

Classification Baseline

Matt Zhang | June 08, 2018



Classification Bug

The classification architecture we used in NIPS had 4 hidden layers. It turns out that all 4 layers were constrained to have the same weight. This has been fixed and resulted in slightly better accuracy using NIPS architecture.

Before, we showed that scanning over `n_hidden_layers` had little effect. Due to the discovery of this bug, this result is obviously called into question.

```
class Classifier_Net(nn.Module):
    def __init__(self, hiddenLayerNeurons, nHiddenLayers, dropoutProb, windowSize):
        super().__init__()
        self.windowSize = windowSize
        self.nHiddenLayers = nHiddenLayers
        self.input = nn.Linear(windowSize * windowSize * 25, hiddenLayerNeurons)
        self.hidden = nn.Linear(hiddenLayerNeurons, hiddenLayerNeurons)
        self.hidden.cuda()
        self.dropout = nn.Dropout(p = dropoutProb)
        self.dropout.cuda()
        self.output = nn.Linear(hiddenLayerNeurons, 2)
    def forward(self, x, _):
        lowerBound = 26 - int(math.ceil(self.windowSize/2))
        upperBound = lowerBound + self.windowSize
        x = x[:, lowerBound:upperBound, lowerBound:upperBound]
        x = x.contiguous().view(-1, self.windowSize * self.windowSize * 25)
        x = self.input(x)
        for _ in range(self.nHiddenLayers-1):
            x = F.relu(self.hidden(x))
            x = self.dropout(x)
        x = F.softmax(self.output(x), dim=1)
        return x
```

NIPS architecture (when windowSize == 25)

```
class Classifier_Net(nn.Module):
    def __init__(self, hiddenLayerNeurons, nHiddenLayers, dropoutProb, windowSize):
        super().__init__()
        self.windowSize = windowSize
        self.nHiddenLayers = nHiddenLayers
        self.input = nn.Linear(windowSize * windowSize * 25, hiddenLayerNeurons)
        self.hidden = [None] * self.nHiddenLayers
        self.dropout = [None] * self.nHiddenLayers
        for i in range(self.nHiddenLayers):
            self.hidden[i] = nn.Linear(hiddenLayerNeurons, hiddenLayerNeurons)
            self.hidden[i].cuda()
            self.dropout[i] = nn.Dropout(p = dropoutProb)
            self.dropout[i].cuda()
        self.output = nn.Linear(hiddenLayerNeurons, 2)
    def forward(self, x, _):
        lowerBound = 26 - int(math.ceil(self.windowSize/2))
        upperBound = lowerBound + self.windowSize
        x = x[:, lowerBound:upperBound, lowerBound:upperBound]
        x = x.contiguous().view(-1, self.windowSize * self.windowSize * 25)
        x = self.input(x)
        for i in range(self.nHiddenLayers-1):
            x = F.relu(self.hidden[i](x))
            x = self.dropout[i](x)
        x = F.softmax(self.output(x), dim=1)
        return x
```

fixed architecture

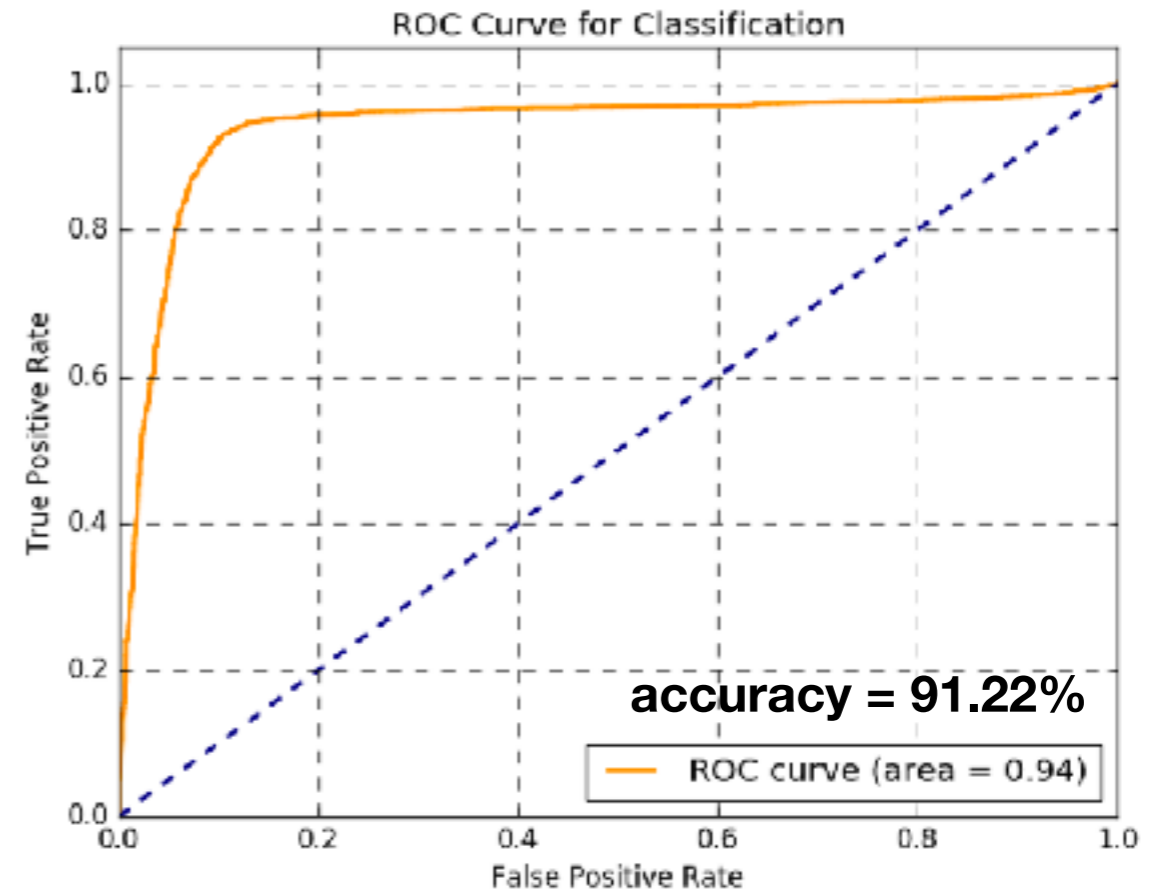
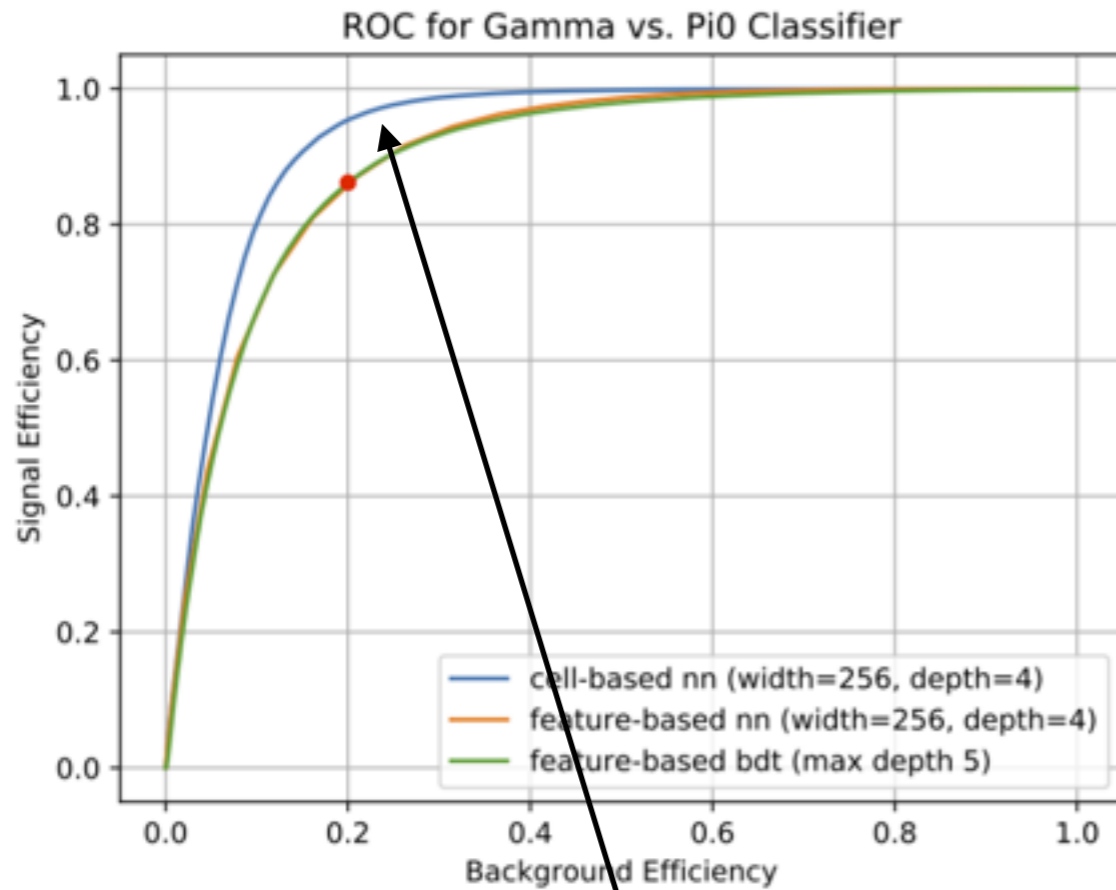
Baseline Classification

Using a window size of 25x25 on new fixed-angle samples with energy between 50-70 GeV, we get an analogue of the data used in the NIPS paper.

When using the original buggy architecture, we got very slightly worse results (sorry, overwrote the plot). With the fixed architecture, our accuracy is better by about 4 percent.

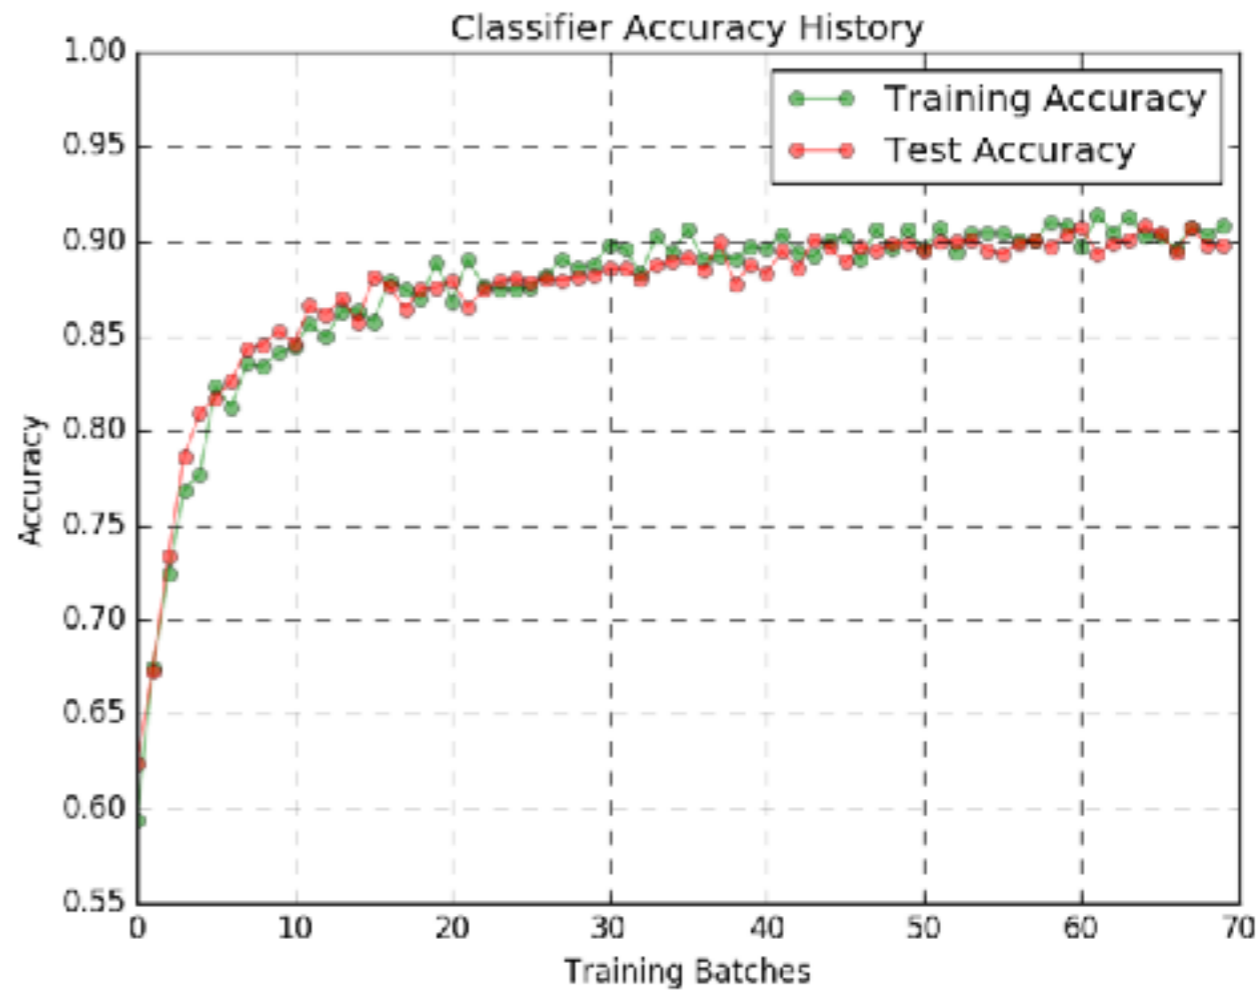

NIPS paper

new samples



Model	γ vs. π^0				e vs. π			
	acc.	AUC	$\Delta\epsilon_{\text{sig}}$	ΔR_{bkg}	acc.	AUC	$\Delta\epsilon_{\text{sig}}$	ΔR_{bkg}
BDT	83.1%	89.8%	-	-	93.8%	98.0%	-	-
DNN (features)	82.8%	90.2%	0.9%	0.95	93.6%	98.0%	-0.1%	0.95
DNN (cells)	87.2%	93.5%	9.4%	1.63	99.4%	99.9%	4.9%	151


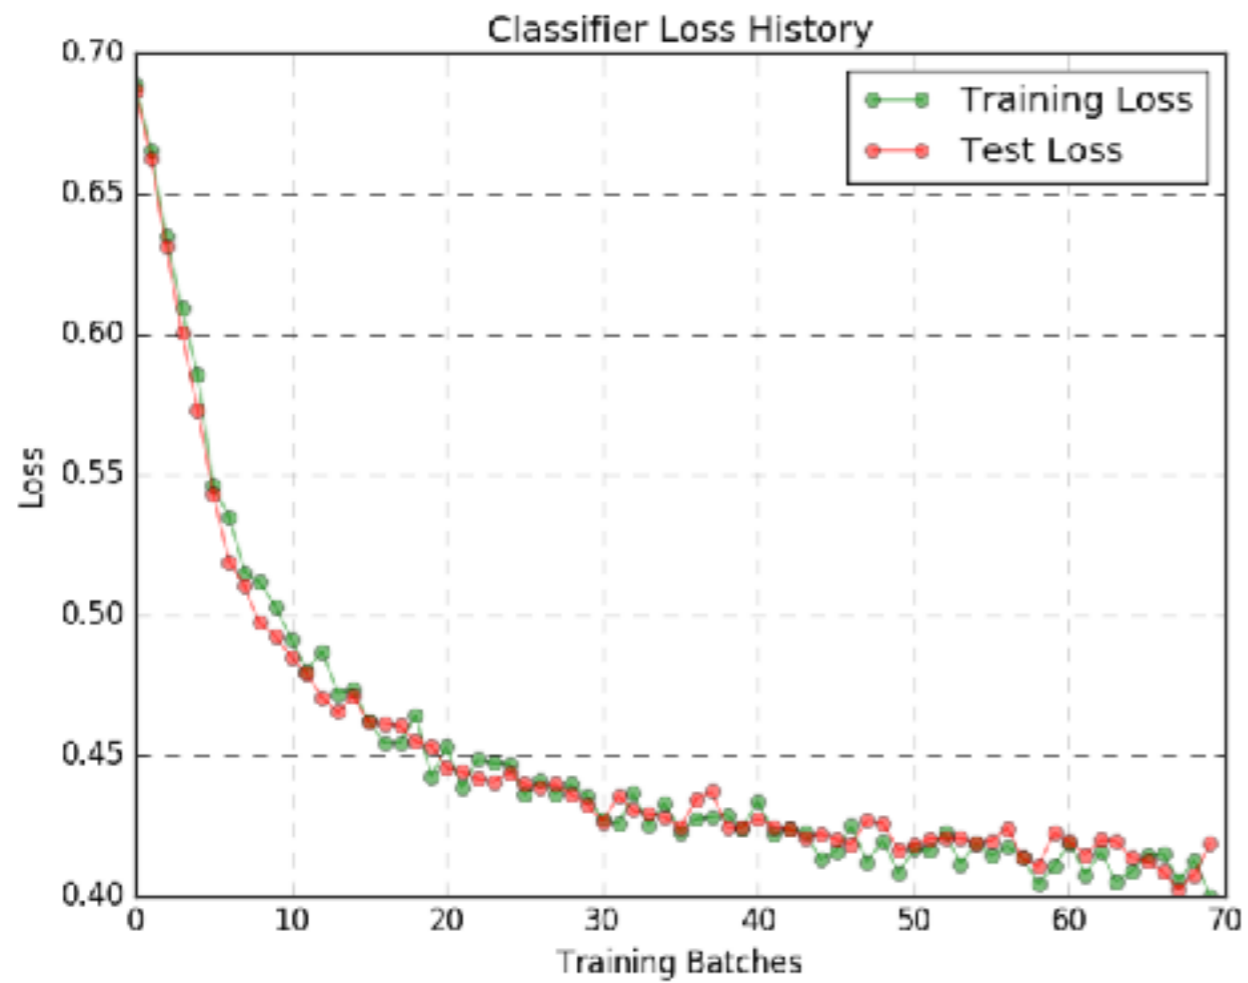
Table 1: Performance parameters for BDT and DNN classifiers.



Baseline w/ Expanded Window

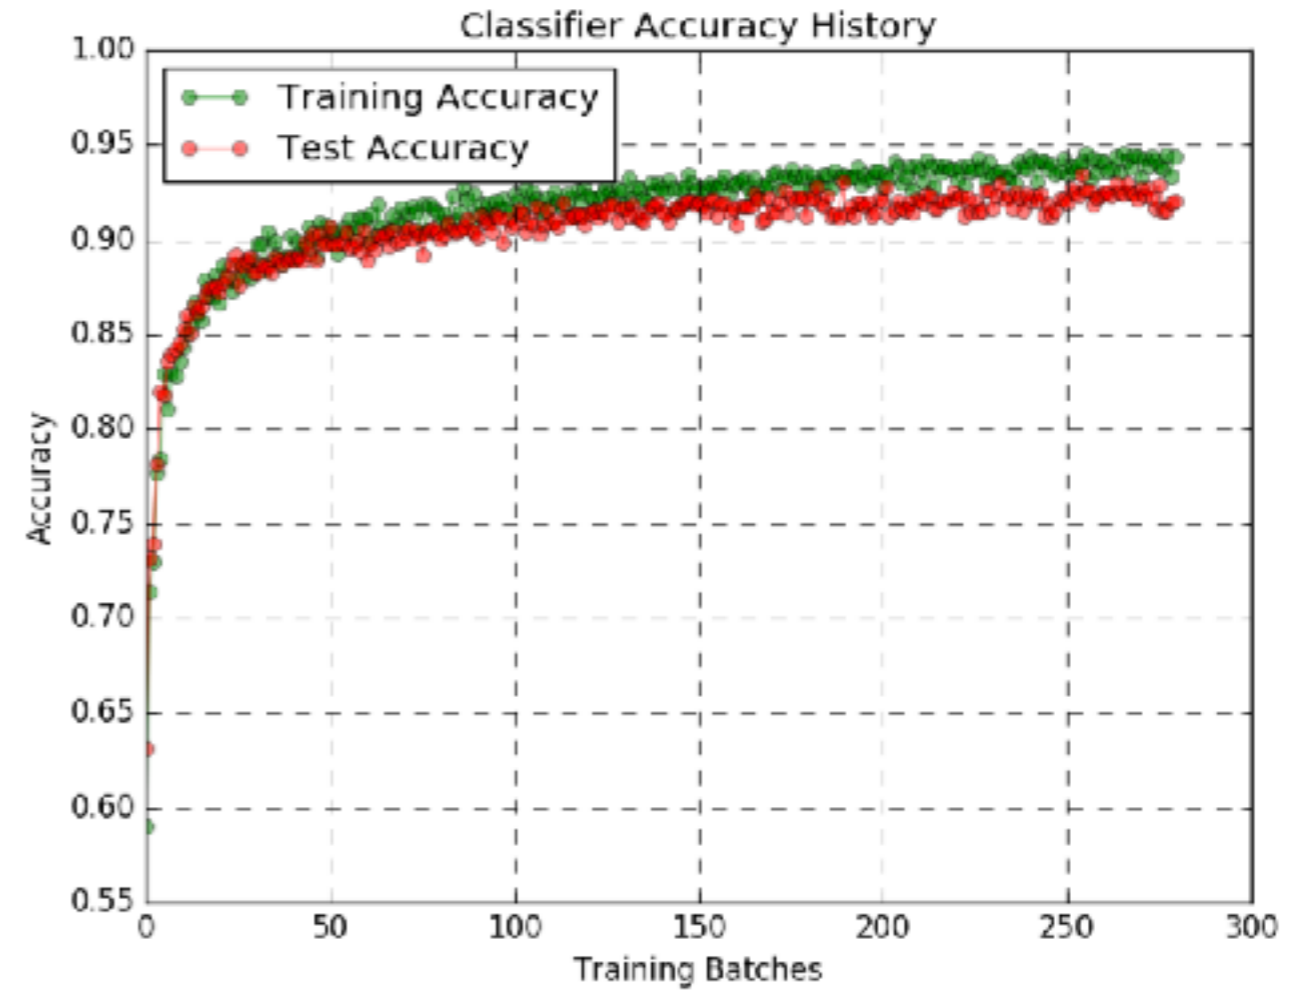

Using a window size of 51x51 on new fixed-angle samples with energy between 50-70 GeV, the same as before except with the larger window.

The results are maybe slightly better than the 25x25 window. Doesn't look like the larger window is either hurting or helping much. It's probably best to keep the larger window since it does contain some info.



Expanded Window, Longer Train

Now training for 20 epochs instead of 5.



Energy-Range Classification

Still using the NIPS architecture with a window size of 25x25, we now use all of the new fixed-angle samples (instead of just those with energy between 50-70 GeV). Note there are more samples, so it takes longer to get through 5 epochs of training.

We also show the results of various hyperparameter scans.

