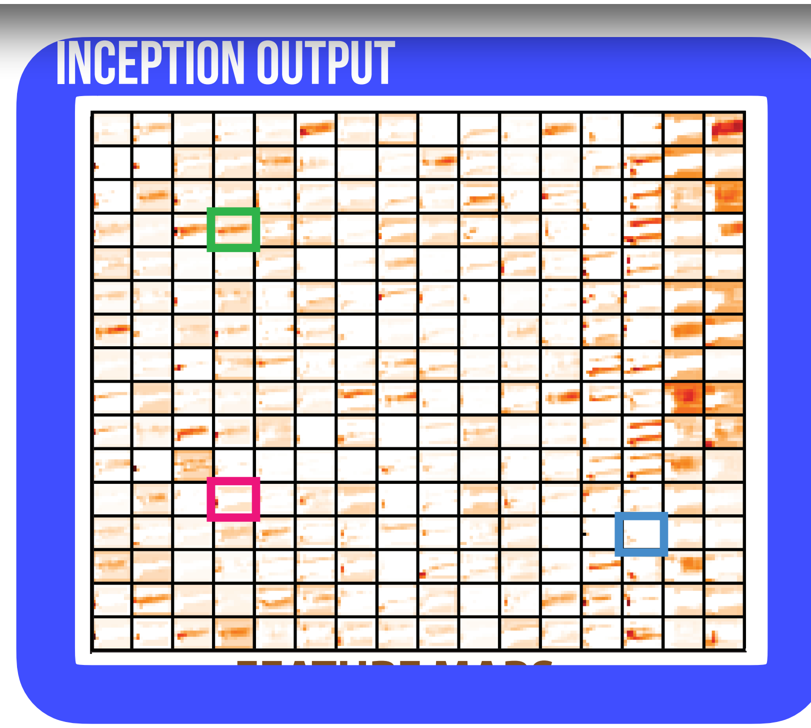

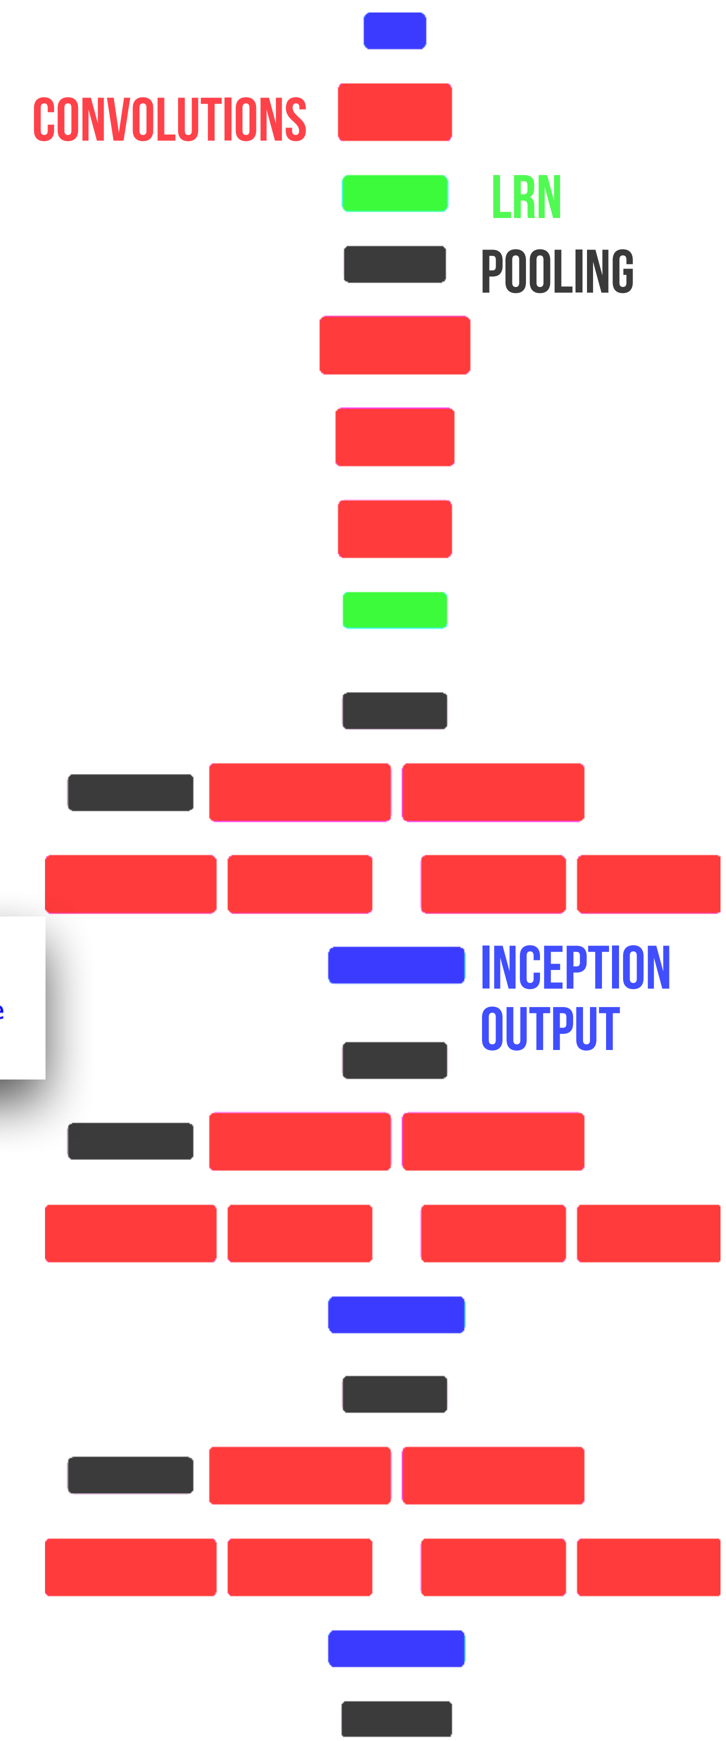

CNN'S FOR NEUTRINO EVENTS

MicroBooNE is exploring CNN implementations on LAr-TPC for:

- *Neutrino interaction detection 85% efficiency*
- *Multi-particle classification 83% efficiency for electrons and 95% efficiency for muons*

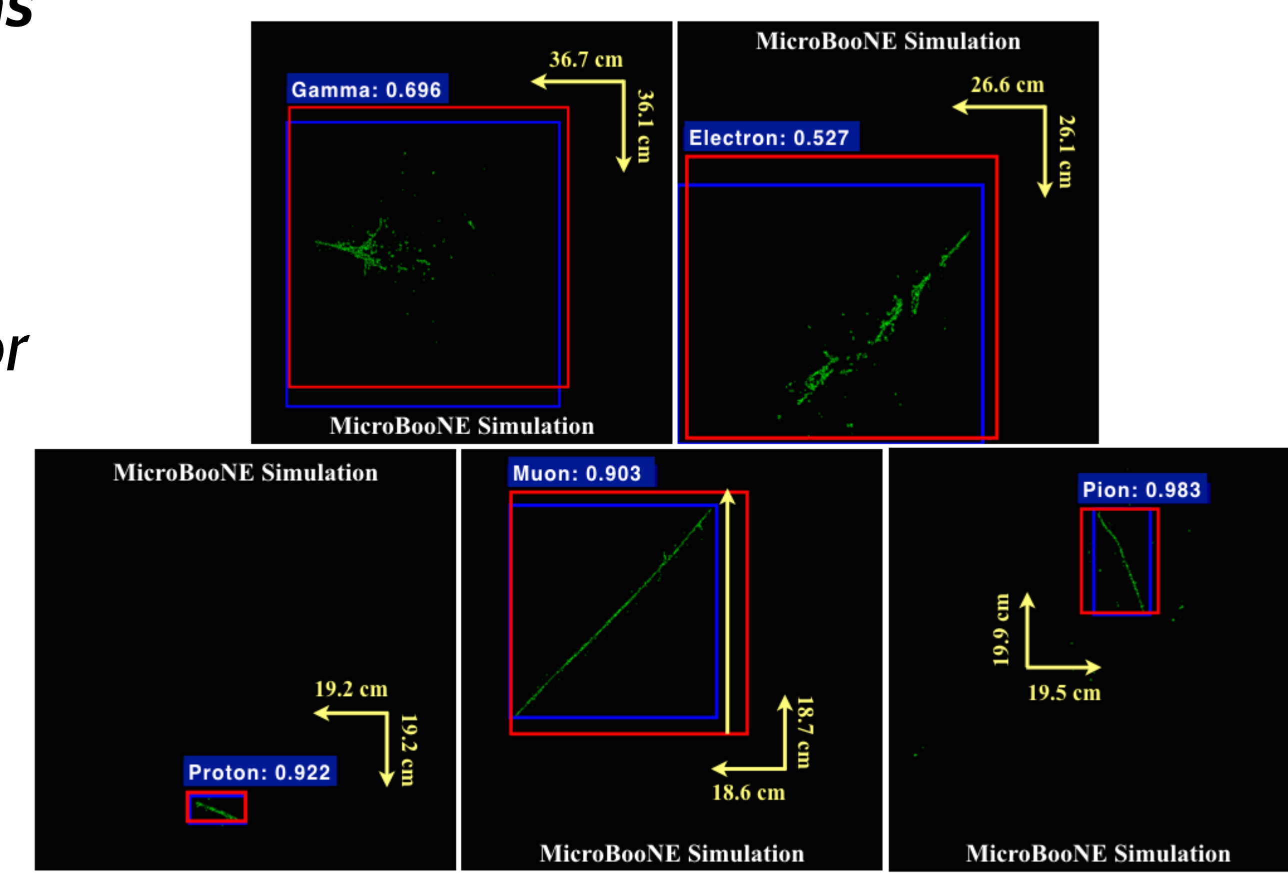

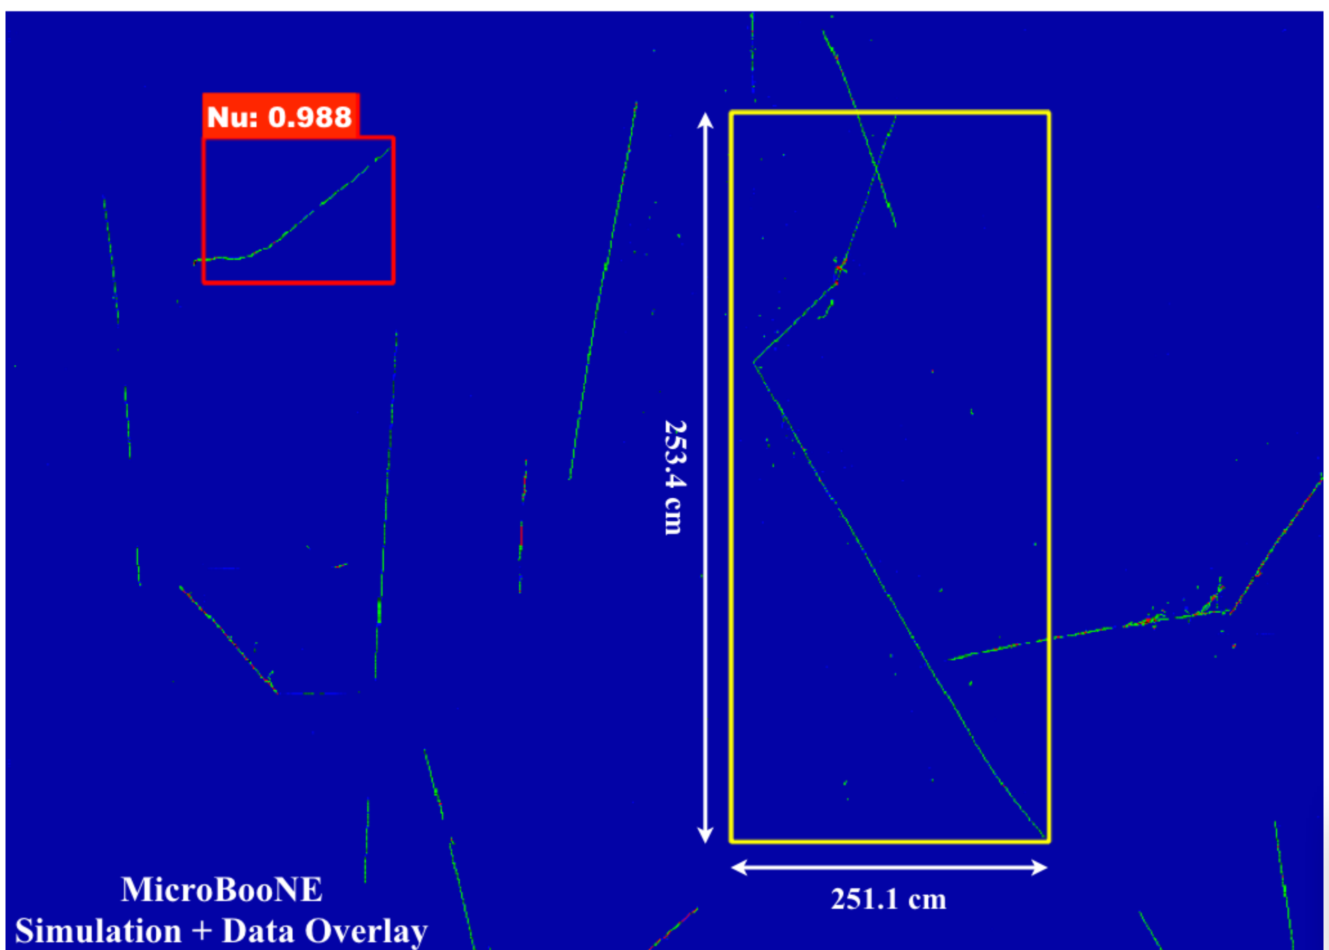

Explored challenges GPU performance vs downsampling effects for large LAr-TPCs

Convolutional Neural Networks Applied to Neutrino Events in a Liquid Argon Time Projection Chamber

MicroBooNE collaboration: R. Acciarri, C. Adams, R. An, J. Asaadi, M. Auger, L. Bagby, B. Baller, G. Barr, M. Bass,

VERTEX FINDING WITH MACHINE LEARNING

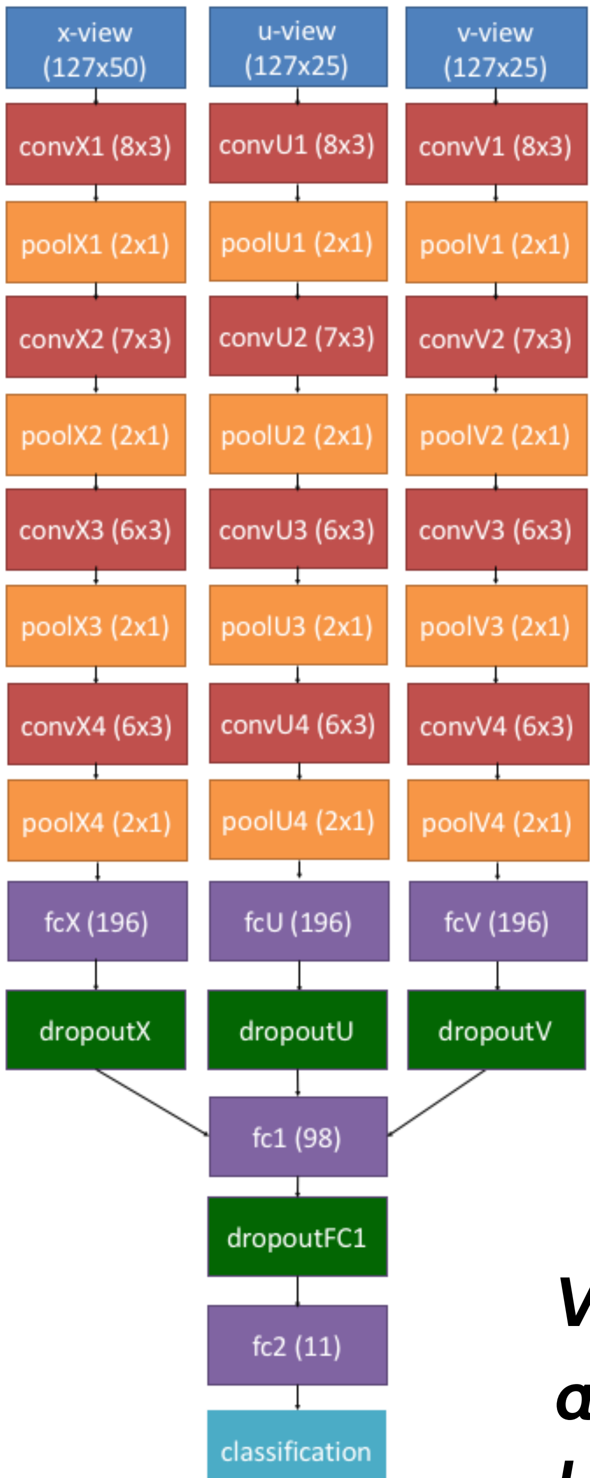

MINERvA uses a CNN with 3 prongs in order to combine information from the X V & U views of the event.

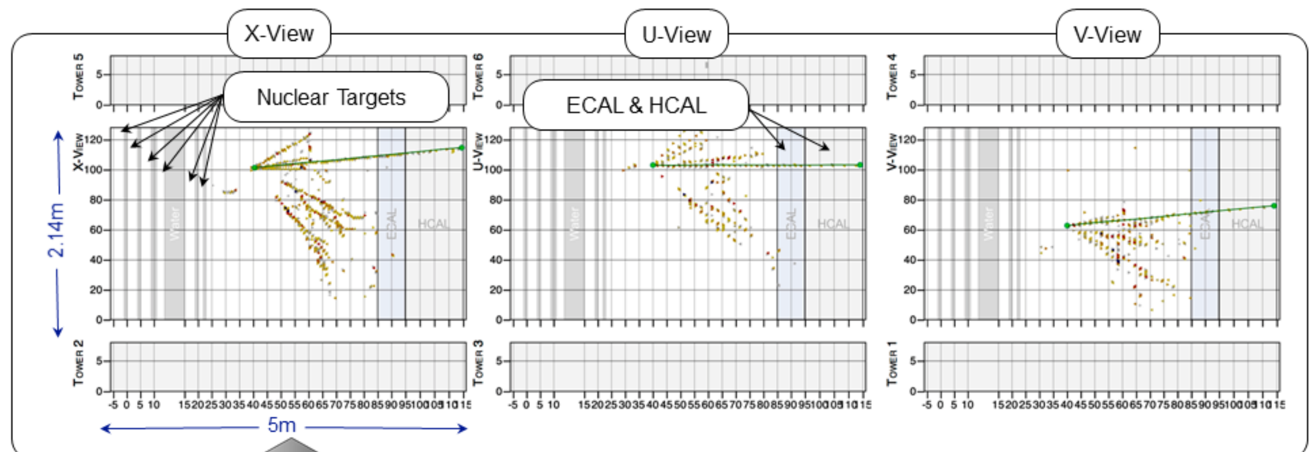

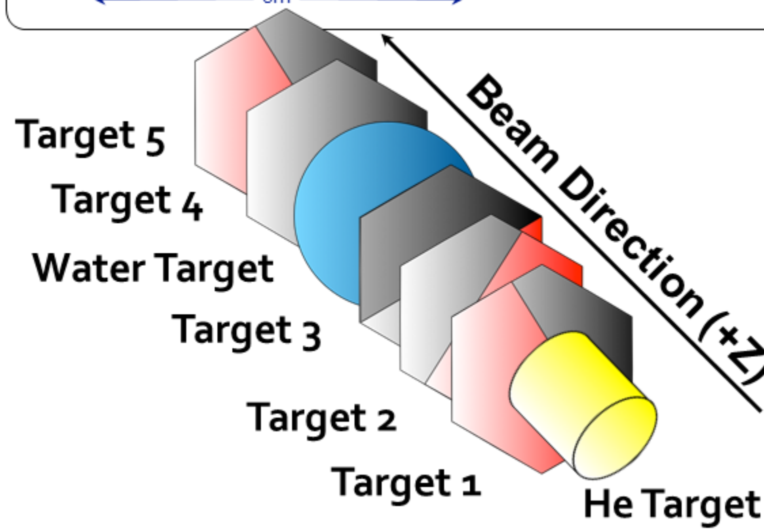

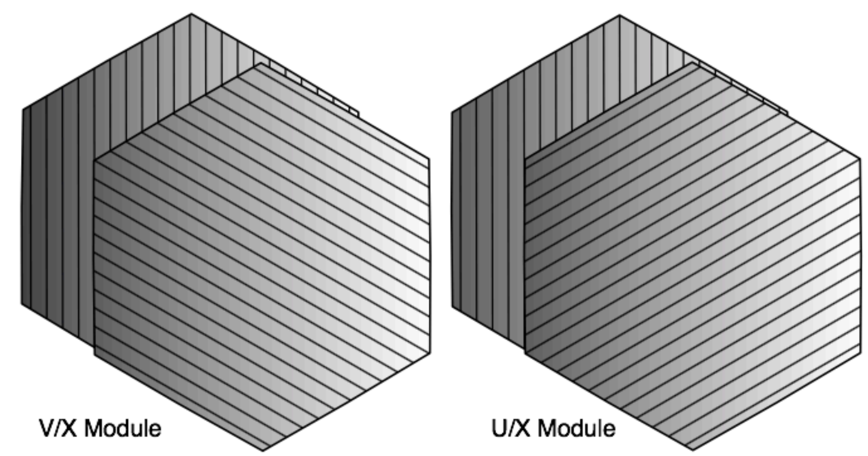

Varying network parameters they accomplish 94% accuracy for vertex location in the Z direction.

# convLayers	Kernel Sizes ($\{h\} \times w$)	Accuracy
Three	$\{6,6,3\} \times 3$	93.58%
Four	$\{\mathbf{8,8,7,6}\} \times \mathbf{3}$	**94.09%**
Five	$\{8,7,7,3,3\} \times 3$	93.55%

neutrino
network
cosmic
net
machine
vision
reconstruction
HEP
quantum
training
dark
algorithm
layer
neural
recursion
CNN
classifier
Science
Fermilab
NOvA
framework
sigmoid
regression
TMVA
bdt
AI
interaction
neutrinos
REL-U
IML
evaluation
MicroBooNE
BDT
LHC
background
amet

CHALLENGES

Performance and Robustness

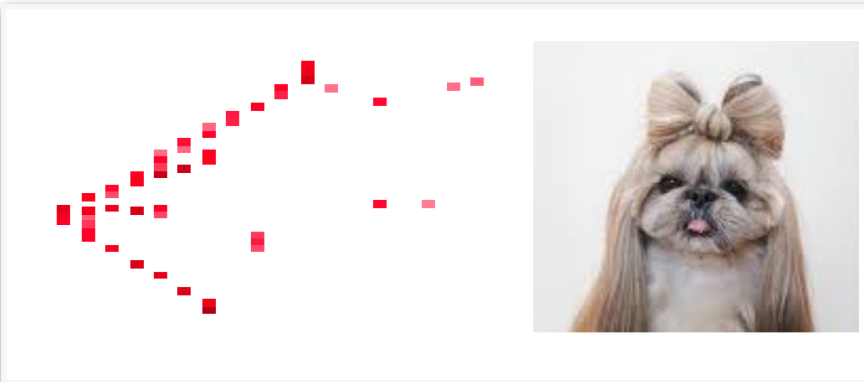

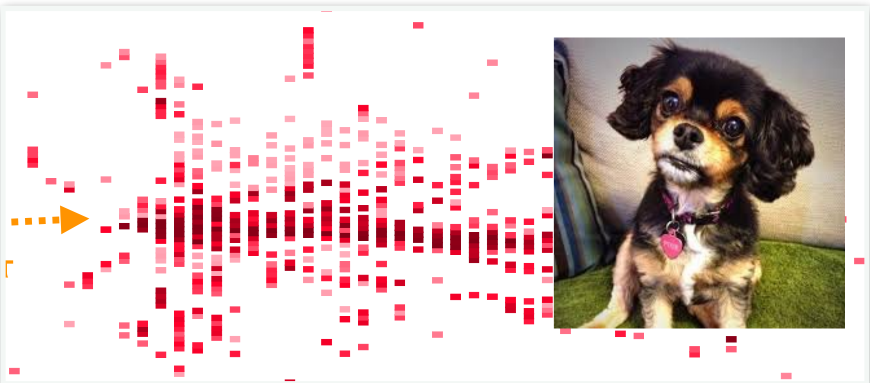

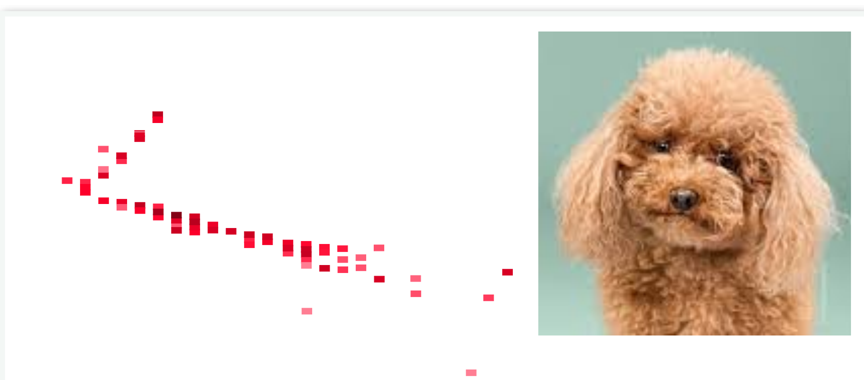

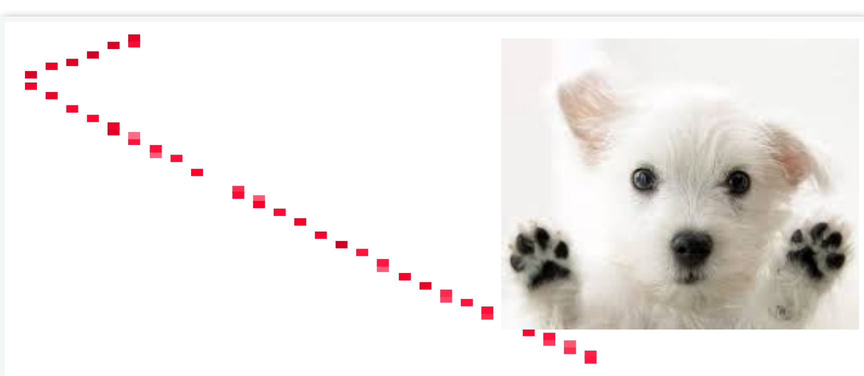

Typical issue is how to show robustness in data.

- *Data driven tests*
- *Training sample composition (to minimize biases which you know of a priori)*

and rough performance…

- *Overall accuracy*
- *Behavior of loss functions, etc*
- *Systematic uncertainties*

How do we find the biases we have introduced in our training?

Ensuring dependencies on the physics

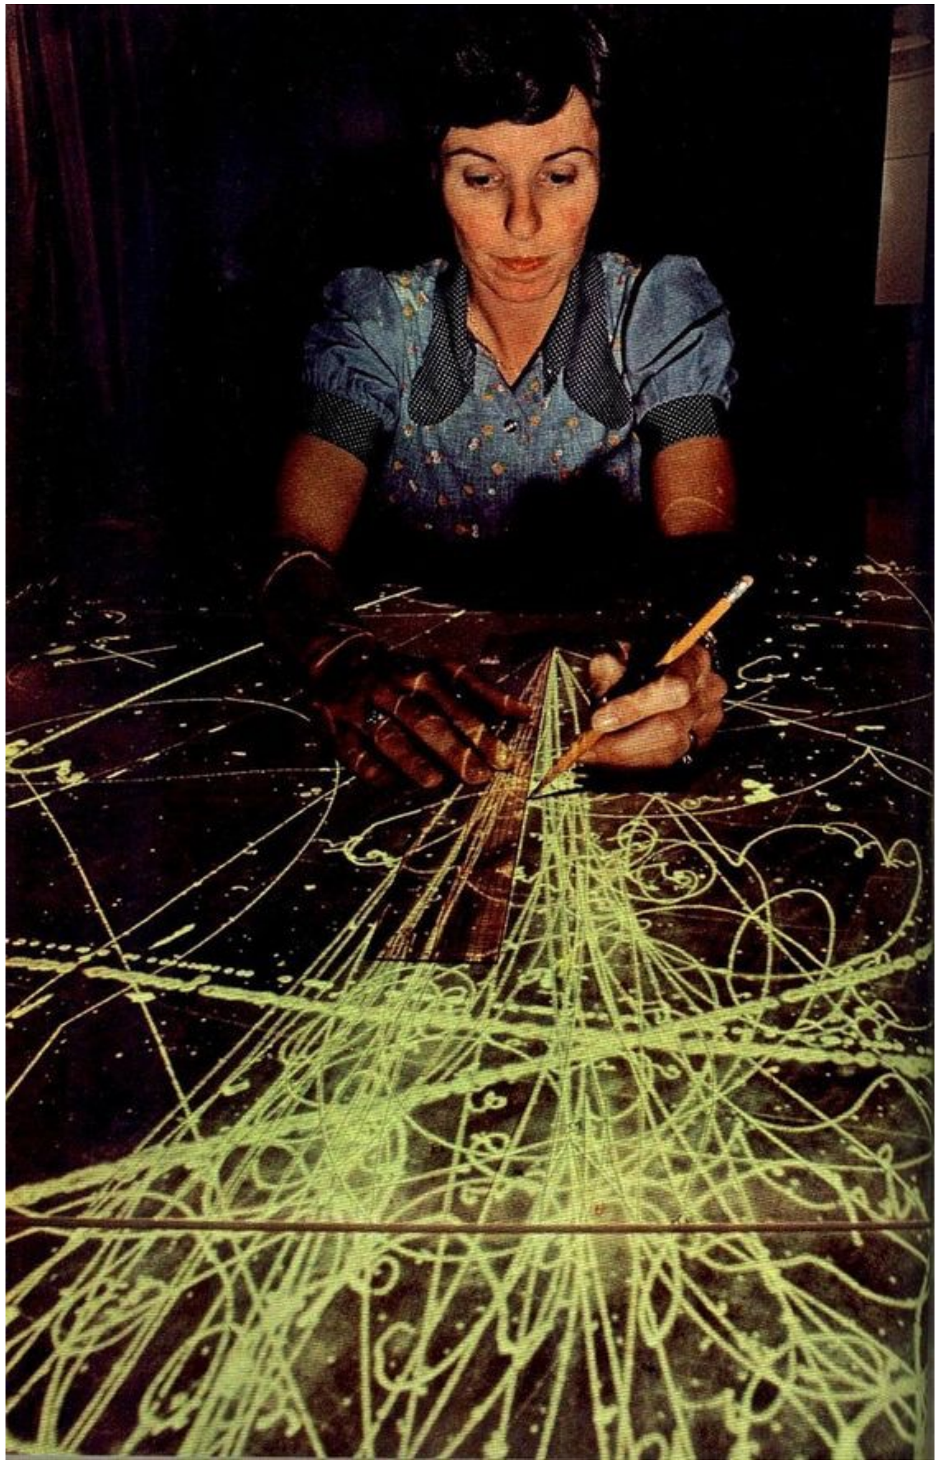

She would know this is not what dogs look like in nature.

How can we make sure these algorithms incorporate the physics that we know?

Can we develop tools to universally optimize (NOT TUNE) for the physics we understand?


algorithm
framework
quantum
reconstruction
HEP
CNN
Science
Fermilab
training
dark
vision
neural
layer
recursion
classifier
NOvA
algorithm
cosmic
net
machine
TMVA
MicroBooNE
framework
sigmoid
neutrino
network
regression
bdt
AI
evaluation
BDT
IML
LHC
background
REL-U
neutrinos
interaction
amet
reconstruction
neural
training
HEP
IML
neutrinos
evaluation
framework
Fermilab
CNN
neutrinos
machine

RESOURCES

EXISTING ML COMMUNITY EFFORTS

CERN openlab

http://openlab.cern/

AMVA 4 NEW PHYSICS

https://amva4newphysics.wordpress.com/

IML

https://iml.web.cern.ch/

Fermilab
50 Years of Discovery
Machine Learning

http://machinelearning.fnal.gov/

DS@HEP

INSIGHTS

FERMILAB MACHINE LEARNING GROUP

Community of analyzers at Fermilab with an interest in ML applications

http://machinelearning.fnal.gov/

Next meeting is this Friday:

June 9 at 10:30 AM
One West

TOPIC: CNN Applications for HEP

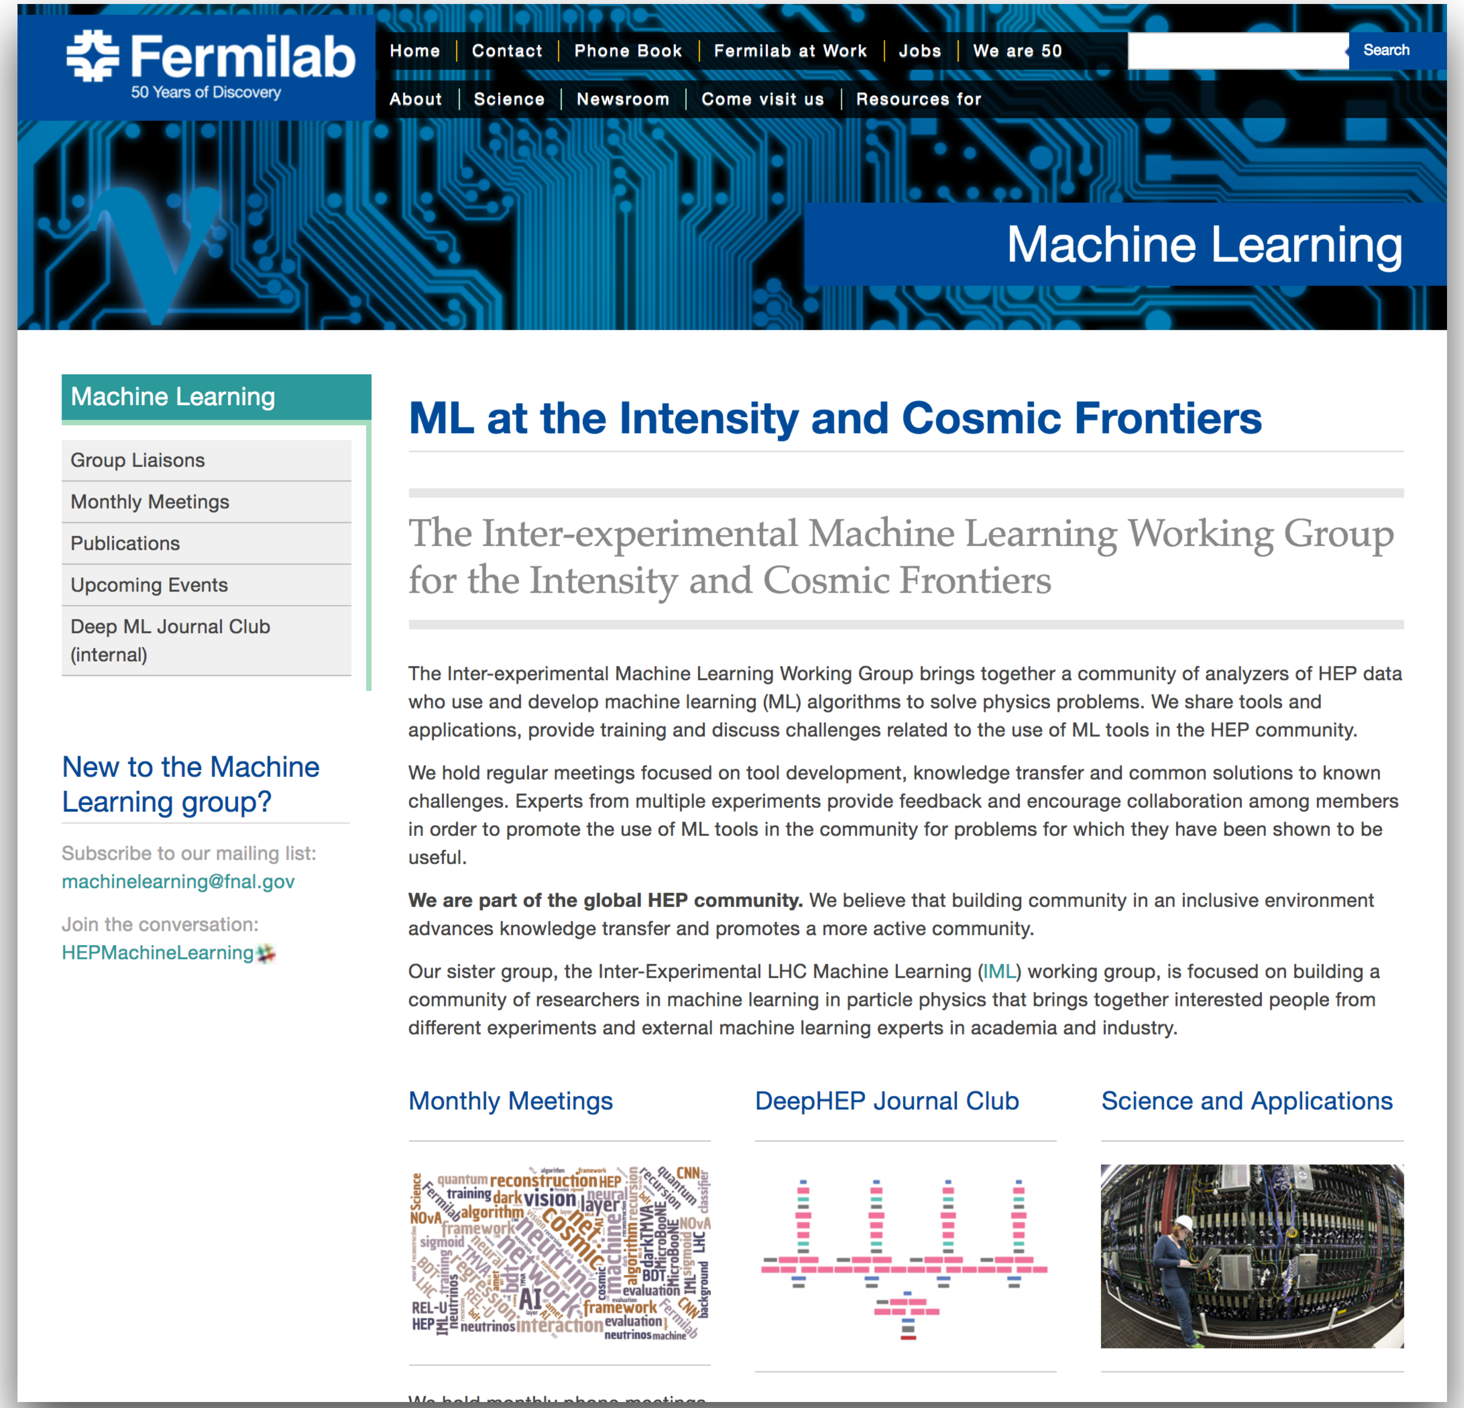

MORE ML APPLICATIONS THIS WEEK!

Electron Neutrino Reconstruction in MicroBooNE Using Deep Learning Technique	*Mr. Victor GENTY*	
One West, Fermilab, Wilson Hall		10:00 - 10:15
Search For Sterile Neutrinos At The NOvA Near Detector	*Mr. Siva Prasad KASETTI*	
One West, Fermilab, Wilson Hall		11:45 - 12:00
NOvA Short-Baseline Tau-Neutrino Appearance Search	*Mr. Rijeesh KELOTH*	
One West, Fermilab, Wilson Hall		12:00 - 12:15
Progress of the Measurement of the Electron Neutrino Charged-current Inclusive Cross Section in NOvA	*Matthew JUDAH*	
One West, Fermilab, Wilson Hall		17:45 - 18:00
Progress of the Charged Pion Semi-Inclusive Neutrino Charged-Current Cross Section in NOvA	*Ms. Jyoti TRIPATHI et al.*	
One West, Fermilab, Wilson Hall		18:00 - 18:15
AstroEncoder: Applications of Deep Learning to Cosmological Data	*Dr. Brian NORD*	
One West, Fermilab, Wilson Hall		14:00 - 14:15
Deep Learning for Hidden Signals—Enabling Real-time Multimessenger Astrophysics	*Mr. Daniel GEORGE*	
One West, Fermilab, Wilson Hall		14:45 - 15:00

Results From the Joint Fit to ν_e Appearance and ν_{μ} Disappearance in NOvA	*Ms. Shiqi YU*	
One West, Fermilab, Wilson Hall		11:00 - 11:15
Sterile neutrino search in the NOvA Far Detector.	*Mr. Sijith EDAYATH*	
One West, Fermilab, Wilson Hall		11:15 - 11:30

“Science doesn't have to be a zero-sum game. The key is to use whatever influence you do have to help your peers, and to trust that your peers will do the same.”

WORLD VIEW *A personal take on events*

J. TREGONING

No researcher is too junior to fix science

If young scientists plan to advance their careers before setting the system right, nothing will change, warns **John Tregoning**.

Thousands of researchers took to the streets last month to march for science. It is time to channel this energy into shaping scientific culture.

We all love to complain how the system for doing science thwarts ideal practice. Researchers reap more rewards for publishing flashy papers than for doing solid work, and the two do not always align. Everyone ends up chasing trends and asking the same questions. Broad, multidisciplinary research might achieve more in terms of advancing science, but it is harder to publish and finance. We end up sticking to the narrow path towards prestigious papers and big grants at the expense of worthier endeavours.

Why don't we change things? After all, science is uniquely self-regulating. The people who hire scientists are scientists, the people who allocate funding are scientists, and the people who decide what gets published are scientists. The tool we hold in highest regard is peer review: we are judge, jury and executioner.

One reason for stasis is that we scientists value consistency. The scientific process requires that variables be controlled as tightly as possible, even those that are unlikely to have any impact on an experiment. I know people who won't change the order in which they use pipette tips; they are unlikely to change scientific practice more broadly.

Another reason is that we are too busy just getting on in this system to pause to fix its flaws. Urgent grant submissions and experimental time points — tasks that reward the individual and have strict deadlines — will always win against some important but nebulous effort for the common good. It can feel as if those who spend their time on anything but their own projects and papers will find themselves scooped of the recognition required to win funding and resources.

Worst of all is the sad reality that those who most feel the need for change have the least power to effect it. The best time to fix the system, we tell ourselves, is after we have gained influence. If a PhD student shouts in frustration, are things going to change, or will she or he just be marginalized as a rabble-rouser?

This pernicious inertia persists at every rung of the career ladder — the higher scientists rise, the smaller seem the problems of those at the level below. Gaining a tenured post puts researchers in a position to make change, yet insulates them from much of what needs changing. The principal investigator tells the postdoc that finding a permanent position is easy compared to the angst of getting a grant. The postdoc tells the PhD student that defending a thesis is easy compared to the angst of finding a permanent position.

Evolutionary theory suggests a potential way out: reciprocal altruism. Science doesn't have to be a zero-sum game. The key is to use whatever influence you do have to help your peers, and to trust that your peers will do the same.

I have reaped the benefits of this approach. One simple example was relinquishing a key authorship position on a paper to maintain a productive collaboration. At the time, I felt that I was losing out by not fighting hard enough in the struggle for credit. But the small sacrifice paid off. I continued to work with my co-authors, and they invited me to join them in writing what turned out to be a successful grant application. The immediate reward of prime authorship would have been less beneficial in the long run.

More broadly, as an early-career principal investigator, I have sought out a group of like-minded colleagues. We consciously try to be less self-centred and to support each other. In practice, this comes down to small things that even those with pipetting rituals can handle: we read each other's drafts, accept our fair share of committee posts so that no one has an undue burden, and forward on relevant grant announcements. We each try to work a bit more towards a collective good: I happen to be enthusiastic about identifying broken stuff that everyone else ignores (burnt-out lights, squeaky doors, blocked sinks) and seeing that they get repaired. Other colleagues run seminar series, take the lead in teaching, interface between animal-care facilities and researchers, or manage the labs that require special biosafety precautions.

DON'T WAIT ON SENIOR COLLEAGUES, AND DEFINITELY DON'T WAIT UNTIL YOU BECOME ONE.

Reciprocal altruism can work more widely: mentoring postdocs or connecting students with careers outside academia, for example.

Don't wait on your senior colleagues, and definitely don't wait until you become one. Build a network of like-minded people. Identify something that doesn't work and fix it. It can be as small as a leaky tap or as big as peer review. Idealism can be catching.

Science will always be competitive, but too narrow a focus on your own advancement may come back to bite you. Academic promotions and appointments to senior positions require recommendations from colleagues. I'm sure I'm not the only one who has heard of ambitious acquaintances not being considered for promotion because they have stabbed too many people in the back.

Let's strive instead to stand together. One science historian called last month's science march unprecedented in its scale and breadth. That energy and optimism need not dissipate — it should be funnelled into making the system function better. The pay-off might not be immediate, but let's play the long game so that all can win. ■

John Tregoning *is a senior lecturer at Imperial College London, where he studies the immune response to viral infections. He blogs at http://drtregoning.blogspot.co.uk*
e-mail: john.tregoning@imperial.ac.uk

4 MAY 2017 | VOL 545 | NATURE | 7


quantum
reconstruction
HEP
CNN
classifier
Science
Fermilab
training
dark
vision
neural
layer
recursion
algorithm
NOvA
framework
cosmic
net
machine
neutrino
network
sigmoid
TMVA
regression
BDT
LHC
MicroBooNE
IML
background
evaluation
AI
REL-U
neutrinos
interaction
amet
bdt
reconstruction
neural
recursion
layer
machine

THANK YOU!