Data Reconstruction Using Deep Neural Networks for Particle Imaging Neutrino Detectors

François Drielsma

on behalf of the SLAC ML group

drielsma@slac.stanford.edu

April 7, 2020





ML for LArTPCs

SLAC Group



- Deep-learning-based data reconstruction chain for liquid argon time-projection chambers
- μ BOONE, pDUNE, ICARUS, ArgonCube 2x2, DUNE
- Disclaimer: no physics results in this talk





Group consists of three scientists, three postdocs, three grad students



High resolution **picture** of ionizing particle trajectories

- mm-scale spatial res.
- MeV-scale energy res.
- Scalable
- Dense
- High repetition rate



High resolution **picture** of ionizing particle trajectories

- mm-scale spatial res.
- MeV-scale energy res.
- Scalable
- Dense
- High repetition rate

Goal: measure ν_e appearance with unprecedented accuracy

- Track-shower separation
- $e \gamma$ separation
- \rightarrow Flavor separation



Introduction

Edrift

Y wire plane waveforms

Single 2D projection





Single 2D projection





Single 2D projection





Whole-image analysis **SLAC** ACCELERA







High-level Physics variables

- Interaction class (ν_e/ν_μ , CC/NC)
- Neutrino energy





Whole-image analysis **SLAC** ^{NATIC} ACCO

Technically the simplest approach, if it works...

Problems expected in high-res LArTPCs

- Complex topologies, huge phase space
- What if things fail ? Why ?

High-level Physics variables

- Interaction class ($u_e/
 u_\mu$, CC/NC)
- Neutrino energy





High-level Physics variables

- Interaction class ($u_e/
 u_\mu$, CC/NC)
- Neutrino energy

Whole-image analysis **SLAC NATION.**

Technically the simplest approach, if it works...

 $\mathbf{P}\text{roblems}$ expected in high-res LArTPCs

- Complex topologies, huge phase space
- What if things fail ? Why ?

 $\boldsymbol{C} ould$ we enforce physics principles:

- Key features like vertex
- dE/dx, particle ID





High-level Physics variables

- Interaction class ($u_e/
 u_\mu$, CC/NC)
- Neutrino energy

Whole-image analysis **SLAC ACC ACC LADORATO**

Technically the simplest approach, if it works...

 $\ensuremath{\textbf{P}}\xspace{roblems}$ expected in high-res LArTPCs

- Complex topologies, huge phase space
- What if things fail ? Why ?

Could we enforce physics principles:

- Key features like vertex
- dE/dx, particle ID
- ightarrow Yes ! That's our research !

What are the optimal network architectures for our data to maximize the physics output ?







Enforce extraction of hierarchical physics features

1. Pixel feature extraction + key points (particle start/end)





- 1. Pixel feature extraction + key points (particle start/end)
- 2. Vertex finding + particle clustering





- 1. Pixel feature extraction + key points (particle start/end)
- 2. Vertex finding + particle clustering
- 3. Particle type + energy/momentum



- 1. Pixel feature extraction + key points (particle start/end)
- 2. Vertex finding + particle clustering
- 3. Particle type + energy/momentum
- 4. Interaction (particle flow) reconstruction





Enforce extraction of hierarchical physics features

- 1. Pixel feature extraction + key points (particle start/end)
- 2. Vertex finding + particle clustering
- 3. Particle type + energy/momentum
- 4. Interaction (particle flow) reconstruction



Make it for 2D/3D data + whole chain trainable











ICARUS, arXiv:1210.5089

LArPix, arXiv:1808.02969

François Drielsma (SLAC)

ML Reconstruction Chain for LArTPCs



Space points



Algorithms to go from to 2D to 3D:

- BNL's WireCell
- T. Usher's Cluster3D

Cluster3D designed for high efficiency, relies on downstream space point solver

• Traditional likelihood-based

T. Usher, P. Tsang, L. Domine

- Semantic segmentation to discriminate against "ghost" points
- $\leftarrow \textbf{ICARUS simulation} \text{ on } 2.3^3 \text{ m}^3 \text{ region}$



Deghosting





François Drielsma (SLAC)

ML Reconstruction Chain for LArTPCs

Deghosting





François Drielsma (SLAC)

Particle type





François Drielsma (SLAC)

Particle type





Particle type identification accuracy:

Particle type	Voxel fraction	Accuracy
HIP	17 %	98.2 %
MIP	34 %	99.4 %
Showers	47 %	99.2 %
Delta rays	1 %	96 %
Michel	1 %	94.7 %
Total		99 %

Network adapted to very sparse data, see paper for details: arXiv:1903.05663



François Drielsma (SLAC)

ML Reconstruction Chain for LArTPCs

Point proposal





Point proposal





François Drielsma (SLAC)

Point proposal





François Drielsma (SLAC)

Fragment clustering





Network predicts 3 things:

- **Embedding**: space in which fragments are spatially separated
- Seediness: likelihood that a voxel is a cluster centroid in embedding space
- Margin: Size of the cluster in embedding space

D. H. Koh



Fragment clustering



François Drielsma (SLAC)



Adjusted Rand Index (ARI)

1.0

0.8

0.6

0.4

0.2

0.0

-0.2

Object clustering

Architecture



CNN struggles with far detached elements such as shower secondaries.

Graphical Neural Networks (GNN) are ideal for this:

• Based on nodes and edges. Features propagate by message passing (MP)



From previous stages:

• Fragmented EM showers



From previous stages:

• Fragmented EM showers

 \mathbf{N} ode features:

- Geometrical features
- CNN autoencoder
- Aggregated embedding of voxels



From previous stages:

• Fragmented EM showers

 \mathbf{N} ode features:

- Geometrical features
- CNN autoencoder
- Aggregated embedding of voxels

Input graph:

• Connect every node with every other node (complete graph)





From previous stages:

• Fragmented EM showers

Node features:

- Geometrical features
- CNN autoencoder
- Aggregated embedding of voxels

Input graph:

• Connect every node with every other node (**complete graph**)

Edge features:

- Displacement vector (+variations)
- Points of closest approach

Network input





Edge classification



At each message passing:

Node update

 $oldsymbol{x}_i' = oldsymbol{\Theta} oldsymbol{x}_i + \sum_{j \in \mathcal{N}(i)} oldsymbol{x}_j \cdot oldsymbol{h}_{\Theta}(oldsymbol{e}_{i,j})$

NNConv: arXiv:1704.02901

• Edge update

$$oldsymbol{e}_{i,j}'=oldsymbol{\phi}_{\Theta}(oldsymbol{x}_i,oldsymbol{x}_j,oldsymbol{e}_{i,j})$$

MetaLayer: arXiv:1806.01261

At each message passing:

• Node update

 $m{x}_i' = m{\Theta} m{x}_i + \sum_{j \in \mathcal{N}(i)} m{x}_j \cdot m{h}_{\Theta}(m{e}_{i,j})$ NNConv: arXiv:1704.02901

• Edge update

 $oldsymbol{e}_{i,j}'=oldsymbol{\phi}_{\Theta}(oldsymbol{x}_i,oldsymbol{x}_j,oldsymbol{e}_{i,j})$

MetaLayer: arXiv:1806.01261

- After n = 3 node+edge updates:
 - Edge binary classification

Target:

- Predict adjacency matrix $A_{ij} = \delta_{g_i,g_j}$ with \boldsymbol{g} the true partition of the set
- Apply cross-entropy loss
 Francois Drielsma (SLAC)



Inference



The network predicts a score matrix \boldsymbol{S} , estimate of the true adjacency matrix \boldsymbol{A}

• How to recover a set partition \hat{g} ?



The network predicts a score matrix $\boldsymbol{S},$ estimate of the true adjacency matrix \boldsymbol{A}

• How to recover a set partition \hat{g} ?

We want the partition \hat{g} that minimizes the CE loss, given $\hat{A}_{ij} = \delta_{\hat{g}_i,\hat{g}_j}$, e.g.

$$\hat{oldsymbol{g}} = \min_{oldsymbol{g} \in G} \mathcal{L}_{CE}(oldsymbol{S}, \hat{oldsymbol{A}}(oldsymbol{g}))$$



The network predicts a score matrix \boldsymbol{S} , estimate of the true adjacency matrix \boldsymbol{A}

• How to recover a set partition \hat{g} ?

We want the partition \hat{g} that minimizes the CE loss, given $\hat{A}_{ij} = \delta_{\hat{g}_i,\hat{g}_j}$, e.g.

$$\hat{oldsymbol{g}} = \min_{oldsymbol{g} \in G} \mathcal{L}_{CE}(oldsymbol{S}, \hat{oldsymbol{A}}(oldsymbol{g}))$$

 \boldsymbol{G} is the set of all possible partitions

- Bell number, huge ($B_{20} \simeq 5 \times 10^{13}$), cannot brute force, how to optimize?
- Build MST on edge scores
- Down select edges until score cannot be improved further







- \mathbf{P} urity $= rac{1}{N} \sum_{i=1}^{n_p} \max_j |c_i \cap t_j|$
 - c_i predicted cluster
 - t_j true cluster with highest count in c_i

Efficiency = $\frac{1}{N} \sum_{i=0}^{n_t} \max_j |c_j \cap t_i|$

- c_j pred. cluster with highest count in t_i
- t_i true cluster







- \mathbf{P} urity $= rac{1}{N} \sum_{i=1}^{n_p} \max_j |c_i \cap t_j|$
 - c_i predicted cluster
 - t_j true cluster with highest count in c_i

Efficiency = $\frac{1}{N} \sum_{i=0}^{n_t} \max_j |c_j \cap t_i|$

- c_j pred. cluster with highest count in t_i
- t_i true cluster
- Adjusted Rand Index (ARI)
 - Measure of overlap of prediction and truth, adjusted for random chance

$$\begin{aligned} \mathsf{RI} &= \frac{a+b}{a+b+c+d} \\ \mathsf{ARI} &= \frac{RI - E(RI)}{1 - E(RI)} \end{aligned}$$





François Drielsma (SLAC)

April 7, 2020 21 / 25



Start identification





Identifying shower starts is a very useful task

- Start point, direction
- Shower matching for π_0 reconstruction

Given a partition \hat{g} :

- Create a complete graph within each pred. group
- Predict which node comes first in time within each pred. group
- CE loss on primary labels, if only 1 true primary in pred. group





François Drielsma (SLAC)

ML Reconstruction Chain for LArTPCs

April 7, 2020 22 / 25

High rate LArTPCs

DUNE ND





Interaction Clustering

Performance





Interaction Clustering

Performance





Interaction Clustering

Performance











ML Reconstruction Chain for LArTPCs:

- Trend in neutrino detection: high-resolution particle imaging
- \bullet Resulting analysis trend: computer vision \rightarrow Machine Learning
- LArTPC images too information rich to be reduced to simple variables in one pass
- Hierarchical feature extraction very successful so far

Areas we will work on not covered in this talk:

- Data vs simulation domain discrepancy
- Error propagation

Please email for more details or if you want to participate !

Back-up slides

Sizes of current and future LArTPCs:

- µBOONE: 100 t (10x2.5x2.5 m³)
- pDUNE: 200t (6x6x6 m³)
- ICARUS: 400 t (2x(20x3x3) m³)
- ArgonCube 2x2: 10t (4x(0.67x0.67x2) m³)
- DUNE-ND: 150t (35x(1x1x3) m³)
- DUNE-FD: 40 kt (4x(12x12x60) m³)

Numbers



Some example numbers (for ICARUS):

- Wire pitch: 3 mm
- Angle between planes: 60°
- Drift field: 500 V/cm
- Drift velocity: $\sim 0.15\,{\rm cm}/\mu{\rm s}$
- TPC time resolution: $0.4 \,\mu s \; (< 1 \, mm)$
- PMT coverage: ${\sim}2\%$
- Scintillation light: 20 % prompt (6 ns), 80 % late (1.5 µs)
- Photon yield: 24000/MeV



Sparse CNNs

Introduction



Submanifold Sparse Convolutions

- 1. Resources waste of dense convolutions on sparse data
- 2. Dilation problem
 - ▶ One nonzero site leads to 3d nonzero sites after 1 convolution
 - How to keep the same level of sparsity throughout the network?



https://arxiv.org/pdf/1711.10275.pdf

François Drielsma (SLAC)

ML Reconstruction Chain for LArTPCs

Sparse CNNs

Introduction



In more details, two new operations:

- Sparse convolutions (SC)
 - Discards contribution of non-active input sites
 - Output site active if at least one input site is active
- Sparse submanifold convolutions (SSC)
 - Output size = Input size
 - Output site active iff center of receptive field active
 - Only compute features for active output sites

 $\bigcirc \bigcirc$

https://arxiv.org/pdf/1711.10275.pdf

Point proposal





PPN outputs voxel location, position within voxel and point class

Three components to the point proposal loss:

• Pixel classification loss at each of three depth (pixel contains point or not)

$$\mathcal{L}_{\mathsf{class},i} = \frac{1}{N_i} \sum_{k=1}^{N_i} y_k \log(p_k) + (1 - y_k) \log(1 - p_k)$$

• L^1 distance from true point at highest resolution on active voxels

$$\mathcal{L}_{\mathsf{dist}} = rac{1}{N_3^*} \sum_{k=1}^{N_3^*} \min_j ||ec{p_i} - ec{q_j}||$$

• Particle type loss at highest resolution on active voxels

$$\mathcal{L} = \frac{1}{N_3^*} \sum_{k=1}^{N_3^*} \sum_{c} y_{k,c} \log(p_{k,c})$$

Fragment clustering



Instance Segmentation:

• Three component loss: pull together points that belong to the same cluster, keep distance between clusters, and regularization

Loss

$$L = \alpha L_{var} + \beta L_{dist} + \gamma L_{reg},$$

$$L_{var} = \frac{1}{C} \sum_{c=1}^{C} \frac{1}{N_c} \sum_{i=1}^{N_c} [\max(0, \|\mu_c - x_i\| - \delta_v)]^2$$

$$L_{dist} = \frac{1}{C(C-1)} \sum_{\substack{c_A, c_B = 1 \\ c_A \neq c_B}}^{C} [\max(0, 2\delta_d - \|\mu_{c_A} - \mu_{c_B}\|)]^2$$

$$L_{reg} = \frac{1}{C} \sum_{c=1}^{C} \|\mu_c\|$$

arXiv:1708.02551

Domain discrepancies



What can we do about imperfect simulation ?

• Issue: the signal distribution learned by the algorithm may be different in two domains!

GAN

- Mitigation techniques in ML domain ?
 - Can try CNN to locate where it is
 - Can try CNN to fix the discrepancy
 - Can try a training technique to minimize the effect



Maximize the loss to discriminate data vs. simulation, feature extractors are penalized to key on simulation specific information

Domain-Adversarial Training of Neural Networks: J. Mach. Learn. Res. 17 (2016)

Open Source Development Highlights





DeepLearnPhysics: Collaboration for ML technique R&D

- Open simulation sample (used throughout this talk)
 - Open real data ? Soon ! (3D prototype R&D at SLAC)
- Open source container (Singularity)
- Open source code (GitHub)
 - All the code used to make this talk is available
- \rightarrow Reproducible results !
 - Readers have reproduced arXiv:1903.05663

