


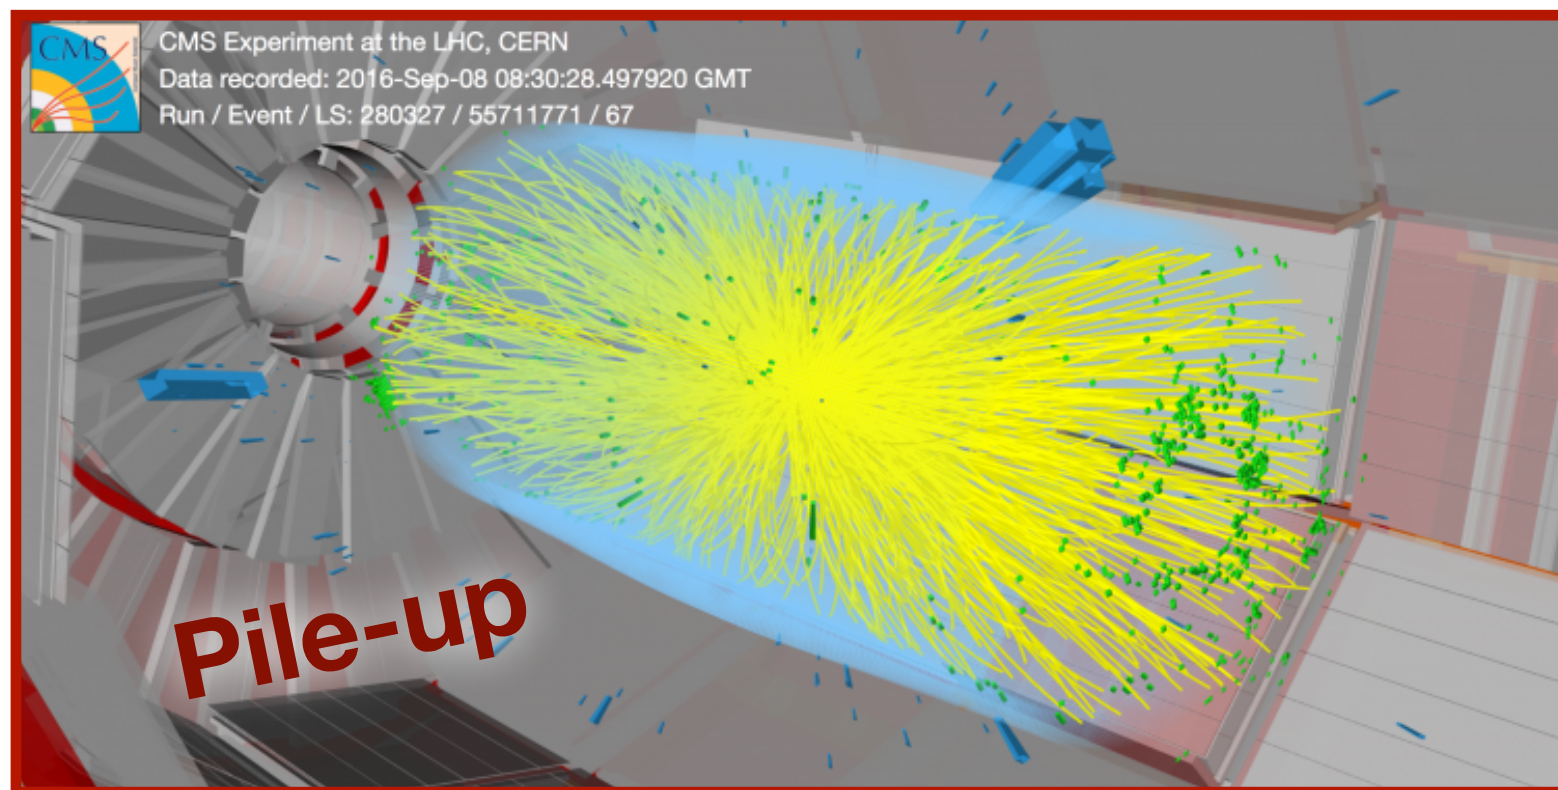
Graph Neural Network reconstruction in HGCAL

Thea Aarrestad, Sitong An, Gianluca Cerminara, Javier Duarte, Sergei Gleyzer, Dejan Golubovic, Shamik Gosh, Lindsey Gray, Phil Harris, Yutaro Iiyama, Jan Kieseler, **Thomas Klijnsma**, Kenneth Long, Jennifer Ngadiuba, Gerrit van Onsem, Joosep Pata, Kevin Pedro, Maurizio Pierini, Shah Rukh Qasim, Marcel Rieger, Sioni Summers, Mary Touranakou, Nhan Tran, Oleksander Viazlo, Kinga Wozniak

7 April 2020



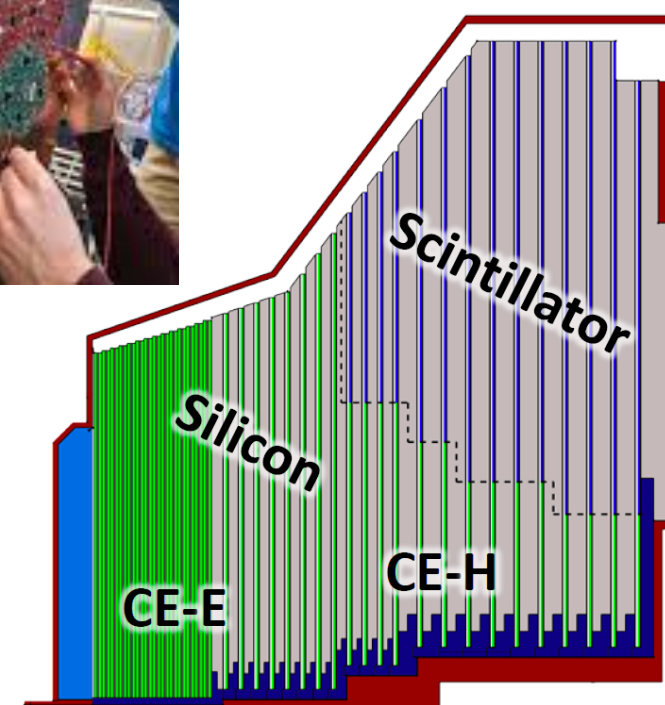

Motivation



- Traditional algorithms scale **combinatorially**
 - **Explodes** going from 32 to **200** PU interactions
- A neural network could reconstruct in **constant time**


About HGCAL

- One event in HGCAL is a large **set of hits** (**x, y, z, E, t**) (i.e. feature space is 5D)
- Dimensionality + geometry not well suited for the CNNs in industry (rectangular, non-sparse, 3 colors)
- Showers and tracks are pretty naturally represented by a graph -> **GNNs**





- **PointNet:** one of the pioneering deep neural network approach to work with point clouds
- Limited use of neighborhood information

Semantic segmentation




Clustering





- **PointNet:** one of the pioneering deep neural network approach to work with point clouds
- Limited use of neighborhood information

Semantic segmentation



Clustering




- **EdgeConv:** Update $x_i \rightarrow x'_i$ by using **edge features**

- i.e. learned features of the edges that connects x_i with its neighbors
- Still independent of ordering of points, but uses **local geometry**
- '**Convolutional**' as the operation is applied point by point to obtain \mathbf{x}'

Dynamic Graph Convolutional NN

[1801.07829]



$$\mathbf{x}'_i = \bigoplus_{j:(i,j) \in \mathcal{E}} h_{\Theta}(\mathbf{x}_i, \mathbf{x}_j)$$

$h_{\Theta}(\mathbf{x}_i, \mathbf{x}_j) = h_{\Theta}(\mathbf{x}_i)$ No neighborhood info (only global)

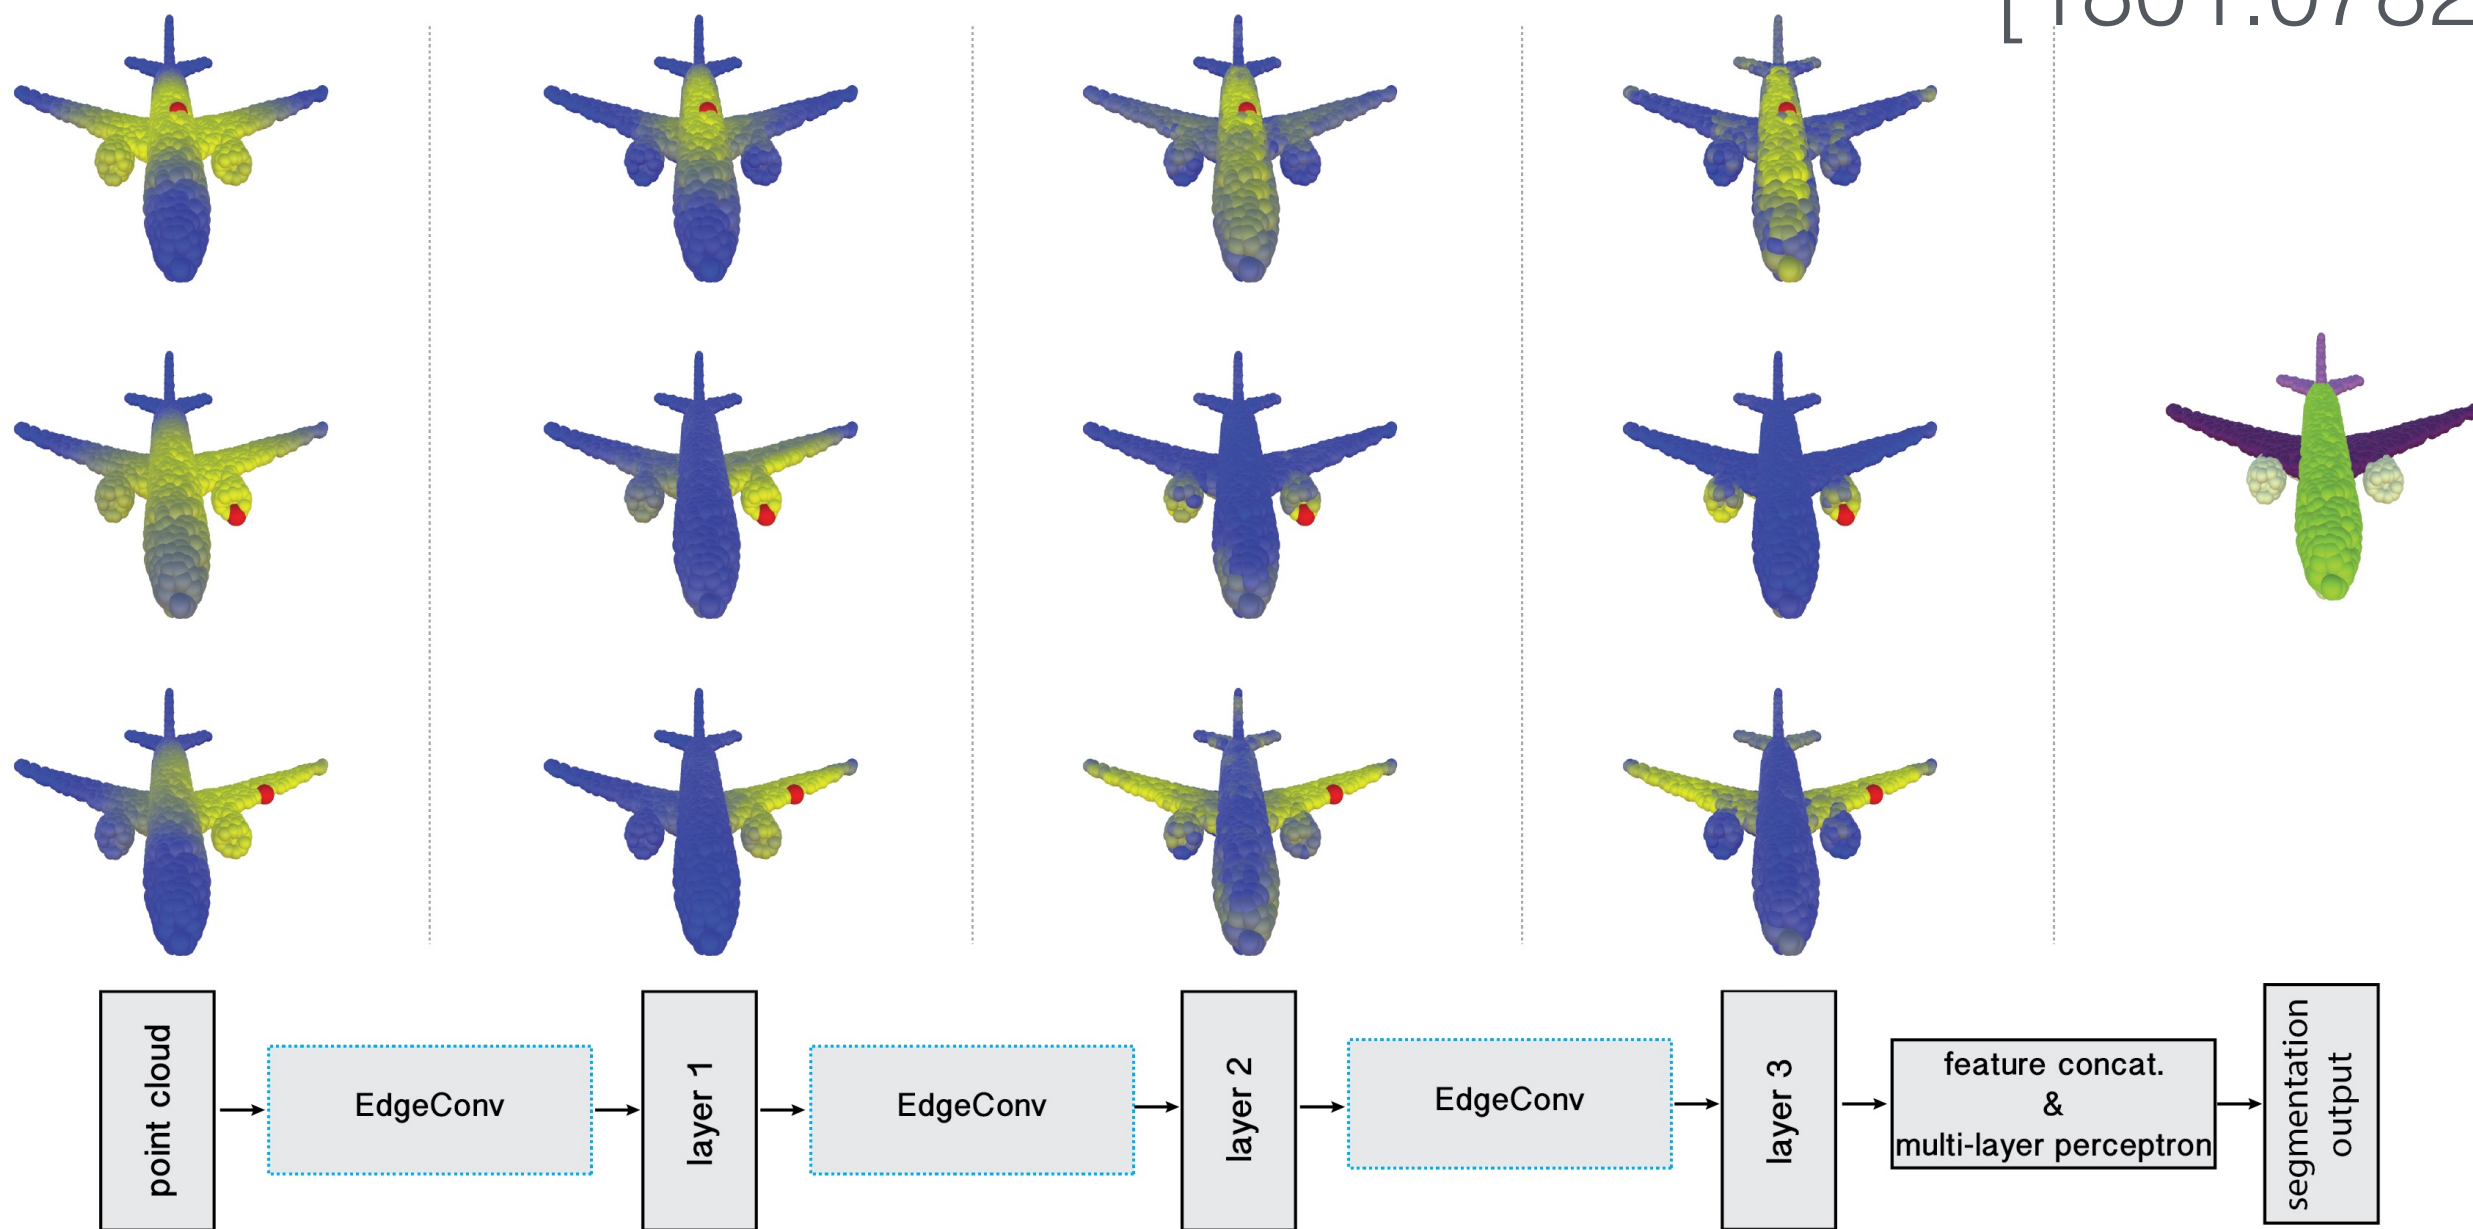
$h_{\Theta}(\mathbf{x}_i, \mathbf{x}_j) = h_{\Theta}(\mathbf{x}_j - \mathbf{x}_i)$ Only local information

$h_{\Theta}(\mathbf{x}_i, \mathbf{x}_j) = \bar{h}_{\Theta}(\mathbf{x}_i, \mathbf{x}_j - \mathbf{x}_i)$ Combination of both

- **Dynamic:** Redo kNN after every update
 - The connectivity matrix changes after every update


Dynamic Graph Convolutional NN

[1801.07829]





Dynamic Graph Convolutional NN

[1801.07829]




CMS Phase-2 Simulation Preliminary


truth





Graphical example




Generate edges
between nodes
(preprocessing, kNN)




Calculate
edge features







Graphical example



Generate edges
between nodes
(preprocessing, kNN)




Calculate
edge features




Update the
node features
using the
edge features

'Message-passing' takes place
Important information from a
neighbor spreads via x_i to
another neighbor in the next
iteration






Graphical example




Generate edges
between nodes
(preprocessing, kNN)




Calculate
edge features



(repeat)




Classify into 'true'
edges vs 'false' edges
(ML, 2 edge features)




Update the
node features
using the
edge features






Graphical example



Generate edges
between nodes
(preprocessing, kNN)




Calculate
edge features



(repeat)





Classify into 'true'
edges vs 'false' edges
(ML, 2 edge features)



Update the
node features
using the
edge features



↑
DGCNN
would *re-kNN*
at this point

Proof of concept: Single pion



Proof of concept: Single photon




GNNs seem to do well for single particles!



- The real challenge: **Many-particle events**



Proof of concept: tau decays

- Most recent development: **clustering** the output of the GNN into **multiple different-type particles** seems to do a good job





Proof of concept: tau decays

- Most recent development: **clustering** the output of the GNN into **multiple different-type particles** seems to do a good job




'Bare' energy prediction (hadronic vs em) looks promising


To be regressed...




Proof of concept: tau decays






- Example event display:
Clear particle-like clusters are constructed
- Clusters are separated by EM ●, Hadronic ▲, MIP ✚, and noise (not plotted)



- Work in progress:
 - Pileup
 - Need better 'truth' definition (data prep)
 - Integration into CMSSW (longterm solution probably PyTorch in CMSSW)
 - Hardware acceleration
- Many problems in common with the PF effort





```

16 class EdgeNetWithCategories(nn.Module):
17     def __init__(self, input_dim=3, hidden_dim=8, output_dim=4, n_iters=1, aggr='add',
18                 norm=torch.tensor([1./500., 1./500., 1./54., 1/25., 1./1000.])):
19         super(EdgeNetWithCategories, self).__init__()
20
21         self.datanorm = nn.Parameter(norm)
22
23         start_width = 2 * (hidden_dim + input_dim)
24         middle_width = (3 * hidden_dim + 2*input_dim) // 2
25
26         self.n_iters = n_iters
27
28         self.inputnet = nn.Sequential(
29             nn.Linear(input_dim, 2*hidden_dim),
30             nn.Tanh(),
31             nn.Linear(2*hidden_dim, 2*hidden_dim),
32             nn.Tanh(),
33             nn.Linear(2*hidden_dim, hidden_dim),
34             nn.Tanh(),
35         )
36
37         self.edgenetwork = nn.Sequential(nn.Linear(2*n_iters*hidden_dim, 2*hidden_dim),
38                                         nn.ELU(),
39                                         nn.Linear(2*hidden_dim, 2*hidden_dim),
40                                         nn.ELU(),
41                                         nn.Linear(2*hidden_dim, output_dim),
42                                         nn.LogSoftmax(dim=-1),
43         )
44
45         for i in range(n_iters):
46             convnn = nn.Sequential(nn.Linear(start_width, middle_width),
47                                   nn.ELU(),
48                                   #nn.Dropout(p=0.5, inplace=False),
49                                   nn.Linear(middle_width, hidden_dim),
50                                   nn.ELU()
51             )
52             setattr(self, 'nodenetwork%d' % i, EdgeConv(nn=convnn, aggr=aggr))
53
54         def forward(self, data):
55             row,col = data.edge_index
56             x_norm = self.datanorm * data.x
57             H = self.inputnet(x_norm)
58             H = getattr(self, 'nodenetwork0')(torch.cat([H, x_norm], dim=-1), data.edge_index)
59             H_cat = H
60             for i in range(1,self.n_iters):
61                 H = getattr(self, 'nodenetwork%d' % i)(torch.cat([H, x_norm], dim=-1), data.edge_index)
62                 H_cat = torch.cat([H, H_cat], dim=-1)
63             return self.edgenetwork(torch.cat([H_cat[row],H_cat[col]],dim=-1)).squeeze(-1)

```

See [github](#)

re constructed

M •, Hadronic
tted)


on (data prep)

longterm solution

with the PF


Intermezzo: The 'history' issue

- Secondary particles are recorded as if they are part of their parent




- This leads to showers that *look* a lot like their type (e.g. photon), but are truth-tagged as their parent-type (e.g. muon)
- This throws off any deep learning

Work in progress, bugs to solve



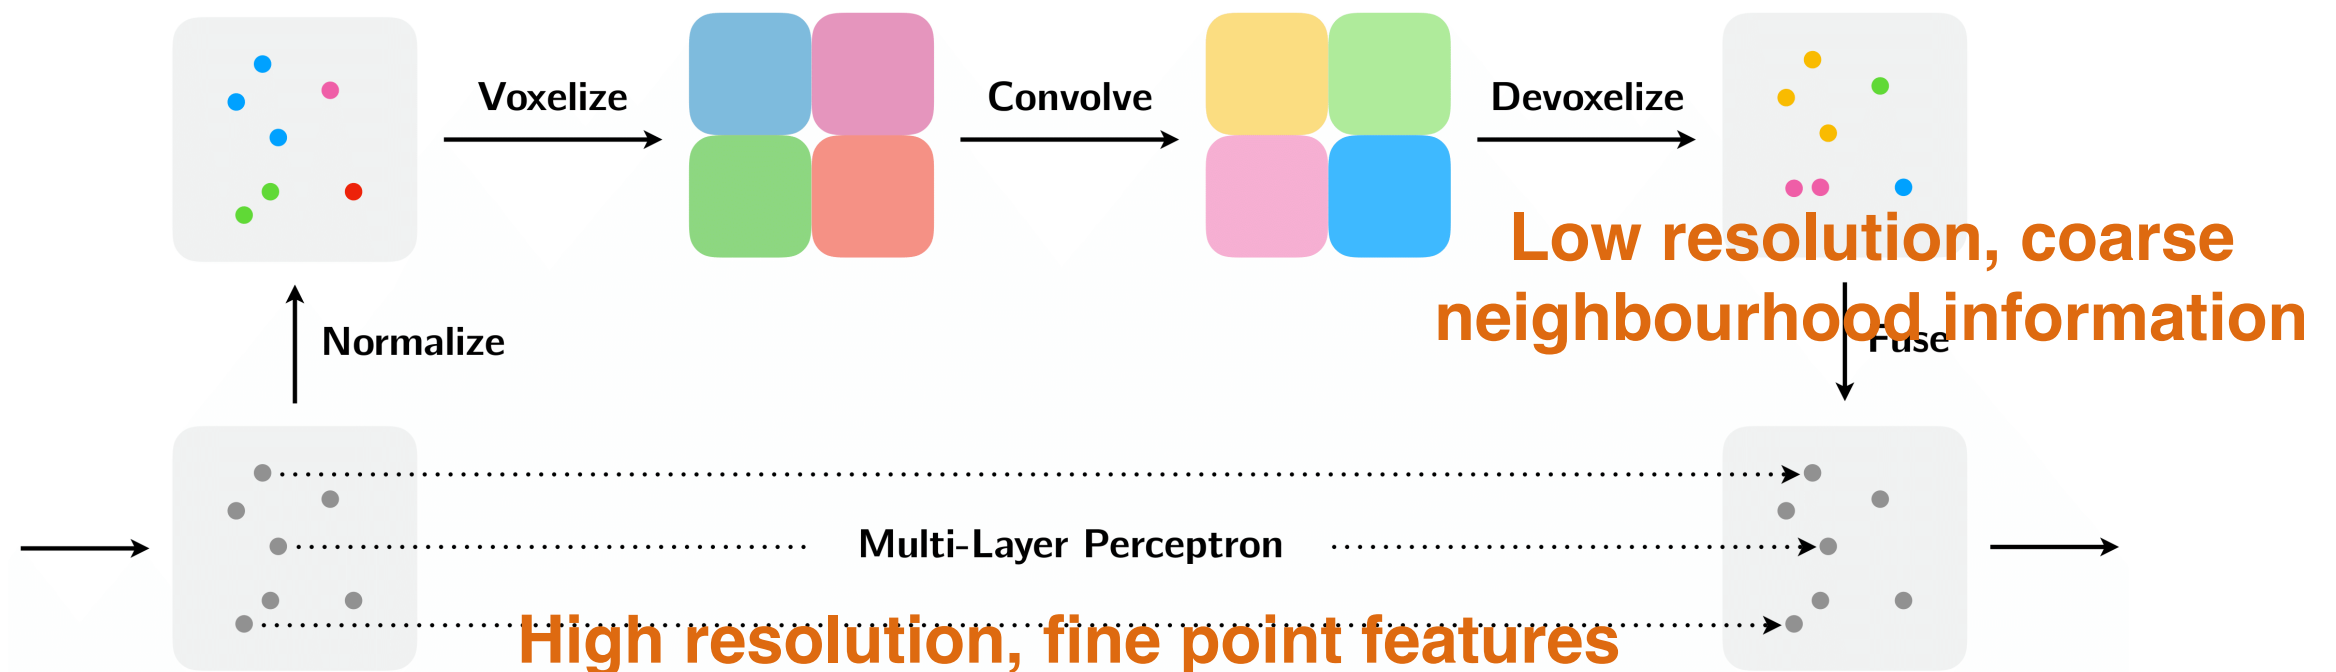
- Presence of a buggy photon

- Still incorrectly labeled hits
- Planning on getting involved with this software ourselves




Point-Voxel CNN

- Use **fine-grained point information** in a small dense layer, **and coarse-grained voxelization** in convolutional layers
 - Mitigate effects from poor memory locality (pure point cloud) and huge memory consumption (pure voxelization)
- In principle a **point-cloud network, not graph based**
 - Working with the authors to get to an edge-classification network



Point-Voxel CNN

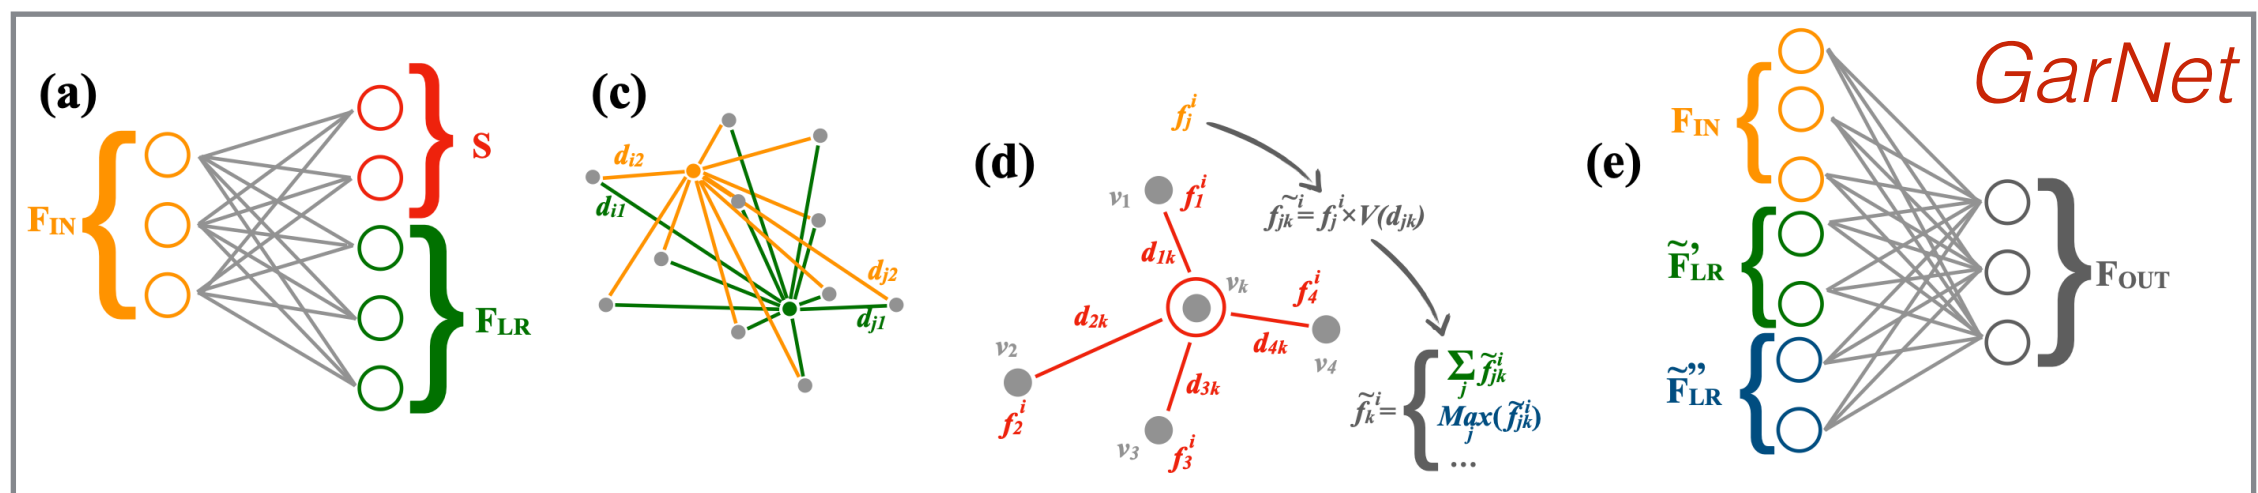



- Does pretty good job out of the box (arch optimized for classifying office furniture)
- Needs further study into
 - Instance segmentation
 - Inference speed
 - Using 'better' PVCNN called SPVNAS

Confusion matix


	Noise	HAD	EM	MIP
Noise	0.9984	0.0268	0.0275	0.0164
HAD	0.0009	0.8004	0.1487	0.0680
EM	0.0007	0.1715	0.8233	0.0089
MIP	0.0001	0.0014	0.0005	0.9068

- **DGCNN** uses *large* amount of *memory* and keeping *inference time* under control is a challenge
- **GravNet/GarNet** greatly reduces computational needs
 - Split coordinate and feature space






CMS Phase-2 Simulation Preliminary

truth




CMS Phase-2 Simulation Preliminary

predicted



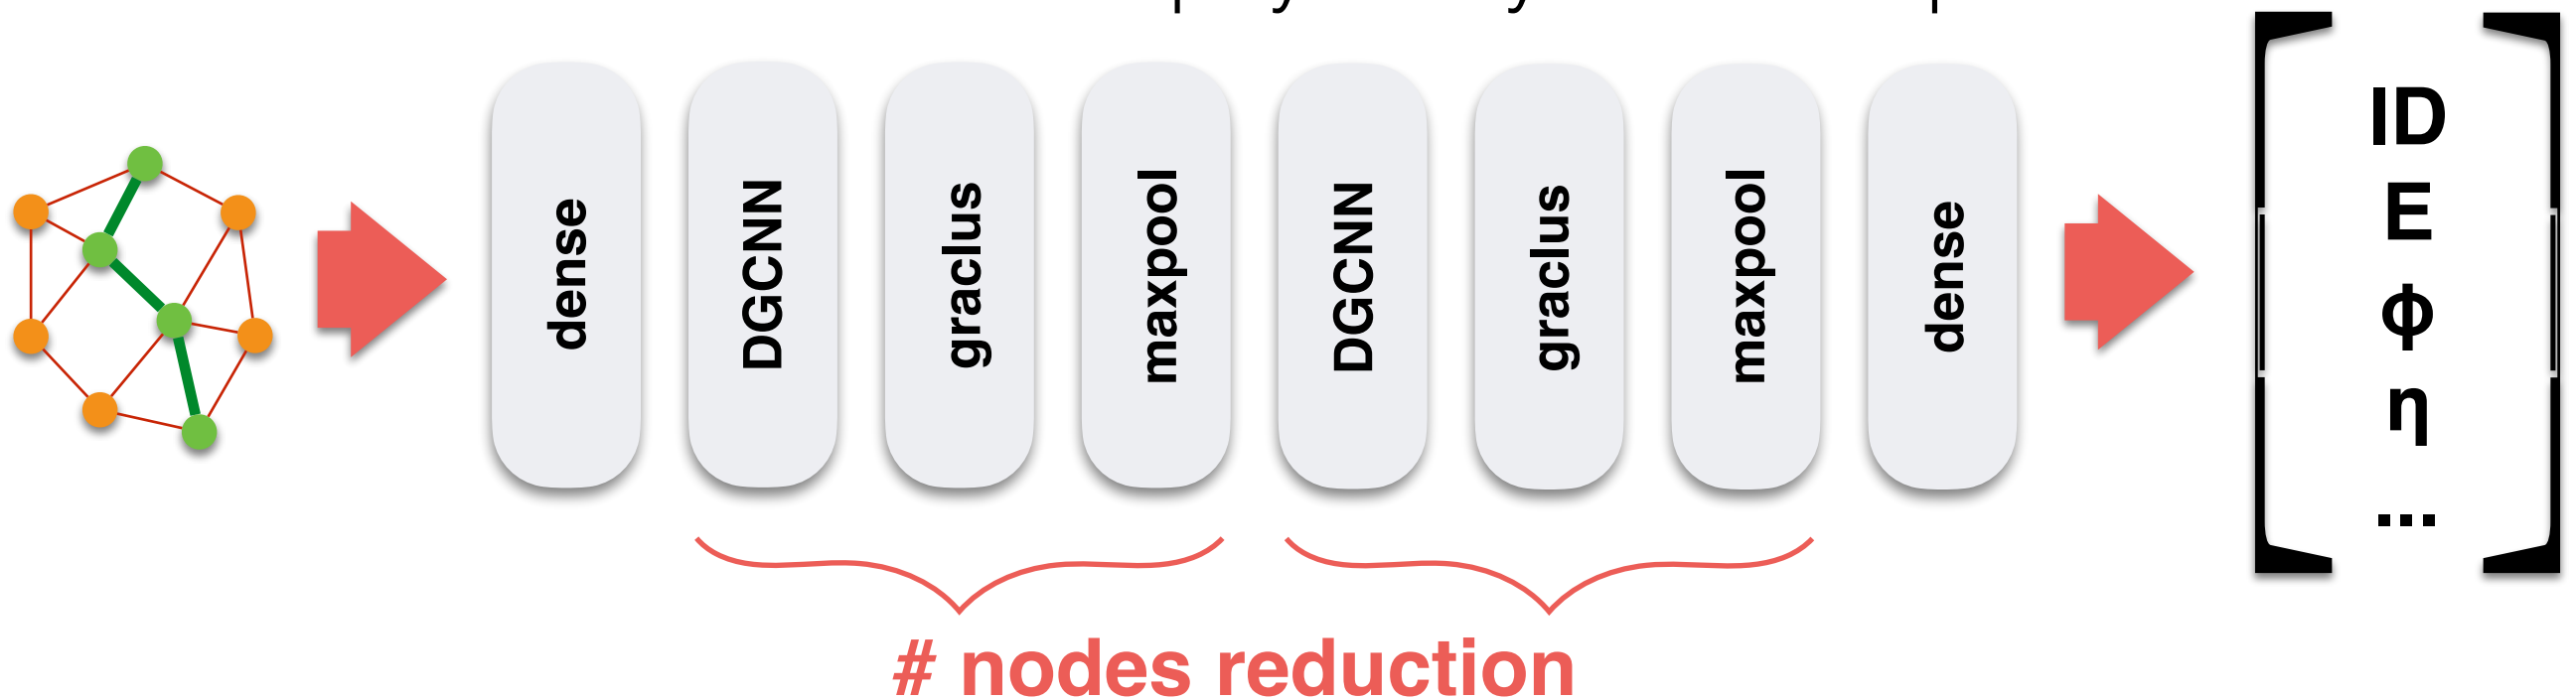
Regression

- In the end, want to get back meaningful properties of cluster




ID, energy, 4-mom, ...

- '**Dynamic Reduction Network**' capable of taking an unordered set and reduce to a vector of physically relevant quantities



- Because of *EdgeConv*, **learns** how to use **organization and weighting** of input data to regress to physics
- Gets 99.55% test accuracy on MNIST (#19-21 on leaderboard)

Regression



ID. energy. 4-

```
def forward(self, data):
    data.x = self.datanorm * data.x
    data.x = self.inputnet(data.x)

    data.edge_index = to_undirected(knn_graph(data.x, self.k, data.batch, loop=False, flow=self.edgeconv1.flow))
    data.x = self.edgeconv1(data.x, data.edge_index)

    weight = normalized_cut_2d(data.edge_index, data.x)
    cluster = graclus(data.edge_index, weight, data.x.size(0))
    data.edge_attr = None
    data = max_pool(cluster, data)

    data.edge_index = to_undirected(knn_graph(data.x, self.k, data.batch, loop=False, flow=self.edgeconv2.flow))
    data.x = self.edgeconv2(data.x, data.edge_index)

    weight = normalized_cut_2d(data.edge_index, data.x)
    cluster = graclus(data.edge_index, weight, data.x.size(0))
    x, batch = max_pool_x(cluster, data.x, data.batch)

    x = global_max_pool(x, batch)

    return self.output(x).squeeze(-1)
```

See github


- Gets 99.55% test accuracy on MNIST (#19-21 on leaderboard)

What's “graclus”????

● Greedy, popularity based graph clustering algorithm

- A node is more popular if it has more and close-by neighbors
- Greedy means that the most popular nodes seed clusters and accumulate neighbors into themselves

Basic idea:



Hadrons - Qualitative Results


● Network learns quickly

- 1 epoch ~3 minutes
- loss settled by epoch 20
- no signs of overtraining


● Clear improvement in first epochs

- avg. scale centers
- width and outliers reduced

Epoch 1





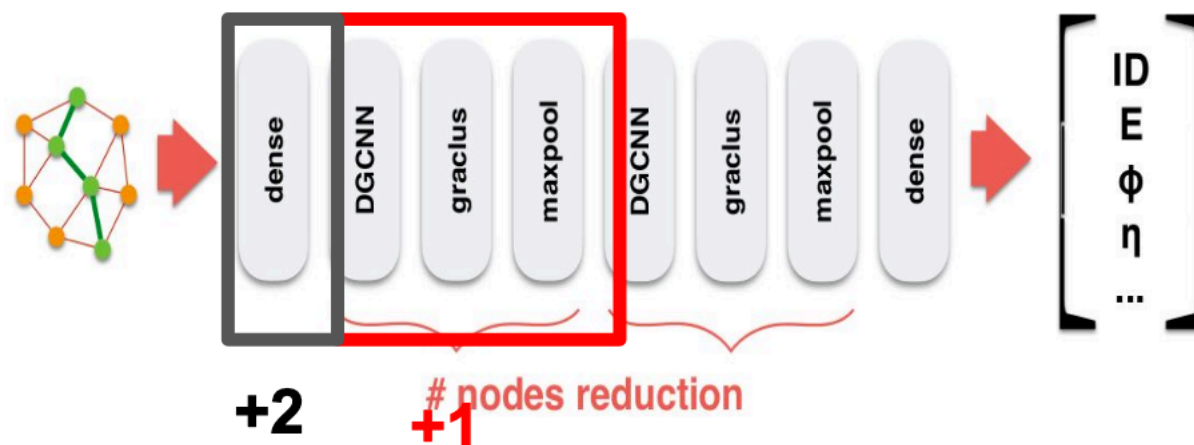
Epoch 9



Hadrons - Best Model Variant So Far




- More ability to encode fluctuations of low energy showers
 - 124% stochastic term
 - 5.3% constant term
 - to improve
 - Scale rather flat (for hadrons)

- Good enough for now
 - next step: flatten scale




S. Ghosh

- Method to perform multiple pattern recognition tasks in one operation
- **Input** is a set of related pixels/points/vertices/edges/...
- **Output** is a **number of objects** (e.g. number of particles in an event) each carrying their high-level object properties (e.g. their four-momenta)



CMS Phase-2 Simulation Preliminary
truth







- Eliminates need for bounding boxes (which don't work well for sparse objects)
- Method generalizes for point clouds, may use to do clustering + regression for HGCal
- Results on this to be produced very soon

Conclusion

- Exciting early-stage results using GNNs for reconstruction in HGCAL
 - Collimated, multi-particle reconstruction within grasp
 - Multiple available methods, progressing on multiple fronts
- Other challenges:
 - Onwards to including pileup
 - Truth definition in dataset
 - Hardware acceleration
 - Integration into CMSSW

Backup

Graphical example



Graphical example

Actual par


Update
node fe
using
edge fe

```


13 class EdgeNet2(nn.Module):
14     def __init__(self, input_dim=3, hidden_dim=8, output_dim=1, n_iters=1, aggr='add'):
15         super(EdgeNet2, self).__init__()
16         convnn = nn.Sequential(nn.Linear(2*(hidden_dim + input_dim), (3*hidden_dim + 2*input_dim) // 2),
17                                nn.ReLU(),
18                                nn.Dropout(),
19                                nn.Linear((3*hidden_dim + 2*input_dim) // 2, hidden_dim),
20                                nn.ReLU())
21     )
22
23     self.n_iters = n_iters
24
25     self.inputnet = nn.Sequential(
26         nn.Linear(input_dim, hidden_dim),
27         nn.BatchNorm1d(hidden_dim),
28         nn.Tanh()
29     )
30
31     self.edgenetwork = nn.Sequential(nn.Linear(2*(n_iters*hidden_dim+input_dim), output_dim),
32                                      nn.Sigmoid())
33
34     self.nodenetwork = EdgeConv(nn=convnn, aggr=aggr)
35
36     def forward(self, data):
37         X = data.x
38         H = self.inputnet(X)
39         data.x = torch.cat([H, X], dim=-1)
40         H_cat = X
41         for i in range(self.n_iters):
42             H = self.nodenetwork(data.x, data.edge_index)
43             H_cat = torch.cat([H, H_cat], dim=-1)
44             data.x = torch.cat([H, X], dim=-1)
45             row, col = data.edge_index
46             return self.edgenetwork(torch.cat([H_cat[row], H_cat[col]], dim=-1)).squeeze(-1)

```

See github




me
(features)
is example




Regression


- Starting to produce physics
 - First ever training of this regression in HEP on only **600** events, 200 testing events, and 14 epochs:




- Soon training with better statistics, but for the limited statistics already does a reasonable job
- Same regression method could be useful for other *unordered set* --> *regressed values* problems

- **PointNet:** one of the pioneering deep neural network approach to work with point clouds
- **Invariance under permutation of points:** Apply a *symmetrizing function* on the point cloud that yields the same output independent of order
- **Invariance under rotation/translation**
- Need both **local** geometry and **global** semantics to do object classification
 - PointNet asks "what is so different about a point far away", but not "what is similar about points near me"





- Update $x_i \rightarrow x'_i$ by using **edge features**
 - i.e. learned features of the edges that connects x_i with its neighbors
 - Still independent of ordering of points, but uses **local geometry**
 - '**Convolutional**' as the operation is applied point by point to obtain \mathbf{x}'
- Calculate edge features simply with e.g. a MLP which takes the node features as input



Can set dim
of edge
feature vector

$$\mathbf{x}'_i = \bigoplus_{j:(i,j) \in \mathcal{E}} h_{\Theta}(\mathbf{x}_i, \mathbf{x}_j)$$