

# Machine Learning, Datascience and Neutrino Physics at Argonne's Leadership Computing Facility



Corey Adams

Data Science Group

Argonne Leadership Computing Facility

*corey.adams@anl.gov*

\* With content from the Datascience  
Team at ALCF, including Venkat  
Viswanath, Rick Zamora, Huihuo Zheng

# Outline

- **Argonne Leadership Computing Facility**
- **Datascience (And machine learning) at ALCF**
- **Scaling Machine Learning for High Performance Computing**
- **Deep Learning for Neutrino Physics**
- **How to get HPC resources at ALCF – we want you to use these resources!**

# ALCF




# Argonne Leadership Computing Facility



The Argonne Leadership Computing Facility provides world-class computing resources to the scientific community.

- Users pursue scientific challenges
- Resources fully dedicated to open science
- In-house experts to help maximize results



## Theta Intel/Cray


4,392 nodes  
281,088 cores  
69 TiB MCDRAM  
824 TiB DDR4  
549 TB SSD  
Peak flop rate: 11.69 PF



**Mira** IBM BG/Q  
49,152 nodes  
786,432 cores  
786 TB RAM  
Peak flop rate: 10 PF

ALCF offers different pipelines based on your computational readiness. Apply to the allocation program that fits your needs.

<https://www.alcf.anl.gov>



# Argonne's path to Exascale is critical to our nation's scientific leadership



**ALCF4**

2007

2013

2017

2021

557 TeraFLOPS

10 PetaFLOPS

11 PetaFLOPS





# Computing Resources

## Focus on Theta




### Theta (Cray XC40)

- **11.69 PF** Peak performance
- **4,392 nodes** (281,088 cores)
  - 2nd Generation Intel® Xeon Phi™ Processor (Knights Landing)
- **843.264 TB DDR4** and **70.272 TB MCDRAM** total memory
- **128 GB SSD** on each node
- Cray Aries high speed interconnect in dragonfly topology
- **10 PB** Lustre file system, **~200 GB/s** throughput

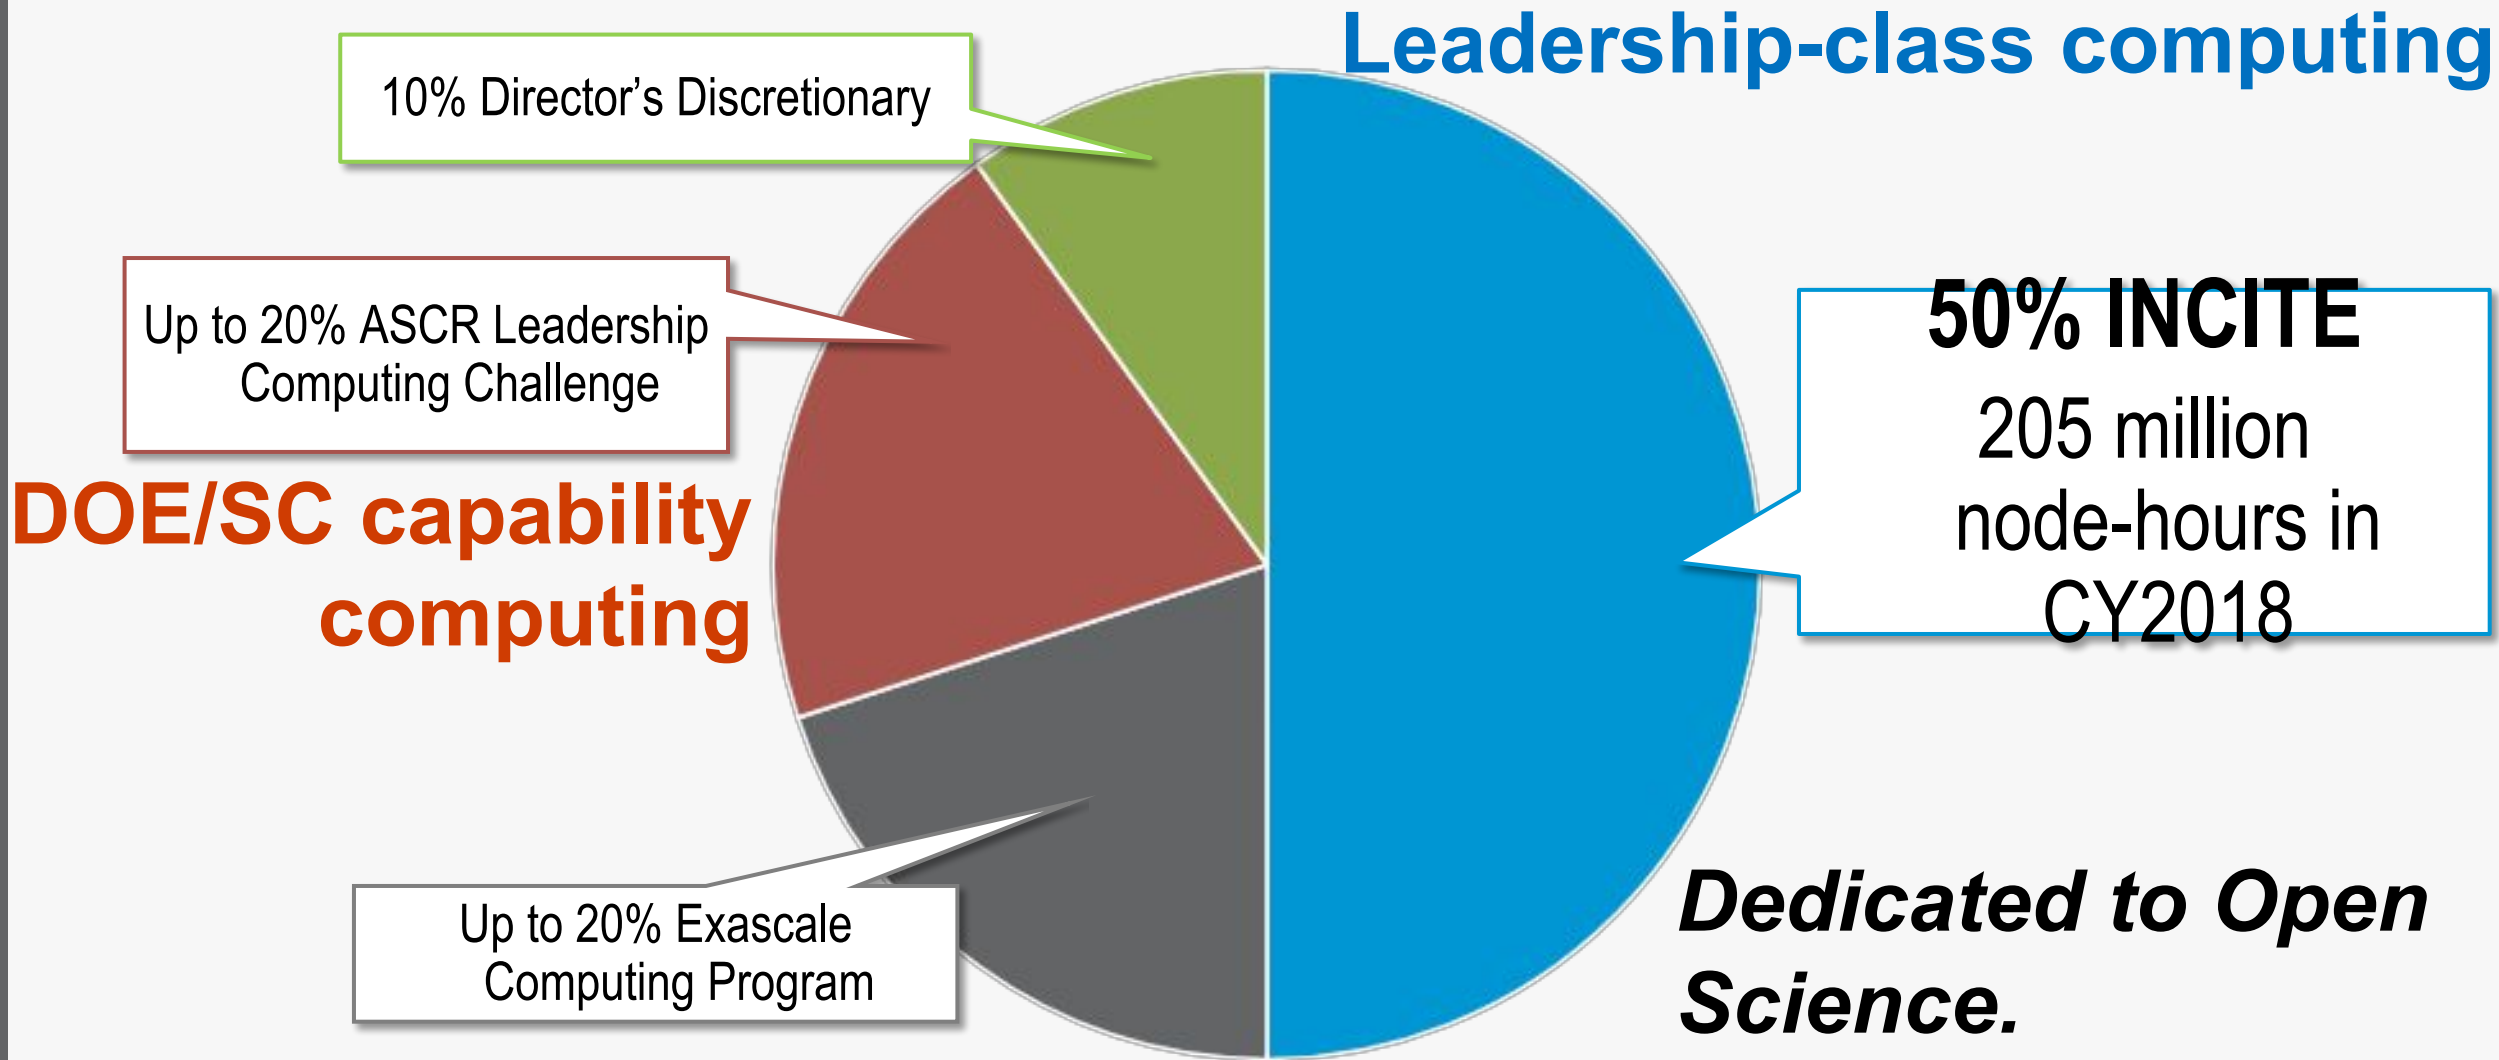


**Also available:** Mira, as well as other visualization and analytics clusters, and storage systems



# Who Uses ALCF?

## Allocation Distributions



# ALCF Data Science Program



# ALCF Data Science Program (ADSP) Overview

- Big Data science problems that require the leadership scale and performance
- Span computational, experimental and observational sciences
- Focus on data science techniques including but not limited to statistics, machine learning, deep learning, UQ, image processing, graph analytics, complex and interactive workflows
- Two-year proposal period and will be renewed annually. Proposals will target science and software technology scaling for data science
- The program started in 2016 and now in the 3rd year
- Yearly call for proposal.

Next deadline – ~June 2019

- <https://www.alcf.anl.gov/alcf-data-science-program>



# ALCF Data Science Program (ADSP)

## Targets Data & Learning Pillars



### Data

- Experimental/observational data
  - Image analysis
  - Multidimensional structure discovery
- Complex and interactive workflows
- On-demand HPC
- Persistent data techniques
  - Object store
  - Databases
- Streaming/real-time data
- Uncertainty quantification
- Statistical methods
- Graph analytics

### Learning

- Deep learning
- Machine learning steering simulations
  - Parameter scans
  - Materials design
  - Observational signatures
- Data-driven models and refinement for science using ML/DL
- Hyperparameter optimization
- Pattern recognition
- Bridging gaps in theory

**Big Data**

# datascience@alcf.anl.gov



Corey Adams



Prasanna  
Balaprakash



Taylor Childers




Murali Emani




Elise  
Jennings



Xiao-Yong Jin



Murat Keceli




Bethany Lusch



Alvaro Vazquez



Adrian Pope



Misha Salim




William Scullin



Ganesh  
Sivaraman



Tom Uram



Antonio Villarreal



Venkat Vishwanath

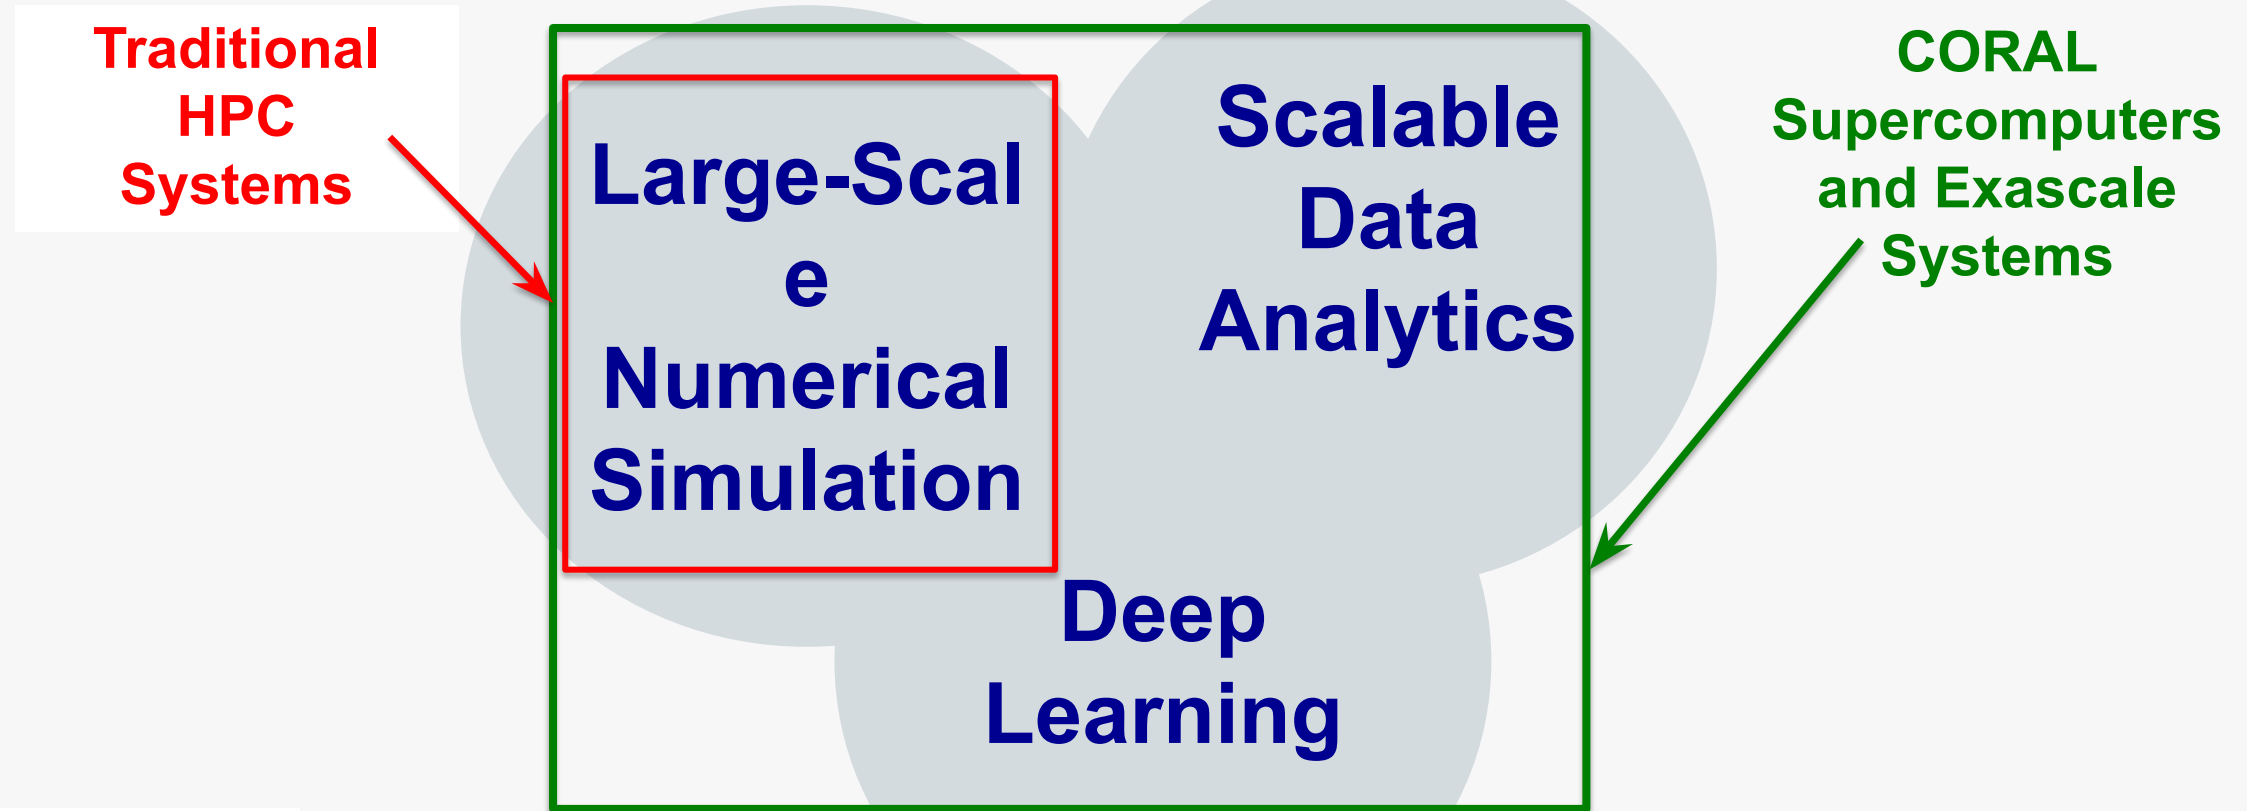


Richard Zamora



Huihuo Zheng

# Integration of Simulation, Data Analytics and Machine Learning on supercomputers



# Aurora 2021 - The first US Exascale System



Architecture supports three ways of computing

- Large-scale Simulation (PDEs, traditional HPC)
- Data Intensive Applications (scalable science pipelines)
- Deep Learning and Emerging Science AI (training and inferencing)



# Machine Learning and High Performance Computing

# Machine Learning and HPC


## Accelerate and improve an application's:

---

**Time to Solution (Training)** – with scalable learning techniques, you can process more images per second, reduce the time per epoch, and reach a trained network faster.

**Quality of Solution** – with more compute resources available, you can perform hyperparameter searches to optimize network designs and training schemes. With powerful accelerators, you can train bigger and more computationally intense networks.

**Inference Throughput** – with high bandwidth IO, it is easy to scale up the throughput of inference techniques for deep learning.



High Performance Computing can improve all aspects of training and inference in machine learning.

# Machine Learning and HPC – KNL Nodes


## Single Node Performance Matters!

Running a model on an HPC node is often not like a standard GPU – **many configuration parameters matter, not all models have the same optimum parameters.**

Intel KNL != Nvidia GPU, but KNL can be powerful.

**intra\_op\_parallelism\_threads**: Nodes that can use multiple threads to parallelize their execution will schedule the individual pieces into this pool.

**inter\_op\_parallelism\_threads**: All ready nodes are scheduled in this pool.



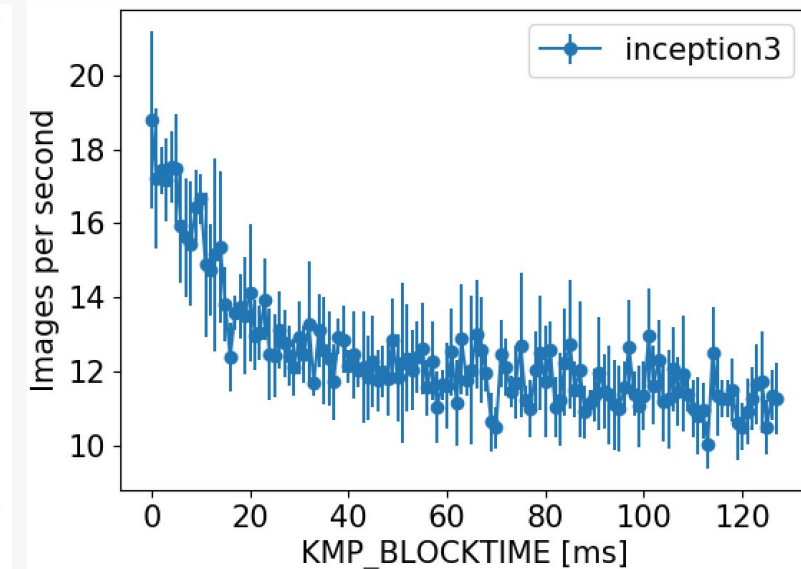

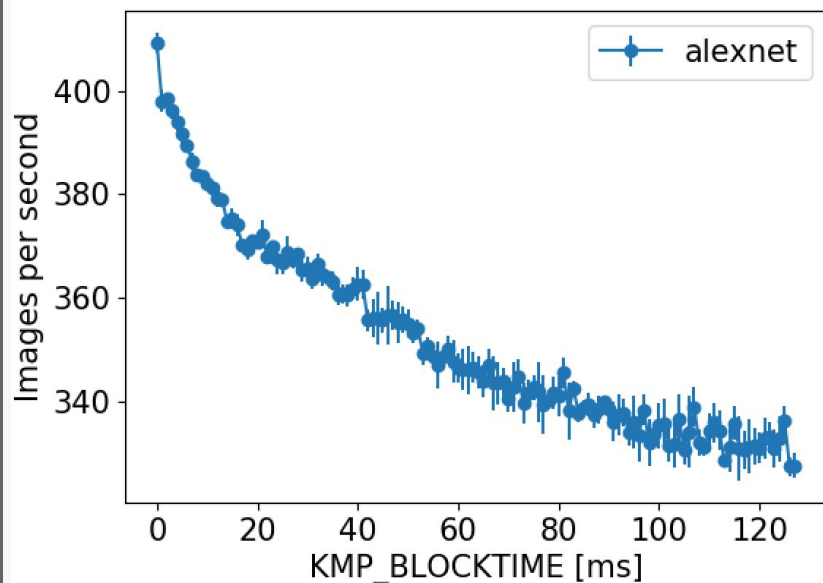
```
config = tf.ConfigProto()  
config.intra_op_parallelism_threads = num_intra_threads  
config.inter_op_parallelism_threads = num_inter_threads  
tf.Session(config=config)
```

<https://www.tensorflow.org/guide/performance/overview>



# Machine Learning and HPC – KNL Nodes

## Affinity can play a large role



`KMP_BLOCKTIME=0` is optimal for KNL Nodes

`KMP_AFFINITY=granularity=fine,verbose,compact,1,0`

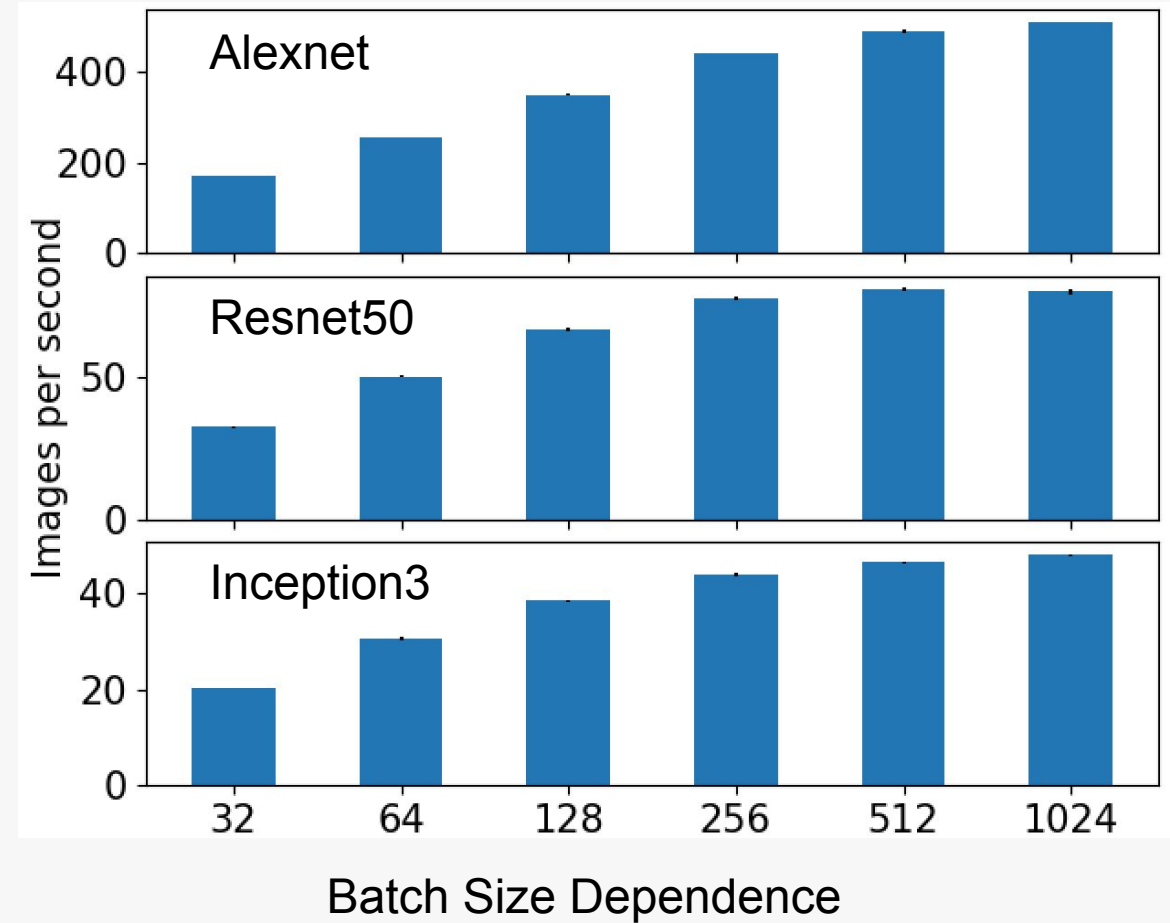
[Intel Affinity Guidelines](#)

# Machine Learning and HPC – KNL Nodes

## Batch Size can be important

Running a model on an HPC node is often not like a standard GPU – **many configuration parameters matter, not all models have the same optimum parameters.**

Bigger batch size often yields more images/second throughput (though not always), but the downside is always more seconds/global step at a large batch size.



# Distributed Learning

Machine learning is a very important workflow for current and future supercomputing systems.

How can you accelerate learning with more computing power?




Image from Uber's Horovod: <https://eng.uber.com/horovod/>

# What is Distributed Learning?

## A technique to accelerate training

**The backpropagation algorithm is unchanged at its heart.**

---

**Data Parallel learning** – with  $N$  nodes, replicate your model on each node. After the forward and backward computations, **average** the gradients across all nodes and use the averaged gradients to update the weights. Conceptually, **this multiplies the minibatch size by  $N$ .**

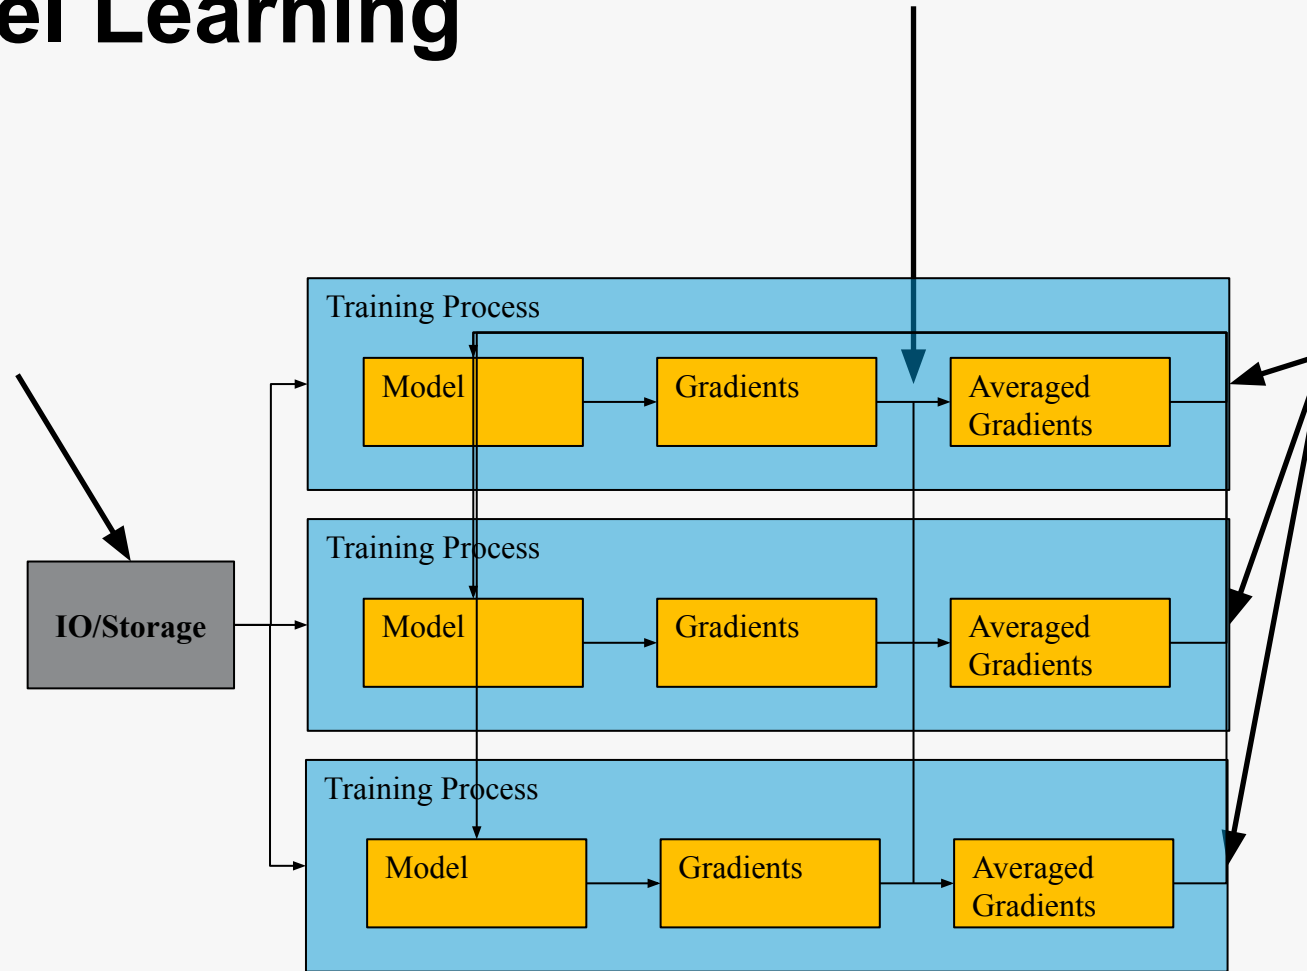
**Model Parallel Learning** – for models that don't fit on a single node, you can divide a single model across multiple locations. The design of distributing a model is not trivial, but tools are emerging.

**Both (“Mesh” training)** – Using  $n$  nodes for a single model, and  $N = k*n$  nodes for distributed training, you can achieve accelerated training of extremely large or expensive models.

# Data Parallel Learning

All nodes communicate to average gradients.

Each Model gets unique input data and performs calculations independently.



Each Node gets it's own copy of the model.

Image from Uber's Horovod: <https://eng.uber.com/horovod/>

# Data Parallel Learning

## Scaling Challenges

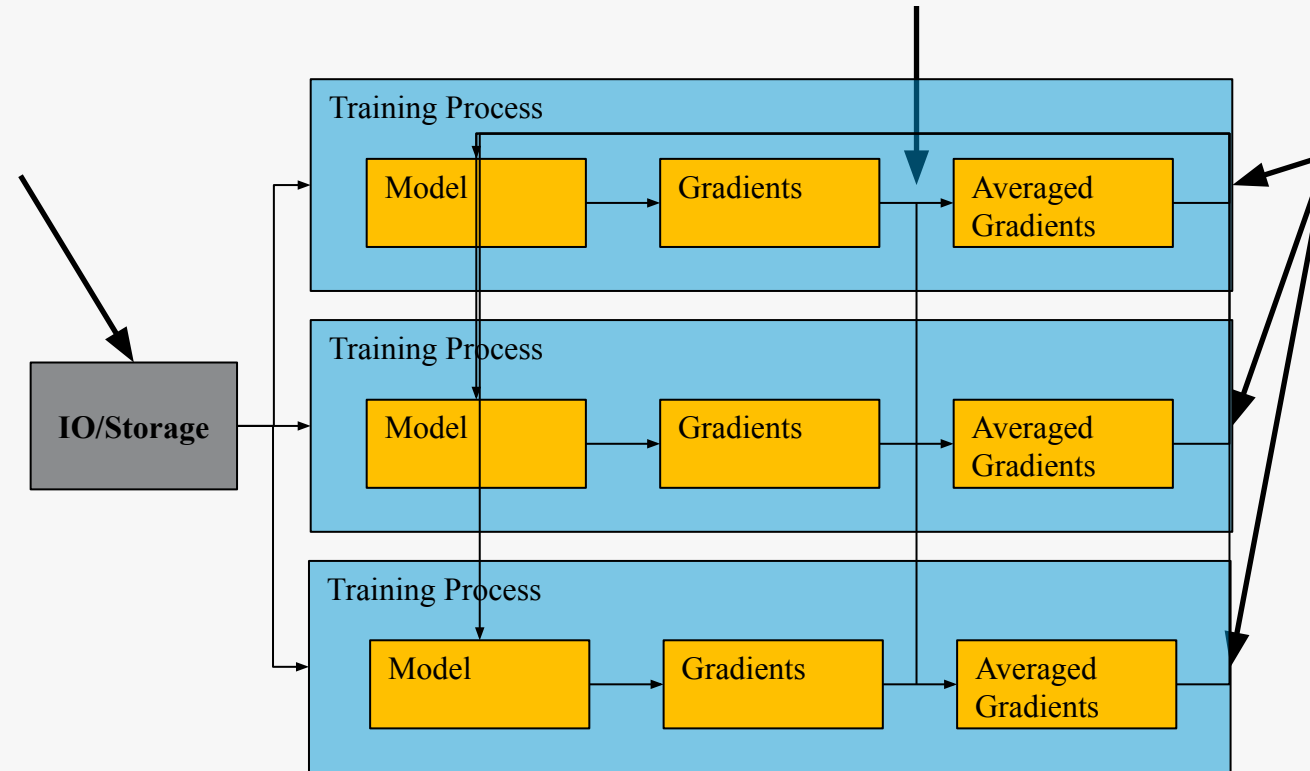
Each Model gets unique input data and performs calculations independently.

IO requires organization to ensure unique batches.

IO contention with many nodes requires parallel IO solutions

All nodes communicate to average gradients.

Computation stalls during communication: keeping the communication to computation ratio small is important for scaling.



Each Node gets it's own copy of the model.

Initialization must be identical or synchronized, and checkpointing/summary information must be managed with just one node.

Image from Uber's Horovod: <https://eng.uber.com/horovod/>

# Data Parallel Learning




Image from Uber's Horovod: <https://eng.uber.com/horovod/>

# Data Parallel Learning

## Horovod

### The simplest technique for data parallel learning

1. Initialize horovod ( `hvd.init()` ).
2. Wrap the optimizer in `hvd.DistributedOptimizer`.
  1. This uses the underlying optimizer for gradient calculations, and performs an averaging of all gradients before updating.
  2. Can adjust the learning rate to account for a bigger batch size.
3. Initialize the networks identically, or broadcast one network's weights to all others.
4. Ensure snapshots and summaries are only produced by one rank.

**Horovod focuses on handling collective communication so you don't have to, but let's you use all of the tools of your favorite framework. Compatible with `mpi4py`.**



Horovod is an open source data parallel training software compatible with many common deep learning frameworks.

[Meet Horovod](#)  
[Github](#)



# Horovod Example Code

## Tensorflow

```
import tensorflow as tf
import horovod.tensorflow as hvd
layers = tf.contrib.layers
learn = tf.contrib.learn
def main():
    # Horovod: initialize Horovod.
    hvd.init()
    # Download and load MNIST dataset.
    mnist = learn.datasets.mnist.read_data_sets('MNIST-data-%d' % hvd.rank())
    # Horovod: adjust learning rate based on number of GPUs.
    opt = tf.train.RMSPropOptimizer(0.001 * hvd.size())
    # Horovod: add Horovod Distributed Optimizer
    opt = hvd.DistributedOptimizer(opt)
    hooks = [
        hvd.BroadcastGlobalVariablesHook(0),
        tf.train.StopAtStepHook(last_step=20000 // hvd.size()),
        tf.train.LoggingTensorHook(tensors={'step': global_step, 'loss': loss},
                                   every_n_iter=10),
    ]
    checkpoint_dir = './checkpoints' if hvd.rank() == 0 else None
    with tf.train.MonitoredTrainingSession(checkpoint_dir=checkpoint_dir,
                                           hooks=hooks,
                                           config=config) as mon_sess
```

# Horovod Example Code

## Pytorch

```
import torch.nn as nn
import horovod.torch as hvd
hvd.init()
train_dataset = datasets.MNIST('data-%d' % hvd.rank(), train=True, download=True,
                               transform=transforms.Compose([
                                   transforms.ToTensor(),
                                   transforms.Normalize((0.1307,), (0.3081,))
                               ]))
train_sampler = torch.utils.data.distributed.DistributedSampler(
    train_dataset, num_replicas=hvd.size(), rank=hvd.rank())
train_loader = torch.utils.data.DataLoader(
    train_dataset, batch_size=args.batch_size, sampler=train_sampler, **kwargs)
# Horovod: broadcast parameters.
hvd.broadcast_parameters(model.state_dict(), root_rank=0)
# Horovod: scale learning rate by the number of GPUs.
optimizer = optim.SGD(model.parameters(), lr=args.lr * hvd.size(),
                       momentum=args.momentum)
# Horovod: wrap optimizer with DistributedOptimizer.
optimizer = hvd.DistributedOptimizer(optimizer, named_parameters=model.named_parameters())
```

# Horovod Example Code

## Keras

```
import keras
import tensorflow as tf
import horovod.keras as hvd
# Horovod: initialize Horovod.
hvd.init()
# Horovod: adjust learning rate based on number of GPUs.
opt = keras.optimizers.Adadelta(1.0 * hvd.size())
# Horovod: add Horovod Distributed Optimizer.
opt = hvd.DistributedOptimizer(opt)
model.compile(loss=keras.losses.categorical_crossentropy,
              optimizer=opt,
              metrics=['accuracy'])
callbacks = [
    # Horovod: broadcast initial variable states from rank 0 to all other processes.
    hvd.callbacks.BroadcastGlobalVariablesCallback(0),
]
# Horovod: save checkpoints only on worker 0 to prevent other workers from corrupting them.
if hvd.rank() == 0:
    callbacks.append(keras.callbacks.ModelCheckpoint('./checkpoint-{epoch}.h5'))
model.fit(x_train, y_train, batch_size=batch_size,
        callbacks=callbacks,
        epochs=epochs,
        verbose=1, validation_data=(x_test, y_test))
```

# Effects of Distributed Learning

1. Increased Batch size means improved estimate of gradients.
  1. Scale by N nodes?  $\text{Sqrt}(N)$ ?
  2. Scale in a layerwise way? Layerwise Adaptive Rate Scaling (LARS)
2. Increased learning rate can require warm up iterations.
  1. Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour
3. Bigger minibatch means less iterations for the same number of epochs.
  1. May need to train for more epochs if another change is not made like boosting the learning rate.

# Mesh Learning

## Tensorflow Mesh

### When data-parallel isn't enough...

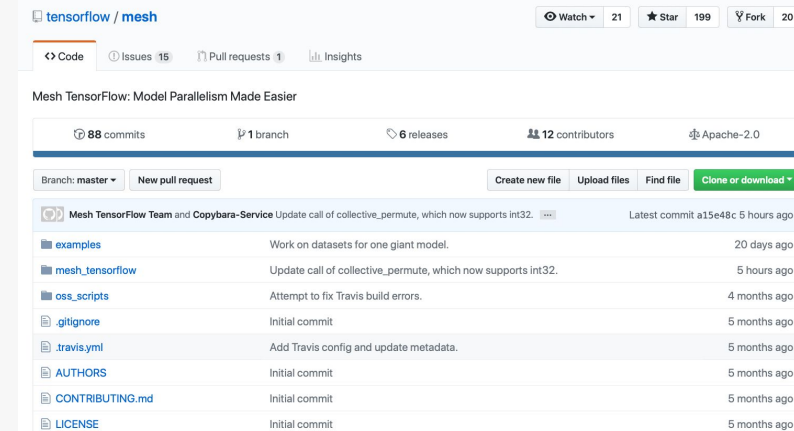
#### Why might you need a Mesh?

- Memory limitations due to network size (number of parameters)
- Memory limitations due to input size (massive images, etc)

#### Mesh Scaling is not trivial:

- Computations need to be distributed in an intelligent way to prevent idle nodes
- Communication needs to happen frequently during both the forward/backward pass
- Message passing organization details arise from forward/backward small-group communications and multi-group communications

**Expect mesh scaling to get easier over the next few years (or wait for bigger, more powerful nodes!)**



tensorflow / mesh

Watch 21 Star 199 Fork 20

Code Issues 15 Pull requests 1 Insights

Mesh TensorFlow: Model Parallelism Made Easier

88 commits 1 branch 6 releases 12 contributors Apache-2.0

Branch: master New pull request Create new file Upload files Find file Clone or download

File	Description	Latest commit
examples	Work on datasets for one giant model.	20 days ago
mesh_tensorflow	Update call of collective_permute, which now supports int32.	5 hours ago
oss_scripts	Attempt to fix Travis build errors.	4 months ago
.gitignore	Initial commit	5 months ago
.travis.yml	Add Travis config and update metadata.	5 months ago
AUTHORS	Initial commit	5 months ago
CONTRIBUTING.md	Initial commit	5 months ago
LICENSE	Initial commit	5 months ago

<https://github.com/tensorflow/mesh>

# Neutrino Physics @ ALCF

The background of the slide features a large, stylized letter 'A' composed of numerous thin, light blue lines. Scattered throughout the scene are many small, semi-transparent blue spheres, resembling particles or data points. The overall color palette is dark blue and black, creating a scientific and digital atmosphere.

# A Crash Course in Neutrino Physics

# Particle Physics - Building Blocks





**"Sewing the fabric of spacetime"**

**Everything** we know about is made from these particles



# Particle Interactions

mass →	≈2.3 MeV/c <sup>2</sup>	≈1.275 GeV/c <sup>2</sup>	≈173.07 GeV/c <sup>2</sup>	0	≈126 GeV/c <sup>2</sup>
charge →	2/3	2/3	2/3	0	0
spin →	1/2	1/2	1/2	1	0
	<b>u</b> up	<b>c</b> charm	<b>t</b> top	<b>g</b> gluon	<b>H</b> Higgs boson
<b>QUARKS</b>					
	≈4.8 MeV/c <sup>2</sup>	≈95 MeV/c <sup>2</sup>	≈4.18 GeV/c <sup>2</sup>	0	
	-1/3	-1/3	-1/3	0	
	1/2	1/2	1/2	1	
	<b>d</b> down	<b>s</b> strange	<b>b</b> bottom	<b>γ</b> photon	
	0.511 MeV/c <sup>2</sup>	105.7 MeV/c <sup>2</sup>	1.777 GeV/c <sup>2</sup>	91.2 GeV/c <sup>2</sup>	
	-1	-1	-1	0	
	1/2	1/2	1/2	1	
	<b>e</b> electron	<b>μ</b> muon	<b>τ</b> tau	<b>Z</b> Z boson	
<b>LEPTONS</b>					
	<2.2 eV/c <sup>2</sup>	<0.17 MeV/c <sup>2</sup>	<15.5 MeV/c <sup>2</sup>	80.4 GeV/c <sup>2</sup>	
	0	0	0	±1	
	1/2	1/2	1/2	1	
	<b>ν<sub>e</sub></b> electron neutrino	<b>ν<sub>μ</sub></b> muon neutrino	<b>ν<sub>τ</sub></b> tau neutrino	<b>W</b> W boson	
					<b>GAUGE BOSONS</b>



Particle Interactions are really just  
**mediating particle**

# Basic Rules

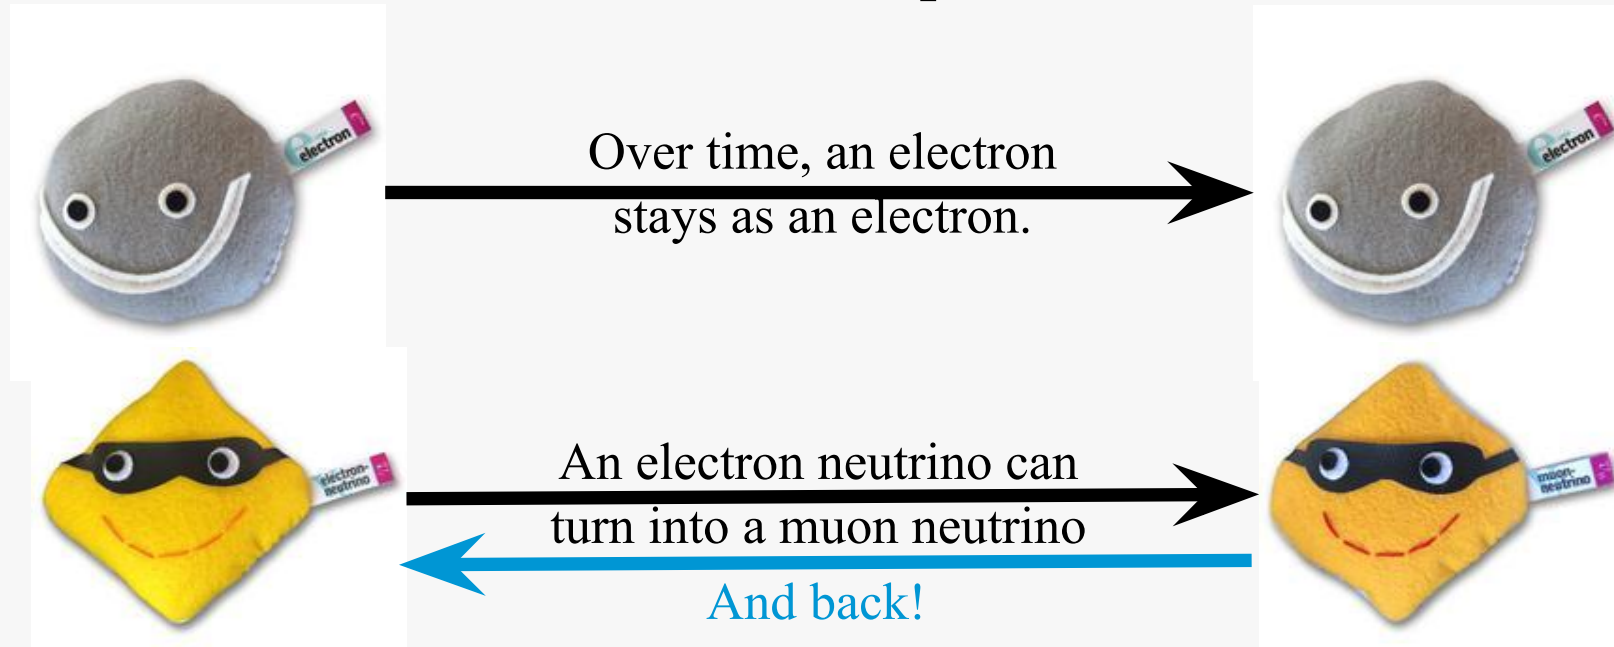
Fundamental Particles follow basic rules:

1. Anything not forbidden is allowed.
2. Things that are forbidden are forbidden by "conservation" laws.
  1. Traditional conservation laws: conservation of momentum, energy, charge, etc...
  2. Particle Physics conservation laws: conservation of "lepton" number, "baryon" number, "flavor", "color", ...

**Neutrinos are one of the biggest rule breakers in the Universe.**

# Neutrino Oscillations

Neutrinos are super weird!



Measuring the frequency and strength of the "oscillations" tells a LOT about neutrinos.

# Why Study Neutrinos?

Are neutrinos their own antiparticles?

Do neutrinos explain the matter/antimatter asymmetry in the Universe?


Are there more than 3 types of neutrinos?

Could new type of neutrinos explain dark matter?



Why is the neutrino mass so much less than other particles?

# Liquid Argon Time Projection Chambers


The modern neutrino detector of choice for high energy physics is the **Liquid Argon Time Projection Chamber**.




# Liquid Argon Time Projection Chambers



# “Projection




# “Projection” Chamber – angle 1








# “Projection” Chamber – angle 2



# 3 Projections of Same Objects



**Liquid argon time projection chambers are the detector-of-choice for high energy physics experiments**

1. Argoneut (Fermilab, Decommissioned)
2. Lariat (Fermilab Testbed, Running)
3. MicroBooNE (Fermilab Running)
4. SBND (Fermilab, Construction)
5. ICARUS (Fermilab/INFN, Assembly)
6. ProtoDUNE SP (CERN, Running)
7. ProtoDUNE DP (CERN, Assembly)
8. DUNE ND (Fermilab, Design) – **pixel readout**
9. DUNE FD (Sanford Lab, SD, Design)

Many other physics experiments that need fine-grained tracking and detailed calorimetry use similar time projection chamber technology.

Novel designs will replace the sense wires with pixel planes for intrinsically 3D imaging.

# Promising Technology for applications of Deep Learning and HPC

Large data size (  $O(10^7)$  pixels per image – for example, 1260 x 2048 pixels, 3 images = 1 Event)

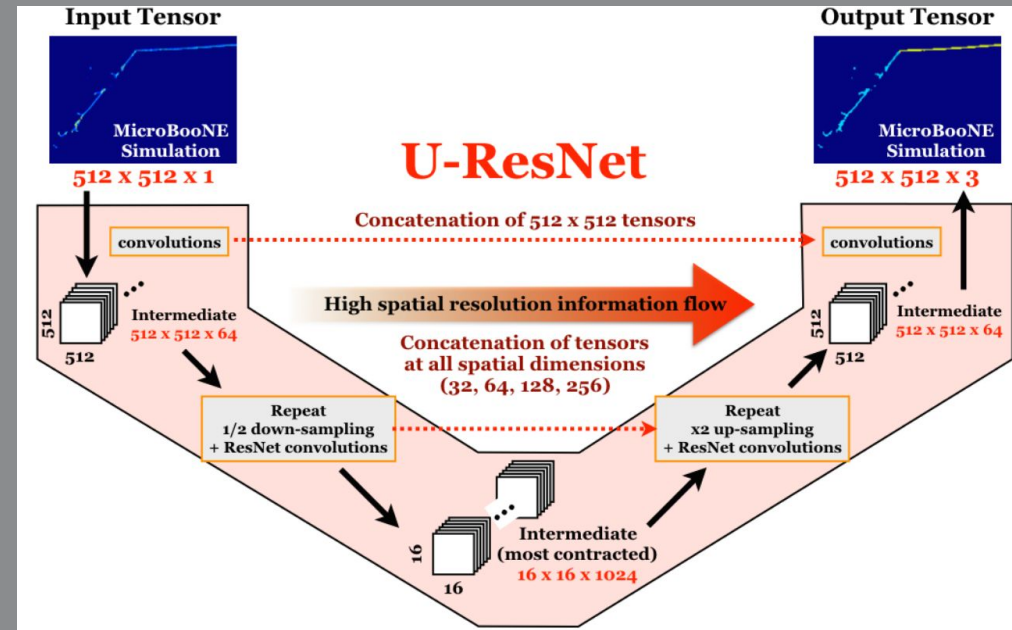
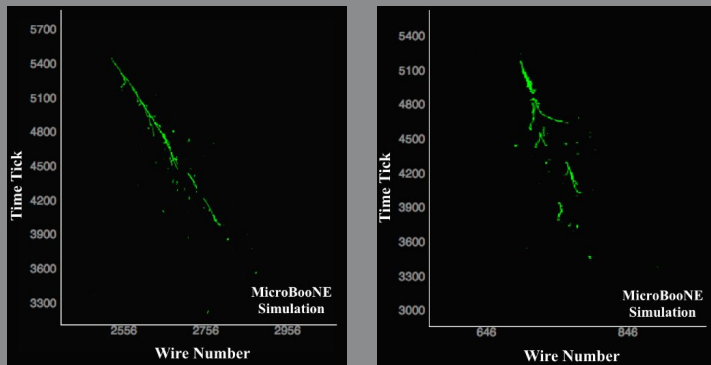
Huge datasets (millions of images of both simulated and real data, per experiment)

Multiple fundamental science applications from neutrino oscillations, proton decay searches, beyond-standard-model physics searches

Multiple deep learning techniques directly applicable

Electron vs Photon (Classification)

JINST 12, P03011 (2017).



A Deep Neural Network for Pixel-Level Electromagnetic Particle Identification in the MicroBooNE Liquid Argon Time Projection Chamber

# More Data, More Problems

Large scale, high resolution detectors don't fit into any successful paradigm of analysis.

For the MicroBooNE detector (Fermilab):  
Each event is  $\sim 40\text{MB}$  of raw data for  $\sim 5\text{ms}$  of data. 8 GB/s of data (**almost 30 TB per hour!**)



**Not feasible!** 

Need to develop a **fast, automated** pattern recognition.



# More Data, More Problems

Large scale, high resolution detectors don't fit into any successful paradigm of analysis




Complex data requires complex pattern recognition.

# Physicist vs. The Machine

Sample	Data	Simulation	Simulation	Simulation
Label	Physicist	Physicist	Simulation	Simulation
Prediction	U-ResNet	U-ResNet	U-ResNet	Physicist
ICPF mean	3.4	2.5	1.8	2.0
ICPF 90%	9.0	5.7	4.6	4.8
Shower	4.8	3.4	3.0	2.6
Track	2.7	2.4	2.2	2.9


ICPF = Incorrectly Labeled Pixel Fraction.

**No classical algorithm achieves this accuracy. Machine Learning out performs physicist hand labels.**





# Cosmic Tagging

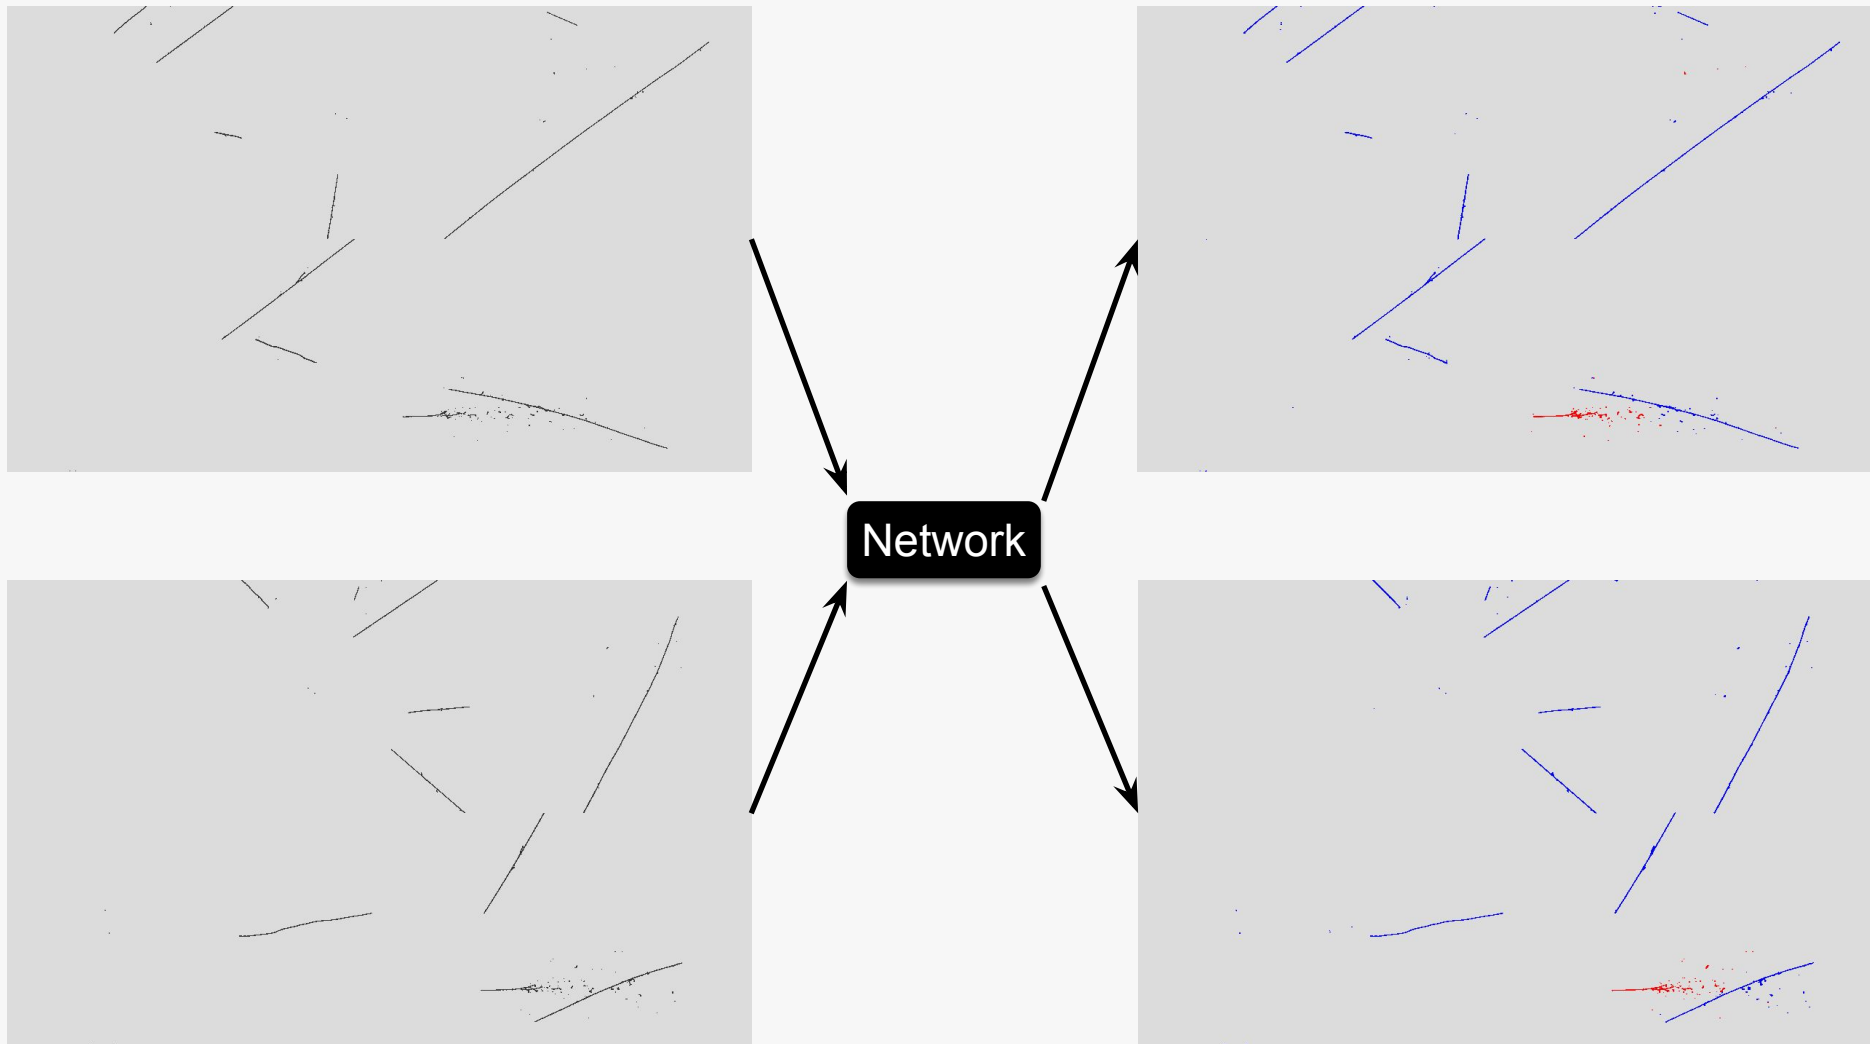


For detectors on the surface of the earth, the particles that produce ionization from neutrino interactions (electrons, protons, muons, photons ...) are the same ones that produce ionization from cosmic rays.

The signatures in the detector are indistinguishable without pattern recognition, but cosmic particles are **much** more common (many orders of magnitude)

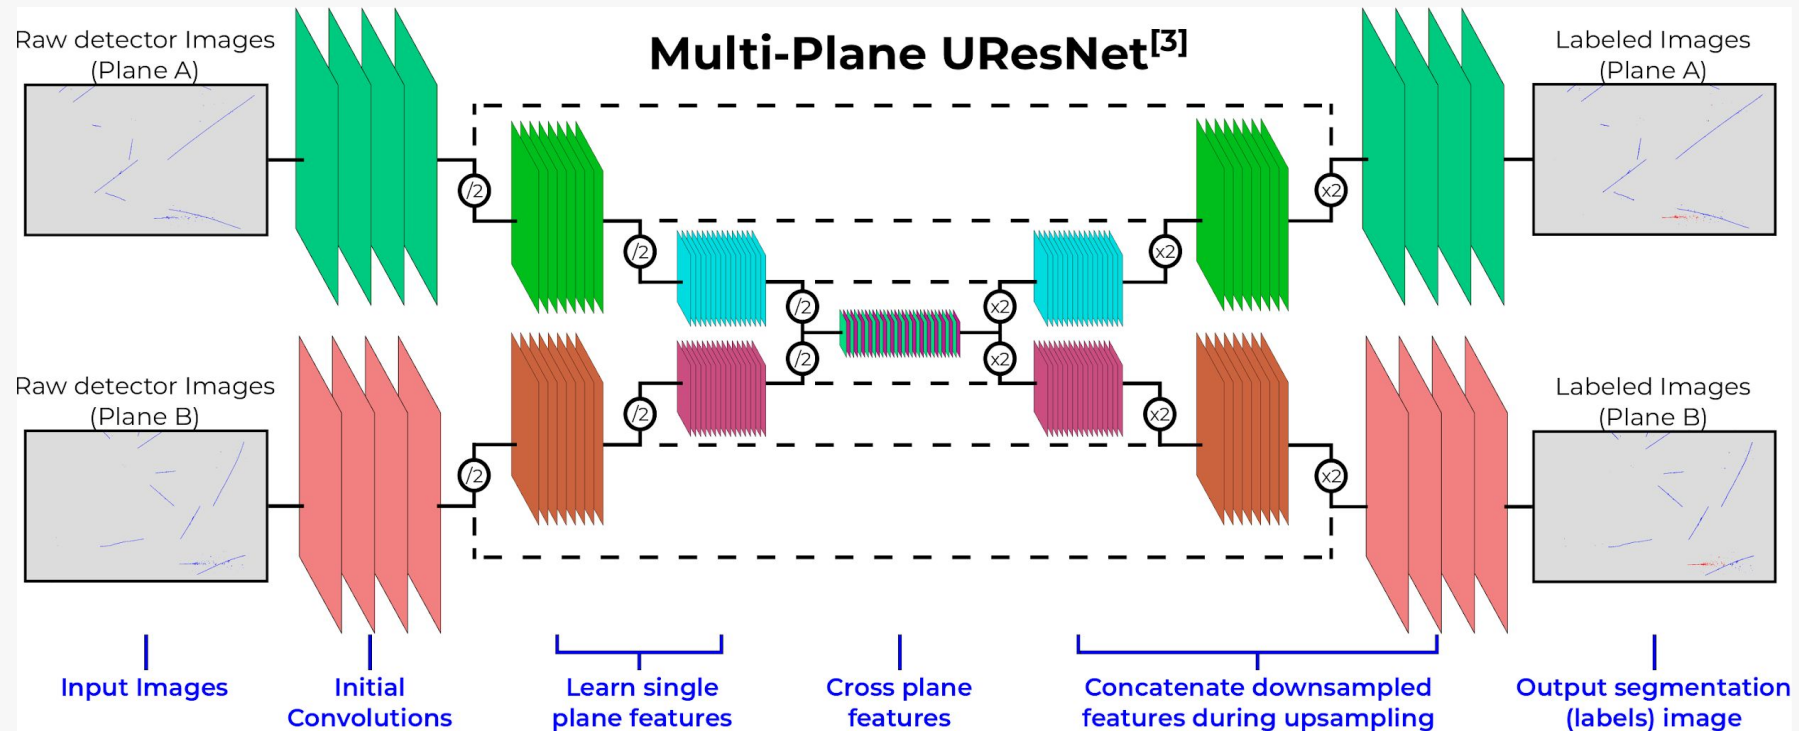


# Cross Plane Information



# Cross Plane Information - Segmentation

How to best utilize multiplane detectors?




Correlations across projections are correlated and learnable.

# Multiplane Segmentation





Network Input



Network Output  
**(No cosmics!)**

# Segmentation: Zoom on the Neutrino



Neutrino passing rate: 80%

Background-only: 6%

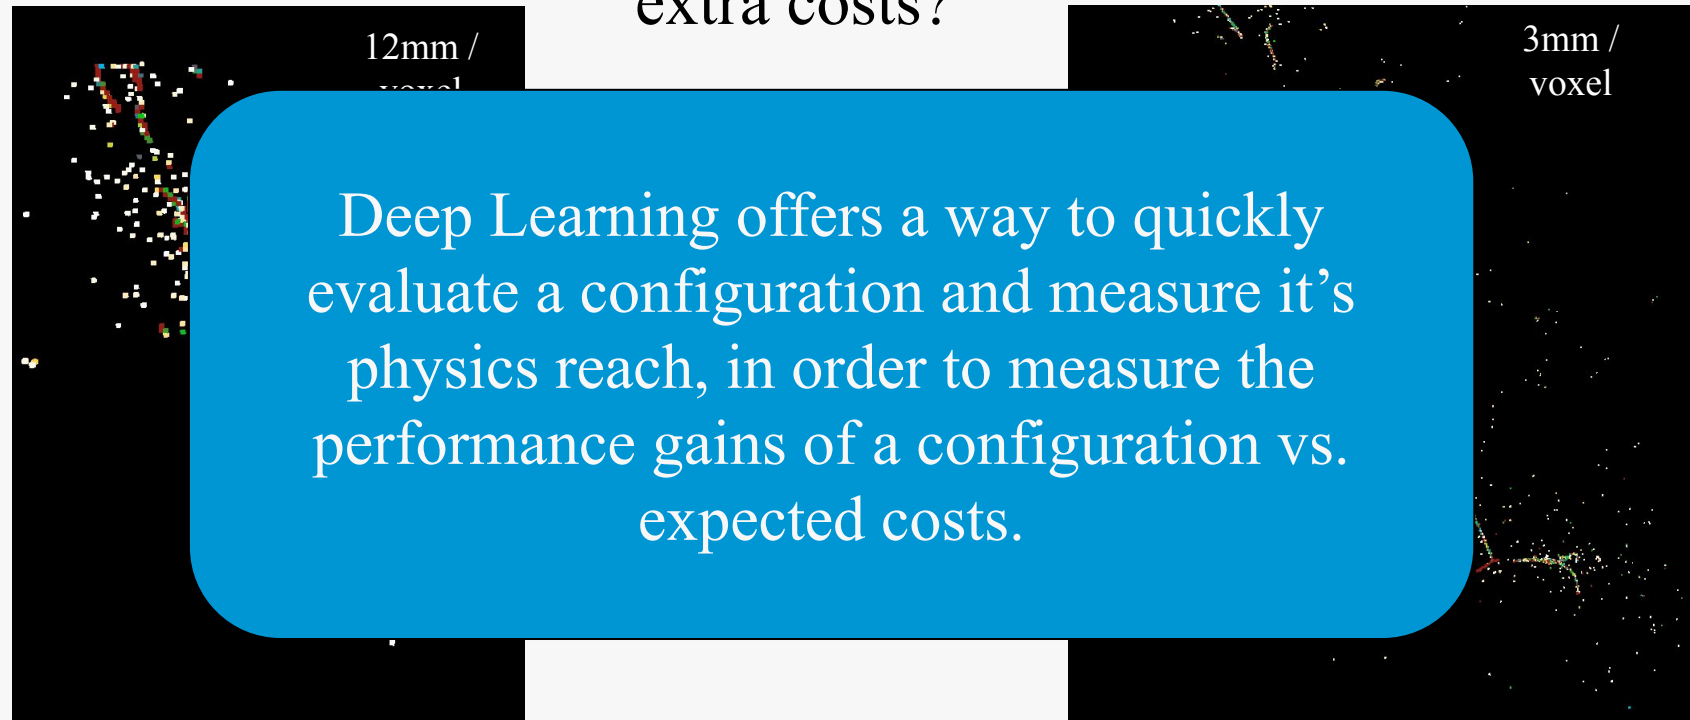
Neutrino Purity: 89%

**There is no corresponding classical algorithm that achieves this result.**

# Natively 3D Detectors

Most LArTPCs at large scale are multi projection, 2D detectors - **but pixelated, 3D detectors are coming.**



An obvious question: are the performance gains worth the extra costs?



# Sparse vs. Dense Convolutions

Deep learning in neutrino datasets could benefit dramatically from an intelligent application of sparse techniques to machine learning.

- Save **memory** by not storing activations on empty pixel sites
- Save **computations** by ignoring non-active sites.



Neutrino images from liquid argon time projections chambers are very sparse.

Occupied pixel fraction is:

< 0.01 in 2D at  $512^2$  pixels

<  $5e-5$  in 3D at  $512^3$  pixels

# Sparse Machine Learning Implementations

**Details vary, but there is some literature on convolutional neural networks on spatially sparse data.**

1. SBNET (from Uber) - inference only gains in speed by only feeding forward active blocks of input.
2. OctNet - sparse convolutions with tree based organization for sparse convolutions
3. Sparse 3D Conv. Nets - hashmap based convolutions for managing sparse representations
4. Submanifold Sparse Convolutional Networks - follow up on 3, deals with dilation as well.
5. (PointNet - graph network that works OK on sparse data)
6. (Dynamic Graph CNN - graph network with convolution-like operation)

# Submanifold Sparse Convolutional Networks

This technique brings three new convolution operators that differ from traditional convolutions:

- **Sparse Convolutions** operate only on active input sites, but behave traditionally on the output layer.
- **Submanifold Sparse Convolutions** maintain the same set of active sites on the input and output layers (and are “valid” - no change to spatial size)
- **Sparse Deconvolutions** invert sparse convolutions by reversing the mapping from input to output of a sparse convolution, useful in UNet or FCN segmentation architectures.

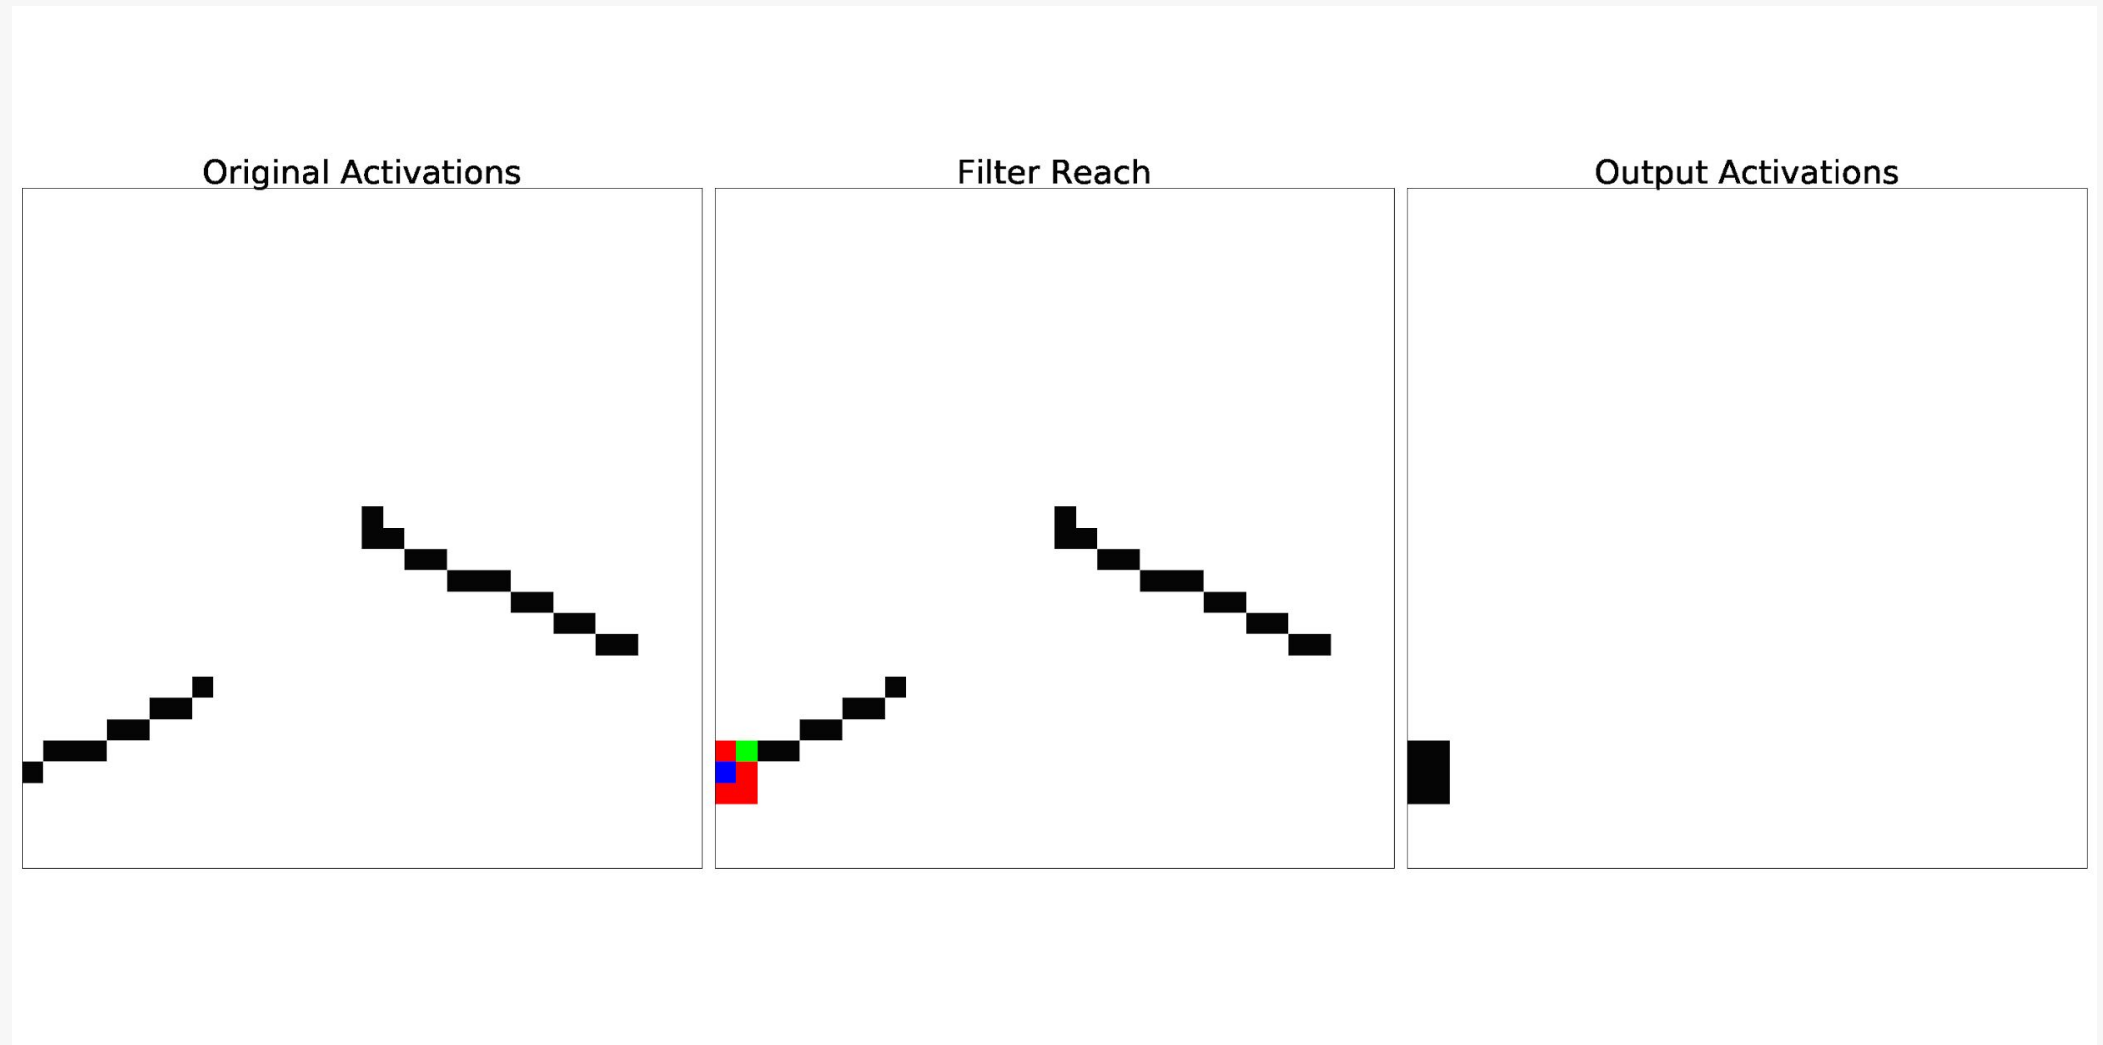


# Submanifold Sparse Convolutional Networks

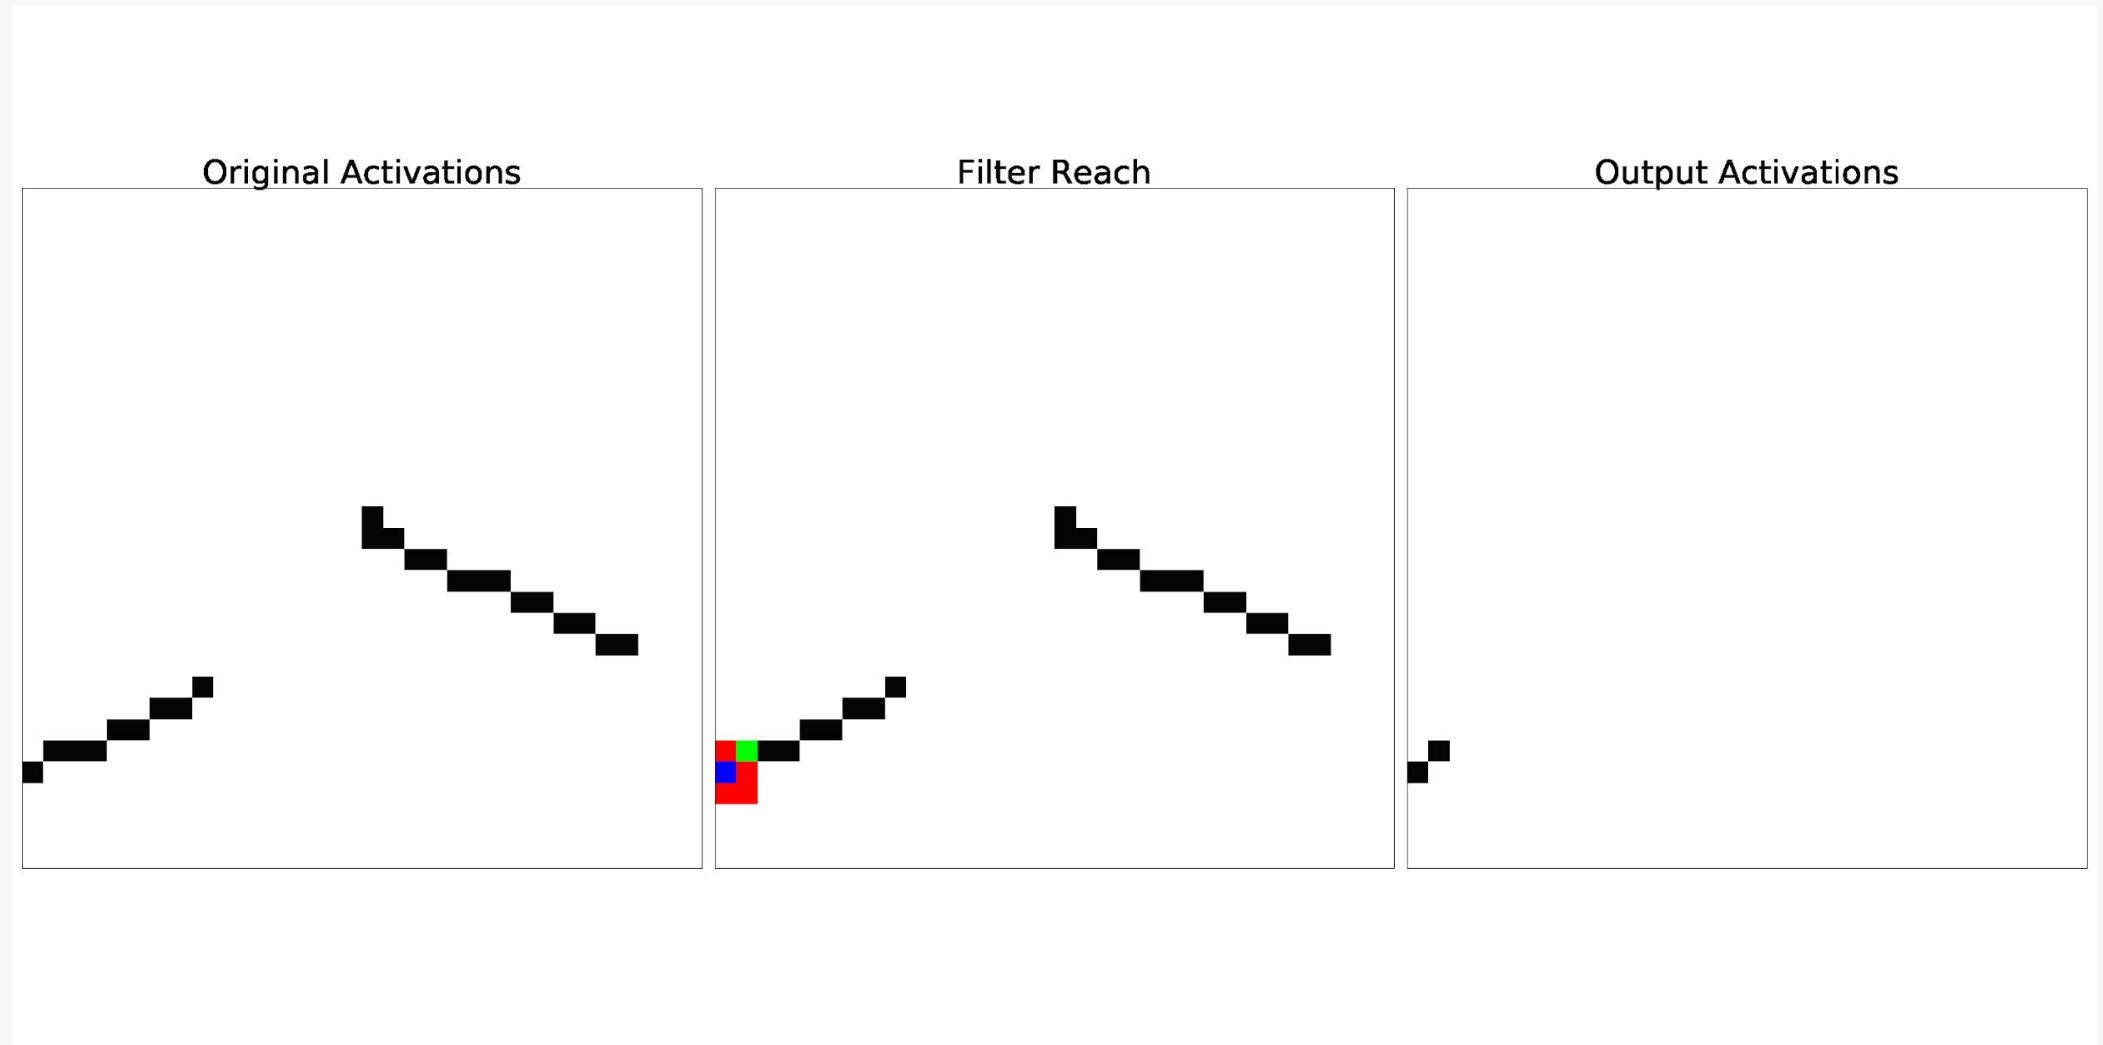
This technique brings three new convolution operators that differ from traditional convolutions:

- **Sparse Convolutions** operate only on active input sites, but behave traditionally on the output layer.
- **Submanifold Sparse Convolutions** maintain the same set of active sites on the input and output layers (and are “valid” - no change to spatial size)
- **Sparse Deconvolutions** invert sparse convolutions by reversing the mapping from input to output of a sparse convolution, useful in UNet or FCN segmentation architectures.



# Sparse Convolutional



# Submanifold Sparse Convolutional





# Sparse/Dense Comparison



Roughly 18x speedup in training on identical batch sizes, 58x speedup in peak single-GPU throughput.

# Sparse/Dense Comparison



60 Epochs (40 in 3D) on 64 GPUs in just 4 hours! Approaches state-of-the-art performance, quantitative study starting to benchmark accuracy of dense vs. sparse implementations.

# Sparse Convolutions


Relatively niche computer vision application could be very impactful in particle physics:

- Dramatic speed up of both training and inference times
- Substantial reduction of memory impact means less powerful machines can perform inference
  - Scalable to CPUs of open science grid
- Neutrino datasets can feed back to the computer vision community for development of new techniques (instance detection?) with sparse convolutions.

# Where is the field going?

# Full Event Pattern Recognition

Instance aware, cross-plane particle segmentation.

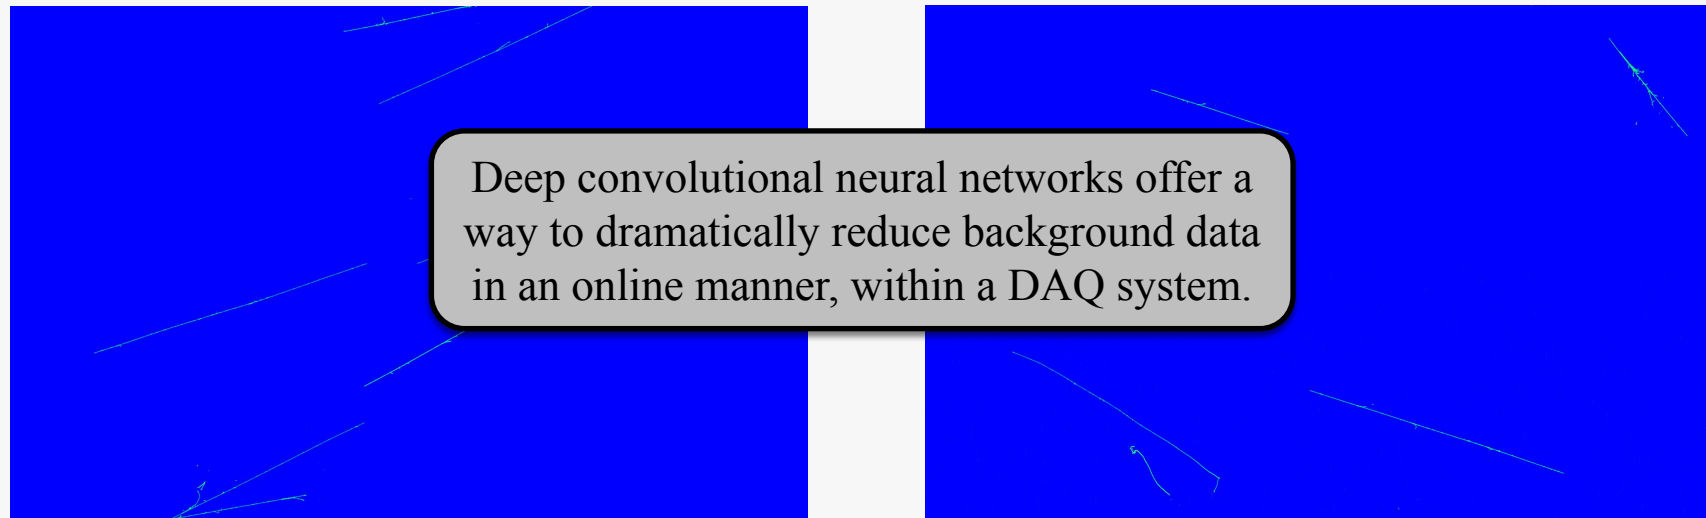


Already, successful demonstrations of segmentation, instance aware predictions, and multi-plane networks have succeeded.



# Online Triggers

MicroBooNE data rate is  $\sim 8$  GB/s. New detectors will only be bigger and faster - how to keep up with the data stream?



Background Only


Neutrino Interaction

**Requires a different data pipeline, but is feasible for future experiments.**

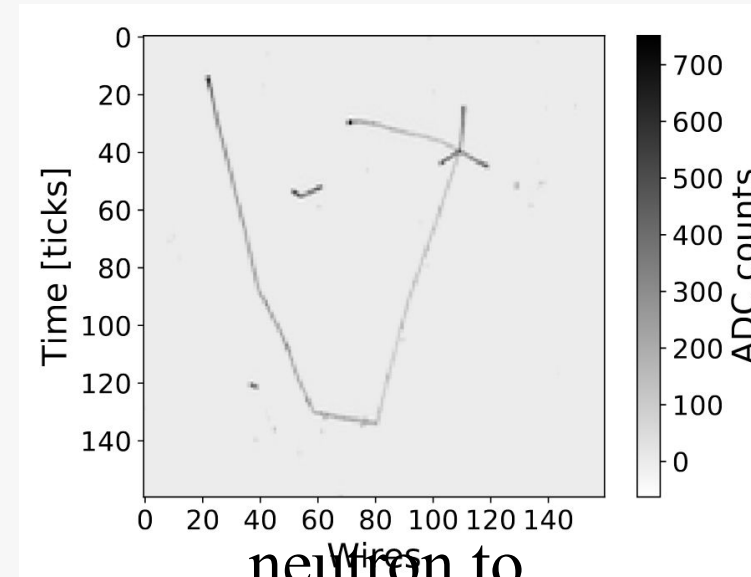
# Rare Event Searches

In the Deep Underground Physics Experiment, there are exciting physics studies to do that require powerful analysis techniques.

J. Hewes Thesis, 2017 (<http://inspirehep.net/record/1662648>)



Proton Decay




neutron to  
anti-neutron  
oscillation

# Collaborative Development

Many experiments (and many students!) are interested in deep learning, but applications to neutrino physics are still emerging and common tools unavailable.

Response: **deeplearnphysics.org** to share open source development:



- Open source data format (and open data set)
- OpenGL Interactive Visualization Tools
- Open source network implementations
- Extensive Tutorials
- Slack channel for news, questions, etc.
- Cross-experimental meetings

# DeepLearnPhysics

When Professor-of-Physics wants their new student to use deep learning in their thesis, where to turn?

## Open Data Set

- Freely available data from authentic neutrino simulation
- Tutorials and walkthroughs on how to generate application-specific datasets

## Model Zoo

- Open source implementations of important networks for physics
- Evaluation and pre-trained weights for open data

Goal: Papers released early 2019

# Open Data Set

The screenshot shows the OSFHOME website interface. At the top, there is a navigation bar with the OSFHOME logo, a search bar, and links for Support, Donate, Sign Up, and Sign In. Below the navigation bar, the page title is "Public Particle Imaging Dataset (PubPAID) by DeepLearnPhysics". The page content includes a description of the project, its contributors (DeepLearnPhysics), creation and update dates, and a list of components. A "Wiki" section provides a detailed description of the project's goals and a list of three key features: publicly available data, publicly available software container, and documented results. A "Files" section shows a table with columns for Name and Modified, listing the dataset and OSF Storage. A "Recent Activity" section shows updates to the license and title. A cookie consent banner is visible at the bottom.

OSFHOME

Public Particle Imaging Dataset (PubPAID) by DeepLearnPhysics

Contributors: DeepLearnPhysics

Date created: 2018-12-03 01:23 PM | Last Updated: 2019-02-06 08:05 PM

Category: Project

Description: PubPAID is a data sharing project organized by DeepLearnPhysics, a group of researchers developing ML techniques and applications for science. This project contains (at least) 2 levels of sub-projects. The lowest level projects contain data files, and intermediate projects define group of applications and/or science domains. See the wiki for more details.

License: CC-BY Attribution 4.0 International

Wiki

This is the top level project for data sharing sub-projects organized by researchers in the DeepLearnPhysics organization. We aim to encourage and maintain highly reproducible research work by other researchers across different domains. We aim to achieve this by providing three things:

1. Publicly available data
2. Publicly available software container
3. Documented results (publication)

This project is...  
[Read More](#)

Files

Name	Modified
Public Particle Imaging Dataset (PubPAID) by D...	
OSF Storage (United States)	


Recent Activity

- DeepLearnPhysics updated the license of Public Particle Imaging Dataset (PubPAID) by DeepLearnPhysics to CC-BY Attribution 4.0 International 2019-02-06 08:05 PM
- DeepLearnPhysics changed the title from Open Samples for Liquid Argon Time Projection Chambers (LArTPCs) to Liquid Argon Time Projection Chambers

This website relies on cookies to help provide a better user experience. By clicking Accept or continuing to use the site, you agree. For more information, see our [Privacy Policy](#) and information on [cookie use](#).


Accept

- Data set is public
- Descriptive Paper is in 2nd draft
- Model zoo is coming!



# DUNE

Neutrino Physics is investing into large, liquid argon imaging detectors.

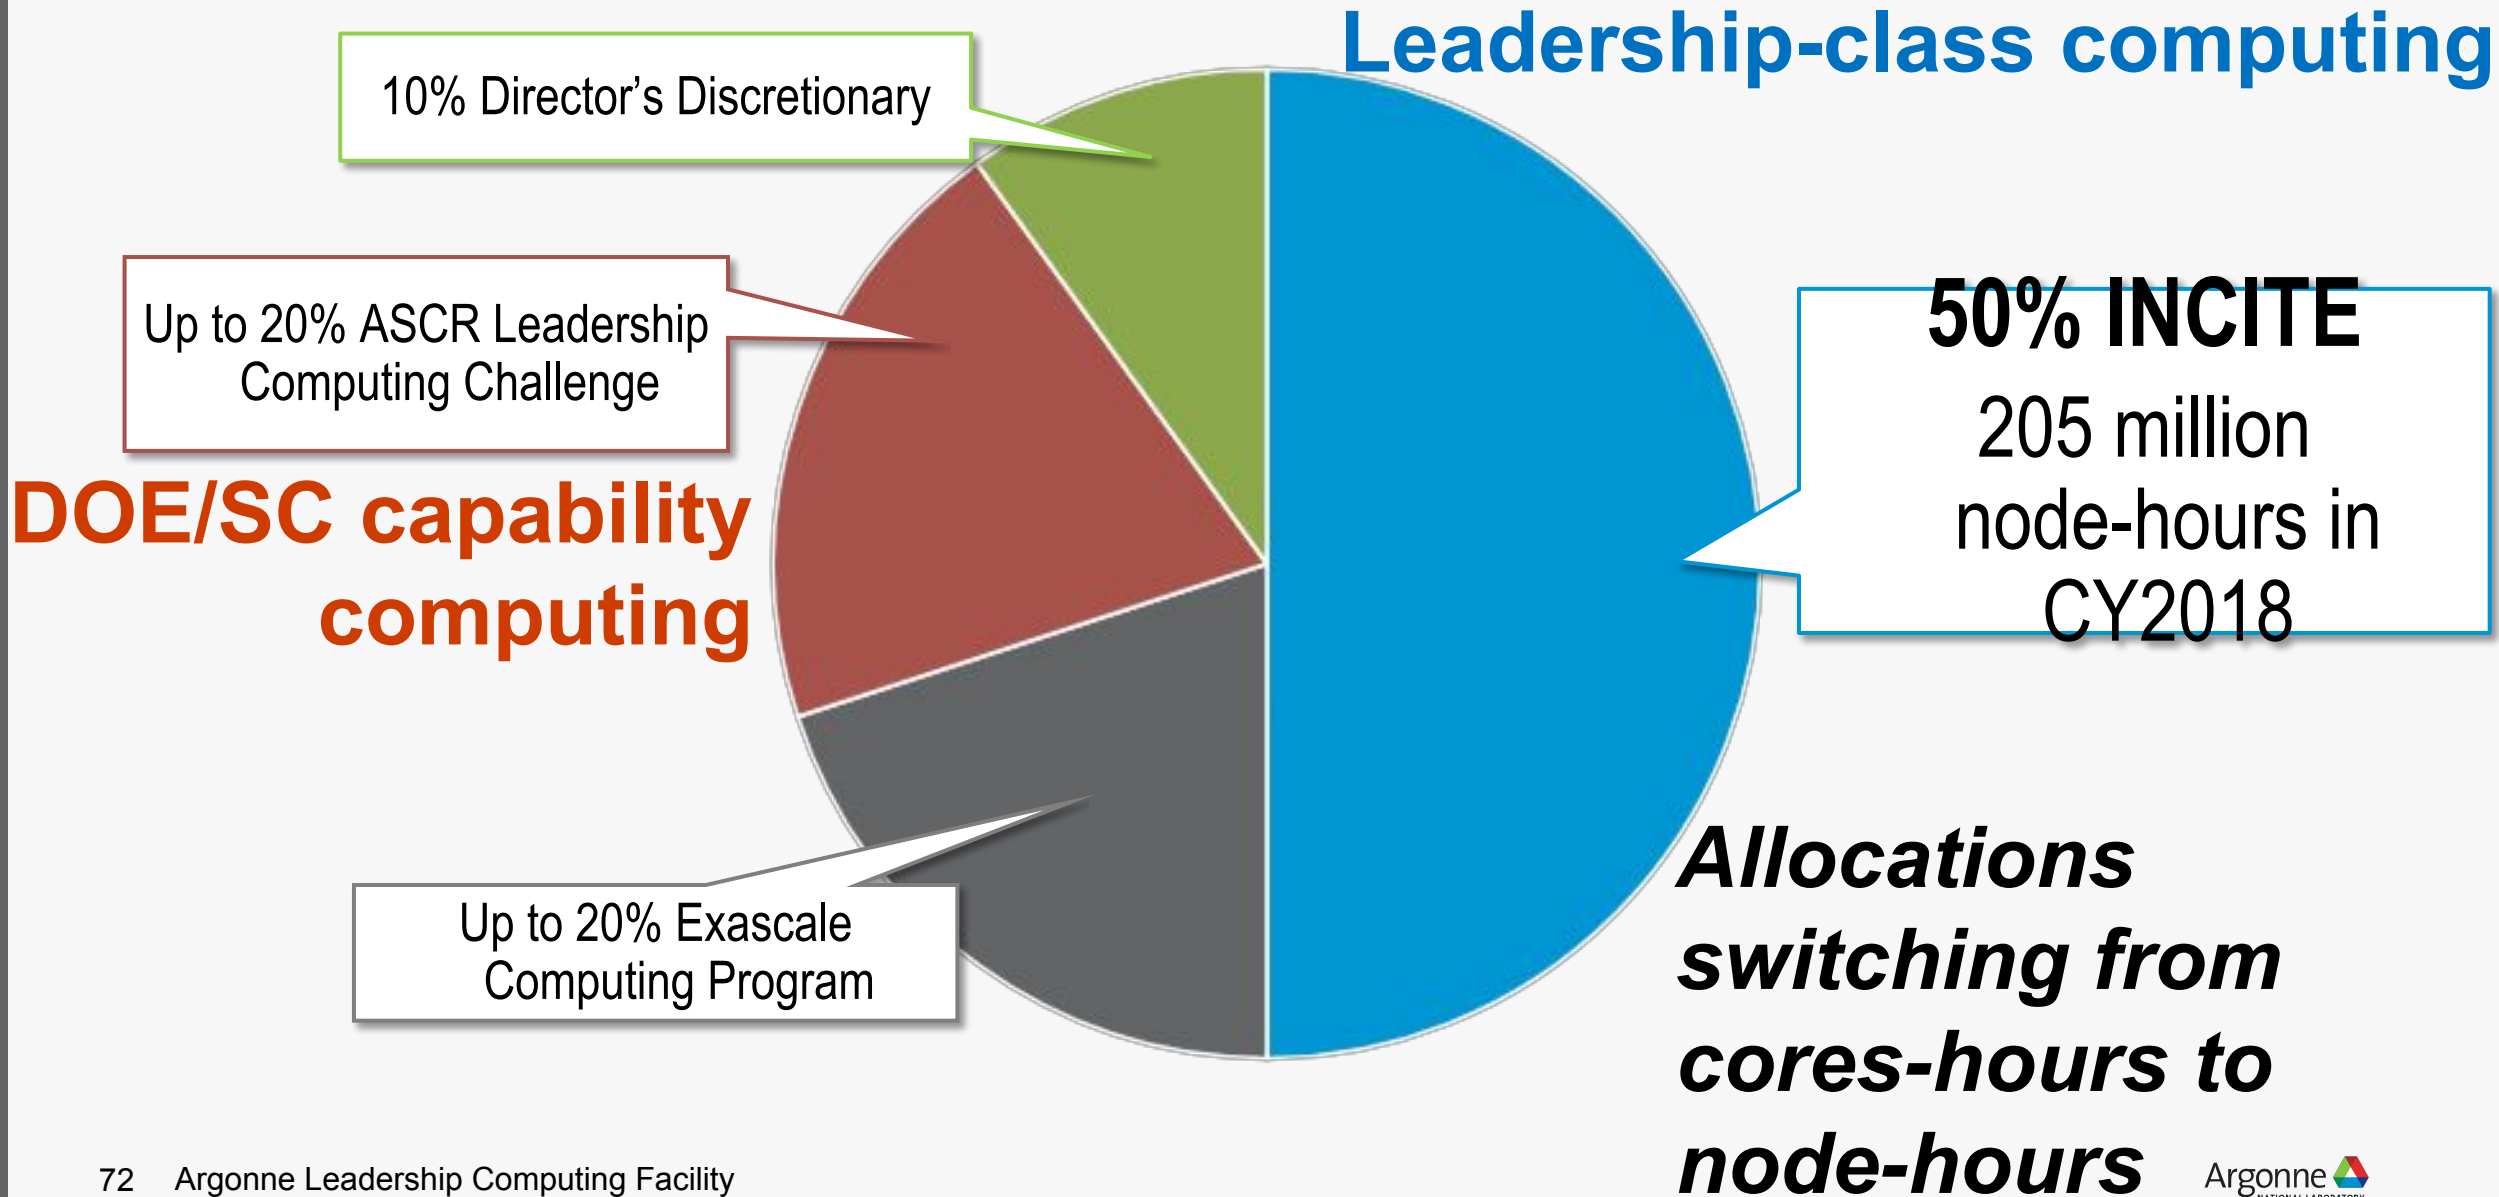


Deep learning will be THE tool that meets the science needs of DUNE.

# Your Science @ ALCF

# Three primary ways for access to LCF

## Distribution of allocable hours





# Software

- ML/DL:

- TensorFlow, Keras, Neon, MXNet, PyTorch, Sci-kit Learn, Graph Analytics (Cray Graph Engine), Horovod...
- With performance libraries e.g. Intel MKL, MKL-DNN, LibXSMM etc enabled
- Intel optimized Tensorflow
  - Conda package on Theta
  - Intel Distribution for Python's optimized numpy

```
# Python 2.7
pip install https://anaconda.org/intel/tensorflow/1.4.0/download/tensorflow-1.4.0-cp27-cp27mu-linux_x86_64.whl

# Python 3.5
pip install https://anaconda.org/intel/tensorflow/1.3.0/download/tensorflow-1.3.0-cp35-cp35m-linux_x86_64.whl

# Python 3.6
pip install https://anaconda.org/intel/tensorflow/1.4.0/download/tensorflow-1.4.0-cp36-cp36m-linux_x86_64.whl
```



PYTORCH



# Software

- **Workflow/Data analysis:**

- Containers

- Singularity container solution for application science workloads
    - Environment imported into container
    - mount additional directories into the container with the -B flag
    - `aprun -n $RANKS -N 1 singularity exec my_image.img ./my_binary`


- Balsam

- Jupyter Hub, MongoDB, Apache Spark, R

- Python

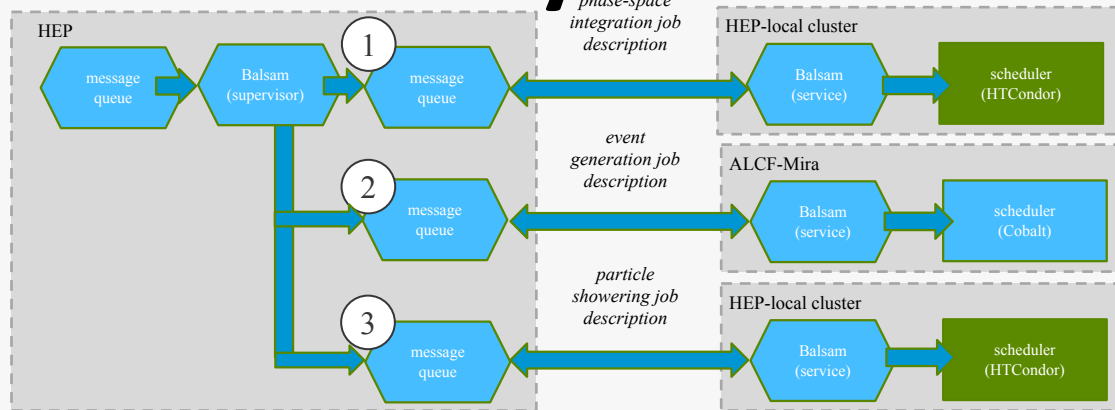
- Intel and Cray modules on Theta
    - ALCF `alcfpython/2.7.14-20180131`

- **Visualization:** Paraview on Theta



# BALSAM Workflow manager and edge services for alcf systems

<https://www.alcf.anl.gov/balsam>



Schematic of ATLAS deployment of Balsam on multiple resources to execute a workflow with alternating serial and parallel stages

Balsam is a workflow manager that simplifies the task of running large-scale job campaigns on ALCF resources while minimizing user involvement and improving productivity. It interacts closely with job schedulers to optimize job throughput via individual jobs and ensemble jobs, staging data in and out as needed. A supervisor component manages execution across multiple compute resources.


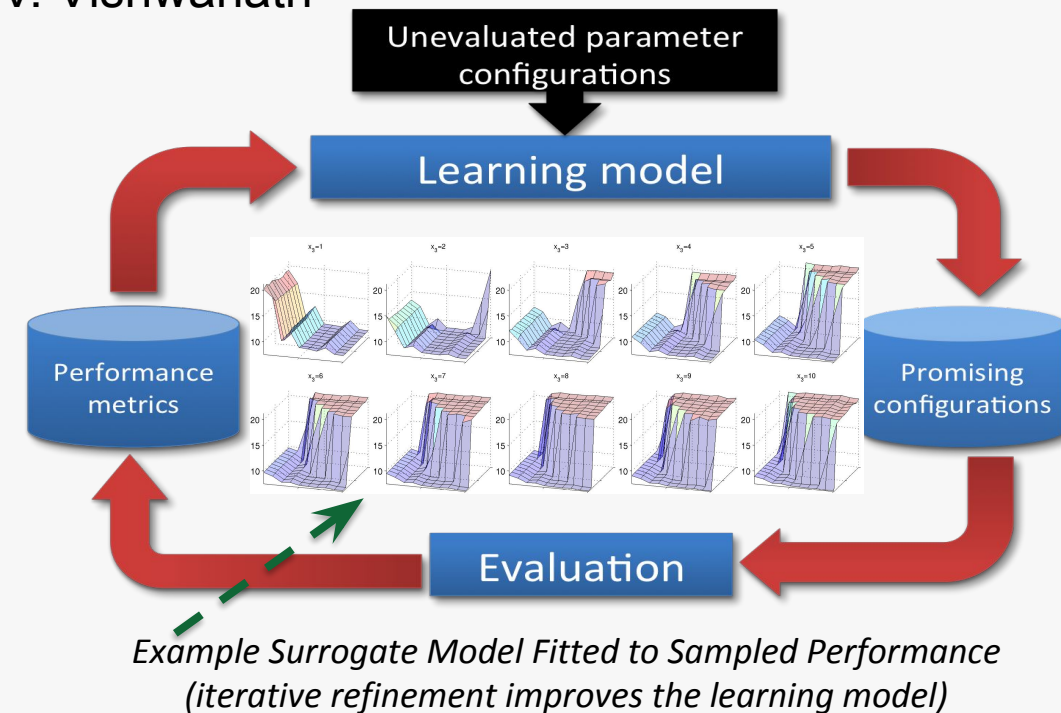
## The Impact

- Delivered >150M Core hours for production science on ALCF systems
- The ATLAS experiment has used Balsam to run hundreds of millions of compute hours of event generation jobs on ALCF systems. *ALCF contribution to ATLAS computing* ranks as the 6<sup>th</sup> largest country in the world.
- The DIII-D National Fusion Facility used Balsam to trigger experiment-time analyses during Tokamak operation, running more complex analyses in less time, leading to higher experiment productivity
- We are investigating use of Balsam for several other projects, including real-time APS/ALCF, and in ADSP and ECP projects

# DeepHyper: Scalable hyperparameter search for deep learning

P. Balaprakash, E. Jennings, M. Salim, T. Uram, S. Wild, V. Vishwanath

- Development of DL algorithms is time consuming and a significant portion is spent on tuning and optimizing the hyperparameters such as number of layers number of units learning rate optimizer epochs,
- Model-based search iteratively refines the model in promising input region by obtaining new outputs at unevaluated input configurations
- Framework:
  - Initialization phase
    - Random or Latin hypercube sampling
  - Iterative phase
    - Fit model
    - Sample using the model
- Integration with the Balsam workflow enables for improved used productivity and performance. This facilitates job submission, data staging, model checkpointing, resource allocation, etc. **Scaled to 1024 Theta nodes** and 64 Cooley nodes



# Map To Allocation Programs

<b>Program</b>		
<b>INCITE</b>	<b>Production</b>	<b>Capability Computing</b>
<b>ALCC</b>	<b>Production</b>	<b>SC Capability Computing missions driven</b>
<b>Discretionary</b>	<b>Production</b>	<b>Development, testing, proposal preparation</b>
<b>Data Science Program</b>	<b>Production</b>	Developing technical and science capability for data and learning based workflows
<b>Early Science Program</b>	<b>Next Generation</b>	Developing technical and science capability for next generation systems
<b>Exascale Computing Projects</b>	<b>Next Generation</b>	The ECP mission-driven

<https://www.alcf.anl.gov/user-guides/how-get-allocation>

# ALCF User Training Workshops and Opportunities



Most Recently: October 2-4, 2018



April 30 - May 2, 2019 [Registration](#)

Argonne Leadership Computing Facility  
an Office of Science user facility



A screenshot of the ALCF User Support page. The page has a dark blue header with navigation links: ABOUT, COMPUTING RESOURCES, SCIENCE AT ALCF, PROJECTS, NEWS &amp; EVENTS, USER SUPPORT, USER GUIDES, and AURORA. The main content area is white and features a 'User Support' sidebar on the left with links for User Support, Machine Status, Presentations, User Survey, and Training &amp; Outreach. The Training &amp; Outreach section lists: ALCF Computational Performance Workshop, ALCF Getting Started Videoconferences, Quantum Computing Workshop, and Simulation.Data.Learning Workshop. The main content area has a header 'Growing the HPC community' and a 'Training' section. The Training section states: 'We offer a variety of videoconferences and workshops each year to provide training and information for new and existing users.' Below this is a '2018 Training Opportunities' section with a list of events: January 16, January 18, January 24, January 26 | Getting Started on Theta and Mira; February 27-March 1, 2018 | Simulation, Data, and Learning Workshop; May 15-May 17, 2018 | ALCF Computational Performance Workshop; July 24, July 26 | Getting Started on Theta and Mira; Ongoing | ALCF Many-Core Developer Sessions; Ongoing | ALCF Many-Core Tools &amp; Techniques Video Series (Video presentations to introduce and refresh users on KNL processors.); and Ongoing | ATPESC 2017 Lecture Videos.

<https://www.alcf.anl.gov/training>

Argonne Training Program on  
Extreme-Scale Computing (ATPESC)  
<https://extremecomputingtraining.anl.gov/application/>

# Thank You!!!

**Feel free to contact us:**

[corey.adams@anl.gov](mailto:corey.adams@anl.gov)


[datascience@alcf.anl.gov](mailto:datascience@alcf.anl.gov)

<https://www.alcf.anl.gov/alcf-data-science-program>


# Extra Materials



# Finding Neutrinos




# First Steps: Nova



The Nova experiment trained a modified version of AlexNet to classify their neutrino interactions by type of neutrino.

The convolutional neural network performed as well or better than traditional pattern recognition in all cases.



The primary physics search saw an efficiency boost of **40%**.



# MicroBooNE: Particle ID

Sample	Electron	Photon	Muon	Pion	Proton
Detection Accuracy (%)	77.8 +/- 0.7	83.4 +/- 0.6	89.7 +/- 0.5	71.0 +/- 0.7	91.2 +/- 0.5
Most frequent MisID (%)	$\gamma$ (19.9)	$e^-$ (15.0)	$\pi^-$ (5.4)	$\mu^-$ (22.6)	$\mu^-$ (4.6)




Using AlexNet, can individual particles be distinguished? Yes...  
JINST 12, P03011 (2017).



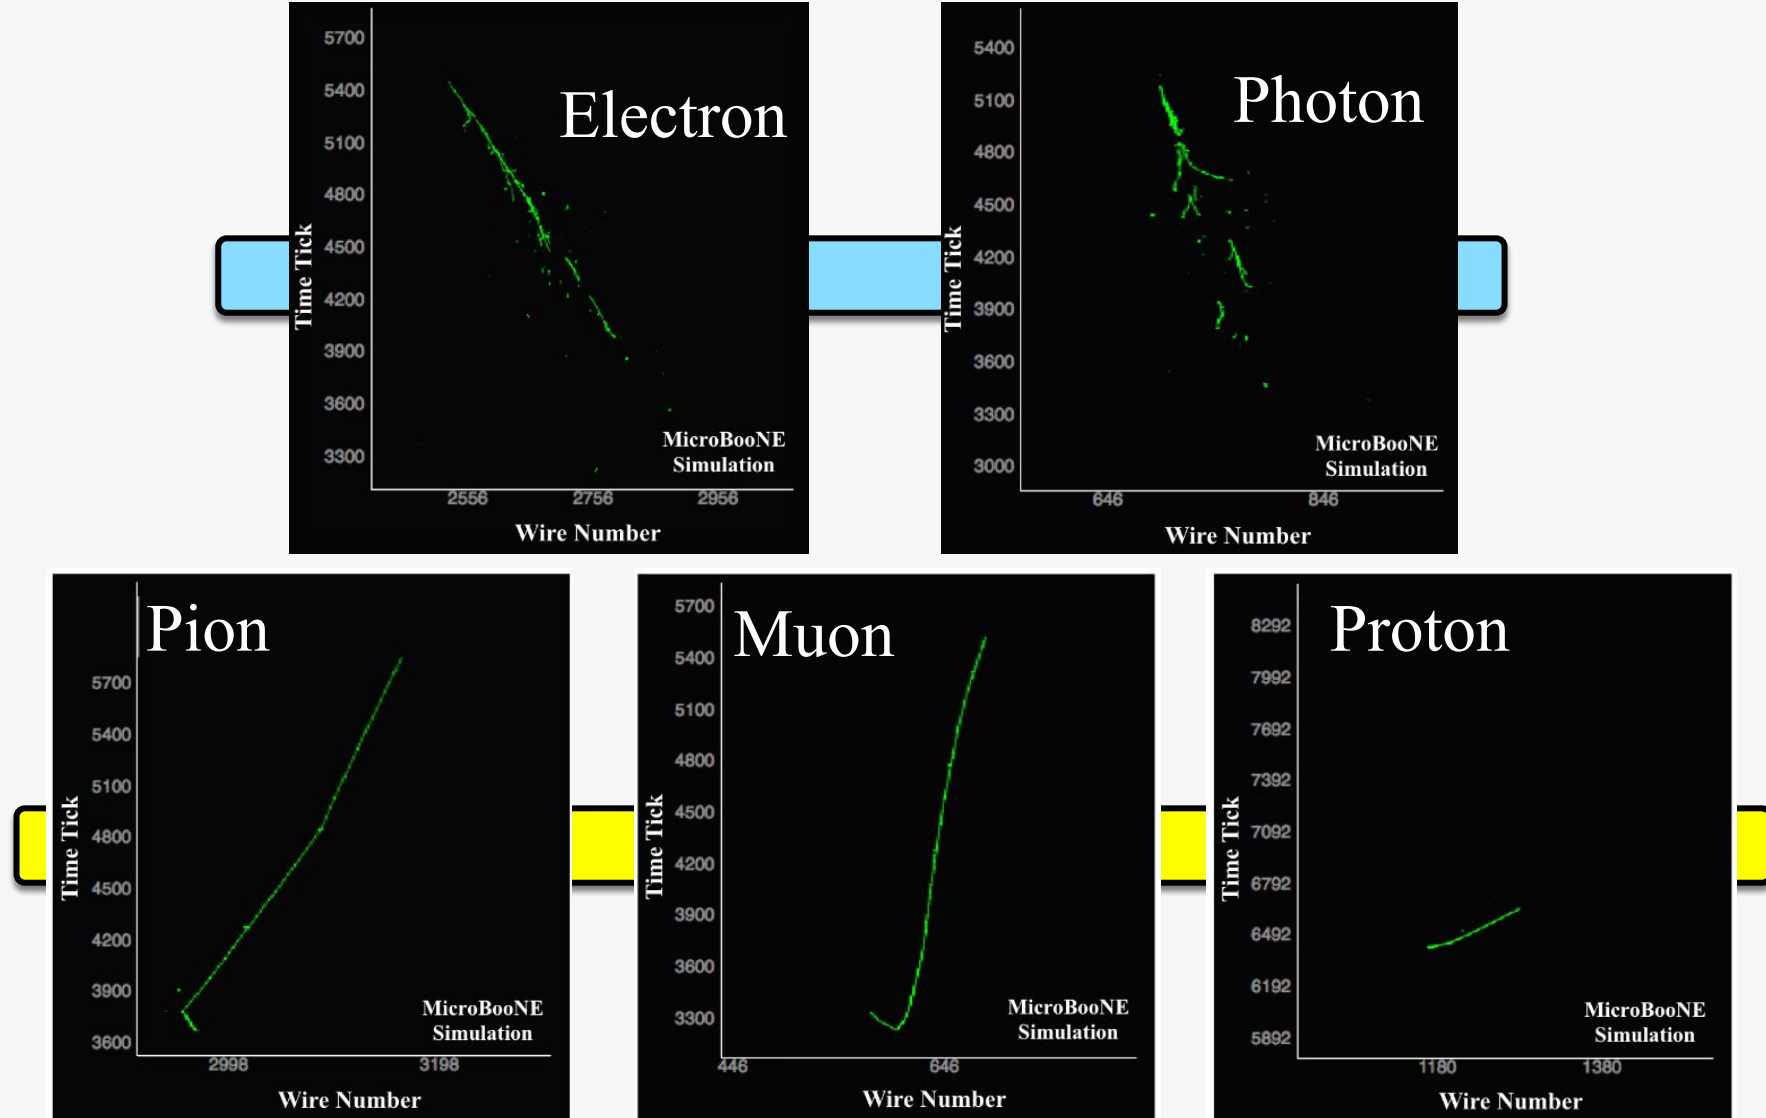
# MicroBooNE: Particle ID

Sample	Electron	Photon	Muon	Pion	Proton
Detection Accuracy (%)	77.8 +/- 0.7	83.4 +/- 0.6	89.7 +/- 0.5	71.0 +/- 0.7	91.2 +/- 0.5
Most frequent MisID (%)	$\gamma$ (19.9)	$e^-$ (15.0)	$\pi^-$ (5.4)	$\mu^-$ (22.6)	$\mu^-$ (4.6)

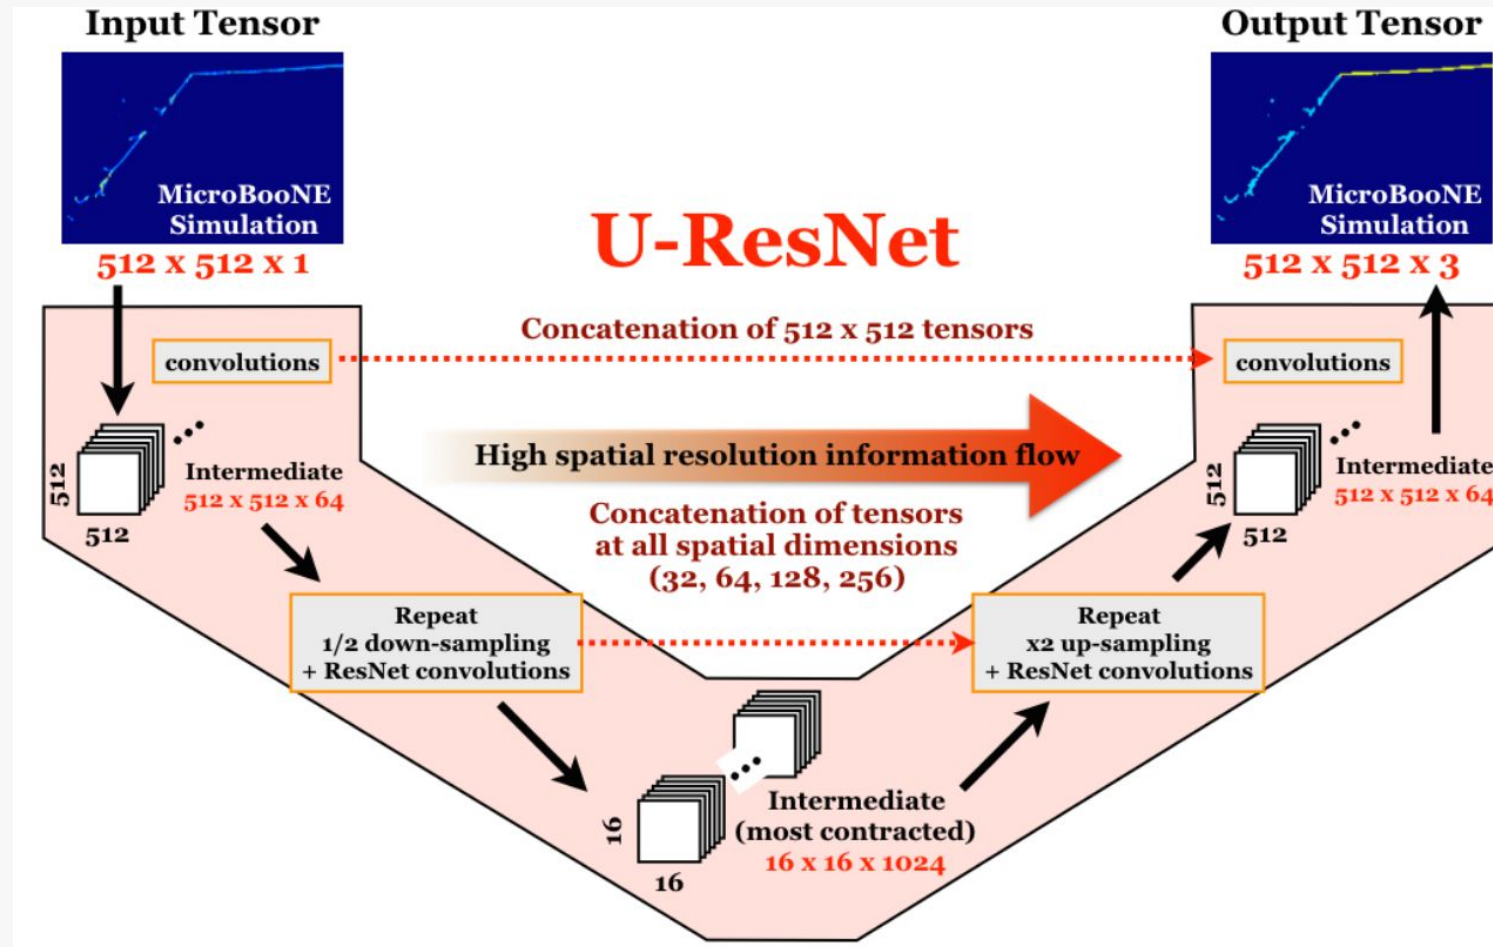
Using AlexNet, can individual particles be distinguished? Yes...  
JINST 12, P03011 (2017).



# MicroBooNE: Particle ID




# Semantic Segmentation




A Deep Neural Network for Pixel-Level Electromagnetic Particle Identification in the MicroBooNE Liquid Argon Time Projection Chamber  
(<https://arxiv.org/pdf/1808.07269.pdf>)

# Semantic Segmentation



In uncrowded interactions, this is enough to augment traditional track-fitting algorithms.


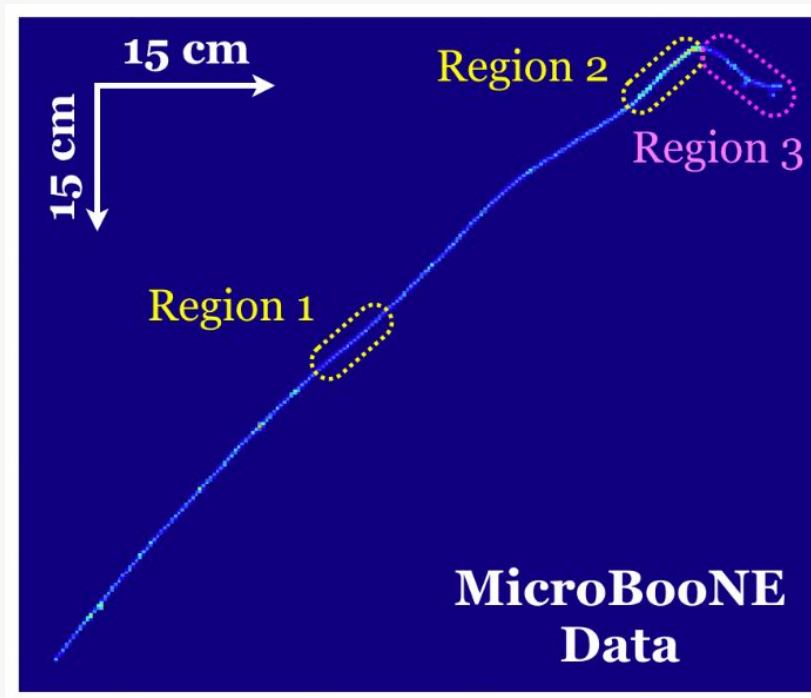


**How do you evaluate performance on data?**

# Data Validation

Even with hand-labeling, how to find a benchmark data set?  
**Leverage existing physics analyses for network validation.**

## Cosmic Muon Decay Events

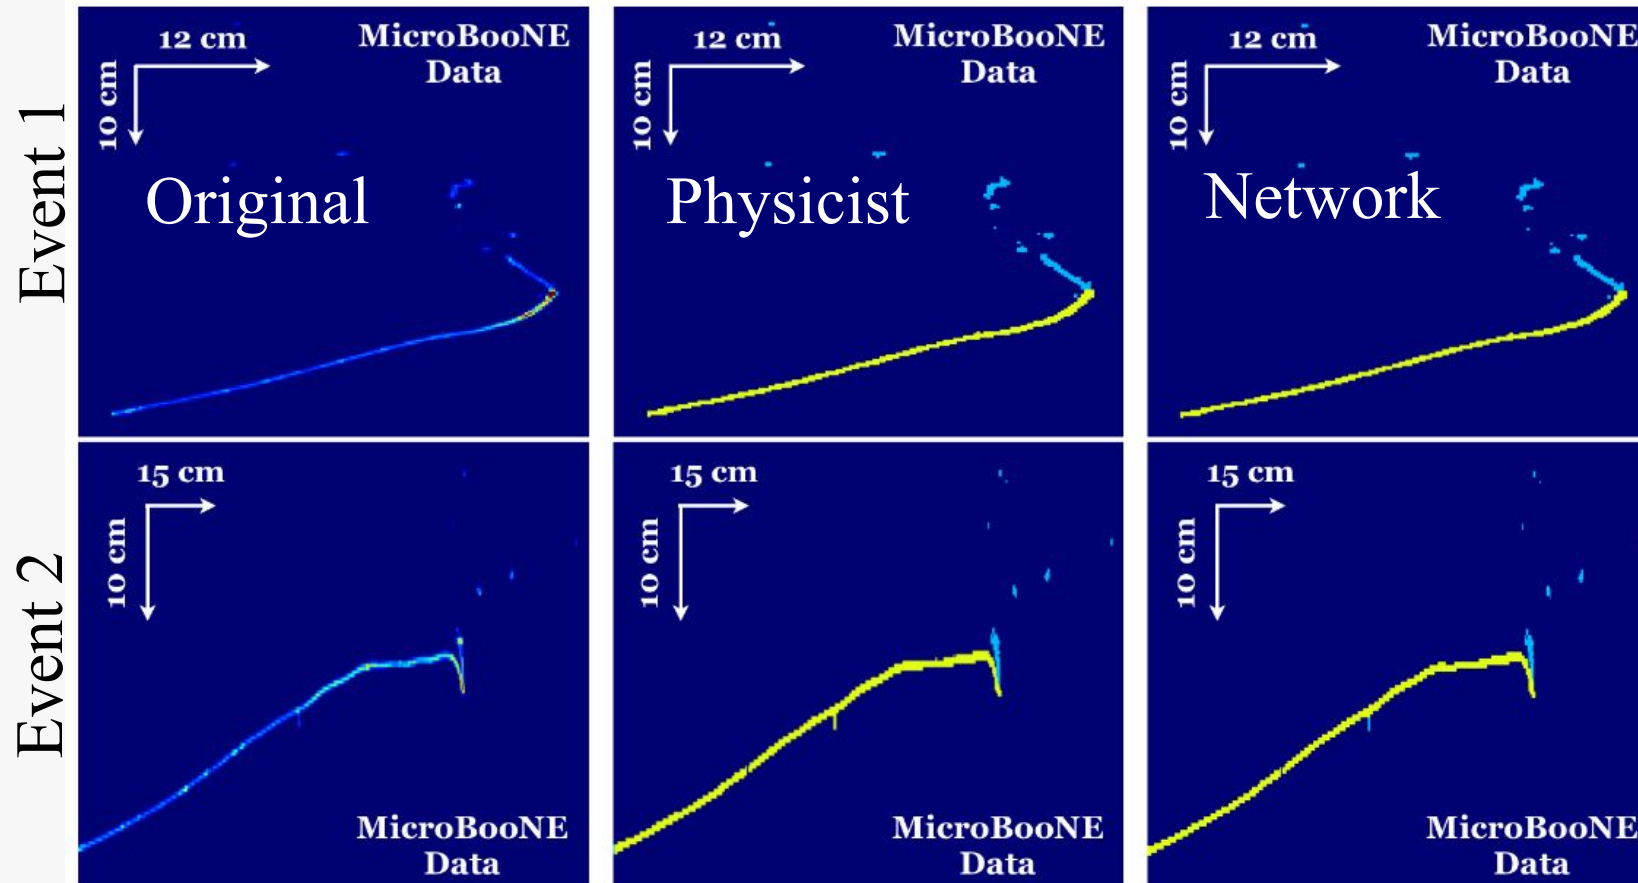


Comparison between network labels and physicist labels.






# Data Validation

Even with hand-labeling, how to find a benchmark data set?  
**Leverage existing physics analyses for network validation.**




# Physicist vs. The Machine

Sample	Data	Simulation	Simulation	Simulation
Label	Physicist	Physicist	Simulation	Simulation
Prediction	U-ResNet	U-ResNet	U-ResNet	Physicist
ICPF mean	3.4	2.5	1.8	2.0
ICPF 90%	9.0	5.7	4.6	4.8
Shower	4.8	3.4	3.0	2.6
Track	2.7	2.4	2.2	2.9



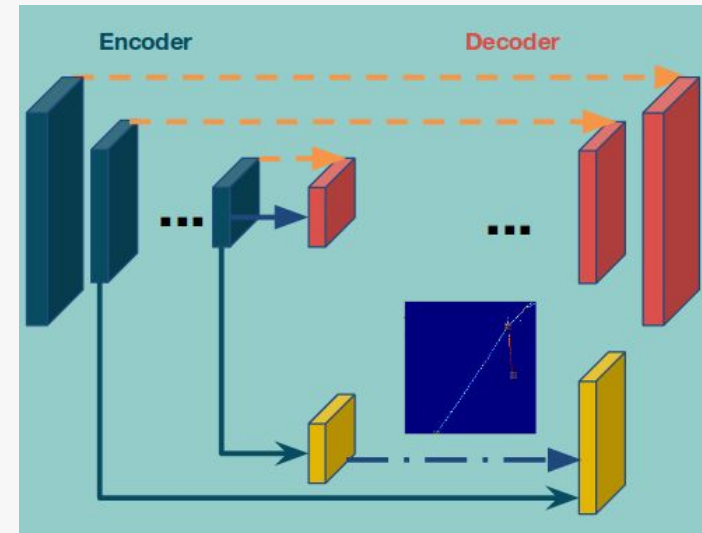
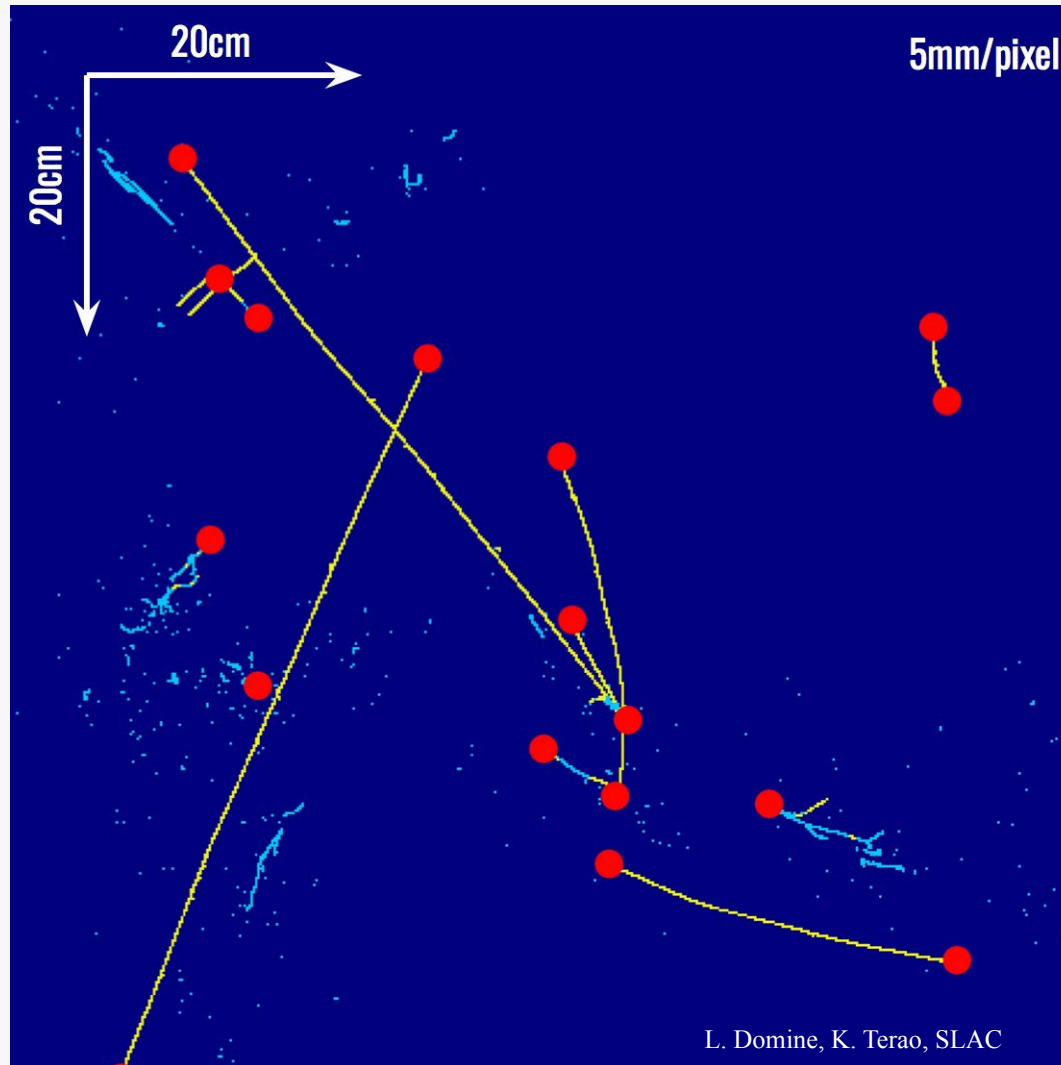
# Point Prediction



Network features for track/shower identification should also be useful for instance aware tasks.

Can a Region Proposal Network predict the start and end points of each particle?

# Point Prediction





Can a Region Proposal Network predict the start and end points of each particle?

**Yes**

# Segmentation in 3D

Meanwhile, extending segmentation networks to 3D has proven feasible and results are encouraging:



Better resolution gives better results, but GPU memory rapidly becomes an issue - difficult to train.

# Convolution Implementation

Given a set of sparse input, in an  $N$  dimensional volume, the input data is represented as  $\mathbf{m}$  features over  $\mathbf{a}$  active locations and  $\mathbf{a}$  coordinates in  $N$  dimensional space.

1. Build a hashmap to convert a spatial coordinate into an array index. Map each  $N$  dimensional active location to a single row index in an  $\mathbf{a} \times \mathbf{m}$  matrix.
2. For each spatial location in a filter ( $3 \times 3 = 9$  locations, for example), create a list for each spatial (input, output) pair for that filter location.
  1. Submanifold convolutions enforce output locations to be in the input feature set, sparse convolutions allow all valid output locations
3. For input location  $i$  that maps to output location  $j$ , perform the matrix multiply of the (input row  $i$ )  $\times$  ( $\mathbf{m} \times \mathbf{n}$ ) filter for that spatial location, and add it to output row  $j$ .