

PAPER • OPEN ACCESS

Deep Neural Networks for Physics Analysis on low-level whole-detector data at the LHC

To cite this article: Wahid Bhimji *et al* 2018 *J. Phys.: Conf. Ser.* **1085** 042034

View the [article online](#) for updates and enhancements.



IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

Deep Neural Networks for Physics Analysis on low-level whole-detector data at the LHC

Wahid Bhimji¹, Steven Andrew Farrell¹, Thorsten Kurth¹, Michela Paganini^{1,2}, Prabhat¹, Evan Racah¹

¹ Lawrence Berkeley National Laboratory, Berkeley, CA 94720 USA

² Department of Physics, Yale University, New Haven, CT 06520, USA

E-mail: wbhimji@lbl.gov

Abstract. There has been considerable recent activity applying deep convolutional neural nets (CNNs) to data from particle physics experiments. Current approaches on ATLAS/CMS have largely focussed on a subset of the calorimeter, and for identifying objects or particular particle types. We explore approaches that use the entire calorimeter, combined with track information, for directly conducting physics analyses: i.e. classifying events as known-physics background or new-physics signals.

We use an existing RPV-Supersymmetry analysis as a case study and explore CNNs on multi-channel, high-resolution sparse images: applied on GPU and multi-node CPU architectures (including Knights Landing (KNL) Xeon Phi nodes) on the Cori supercomputer at NERSC.

We compare statistical performance of our approaches with selections on high-level physics variables from the current physics analyses, and shallow classifiers trained on those variables. We also compare time-to-solution performance of CPU (scaling to multiple KNL nodes) and GPU implementations.

1. Introduction

A major aim of experimental high-energy physics (HEP) is to find rare signals of new particles produced in large numbers of collisions at accelerators such as the Large Hadron Collider (LHC) and the ATLAS and CMS experiments. Improvements in classifying these collisions would improve sensitivity to discoveries that could overturn our understanding of the universe at the most fundamental level. Neural Networks have been used in high energy physics for some time. Recently attention has focused on deep learning to improve sensitivity, as well as tackle the increase in detector resolutions and data rates in HEP experiments.

Here we evaluate the use of deep Neural Network's (NN) on low-level detector data directly for physics analysis without reconstruction of physics objects like jets; without tuning of analysis variables; and using data from the entire calorimeter together with other sub-detectors. We use modern deep learning methods for performance; relate performance to that of physics variables; and ensure results are robust to factors such as pileup and alternative physics models.

The second goal of this work is to ensure these NN architectures run efficiently with popular deep-learning software frameworks on the Cori supercomputer at the National Energy Research Scientific Computing Center (NERSC). This is a major computing resource that is predominately composed of Intel Knights Landing (KNL) XeonPhi CPUs. We seek to improve single node



performance on these CPUs, providing timings, optimisations and recipes, as well as distributed training up to the full scale of the machine (10k KNL nodes).

We provide background, in Section 2.1, on the physics use-case studied and benchmark selections; in Section 2.2, the CNN architecture we employ; and, in Section 2.3 the computing resources at NERSC. We then provide physics performance results in Section 3.1 together with interpreting how the learned network compares to the physics variables used in the benchmark analysis (in Section 3.2). Then in Section 3.3 we show timing results for implementations in various frameworks and on GPU and CPU resources at NERSC, before concluding in Section 4.

2. Background

2.1. Physics and datasets

For this study, we take as a use-case a particular analysis searching for new massive supersymmetric (‘RPV-Susy’) particles in multi-jet final states at the LHC [1]. We focus on ‘gluino-cascade decays’ with gluino and neutralino masses of 1400 GeV and 850 GeV by default (and varied for performance studies).

We use the Pythia event generator [2] interfaced to the Delphes fast detector simulation [3] and using the default Delphes ATLAS detector configuration (‘card’). We generate events for two classes: the RPV-Susy signal and the most prevalent background (‘QCD’). In order to have sufficient quantities of background events at high transverse momentum (p_T) that are harder to separate, the QCD sample is generated in ranges of p_T and then weighted according to the amount expected (cross-section) in both training loss and the evaluation.

Before training our network we apply some of the physics selections used in [1] to filter images to those more challenging to discriminate, resulting in a training sample of around 400k events. We compare the performance of our deep network to our own implementation of the full selections in [1] as a baseline benchmark. To produce the jet variables we use the same jet algorithm as in the physics analysis with parameters given in Appendix A and applied using ‘FastJet’ [4] within Delphes. The preselection and benchmark selections are also summarized in Appendix A. We have verified that our samples and baseline selections give performance comparable to that obtained in [1] providing a meaningful benchmark even though those selections were not tuned for exactly these datasets. We also compare our analysis to shallow classifiers based on these physics variables. As inputs to those classifiers, we use the full four-momentum of five jets with highest transverse momentum as well the jet variables used in the physics analysis selections: namely the sum of jet mass, number of jets, and $\Delta\eta$ between the leading 2 jets.

Data from the surface of the cylindrical detector is represented as a 2D image with coordinates corresponding to azimuthal angle ϕ and pseudorapidity η . For the pixel intensity in this image, we use either the overall energy deposited in the combined calorimeter or split the energy deposited in the electromagnetic and hadronic calorimeters, and add the number of tracks formed from the inner detector in that region to provide a three channel image. This is similar to the approach of [5] [6] except that we use large images covering the entire detector, and use these directly for classifying entire events rather than individual objects. We choose to bin the energy (and number of tracks) into uniform 64x64 bins, which correspond to the approximately 0.1x0.1 ($\eta \times \phi$) resolution of the ATLAS hadronic calorimeter, but also use larger 224x224 images for large-scale studies that could capture the granularity of the ATLAS electromagnetic calorimeter and future detectors.

2.2. Neural Net Architecture

We perform binary classification on the resulting images using a Convolutional Neural Net (CNN) [7] comprised of 3 convolution+pooling units (or 4 for larger images) with rectified linear unit (ReLU) activation functions [8]. These layers output into two fully connected layers which project into a two-dimensional vector on which a softmax function is applied to determine the

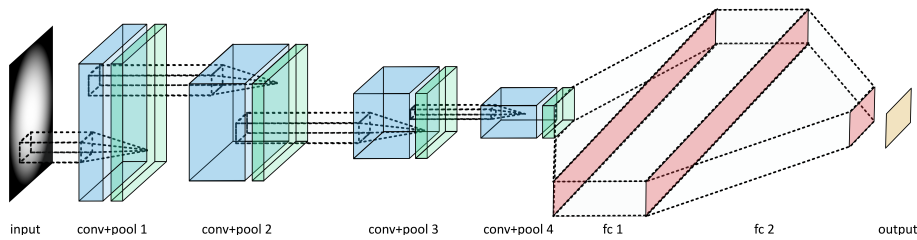


Figure 1: CNN architecture

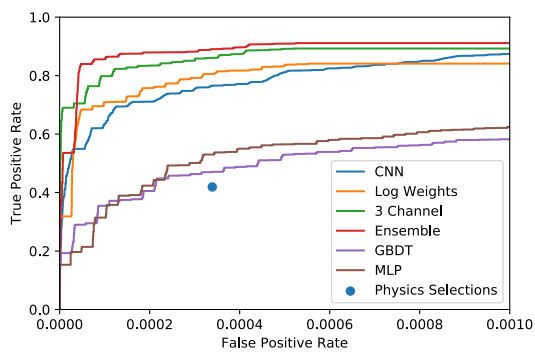


Figure 2: ROC curve comparison of different CNN implementations with physics selections and shallow classifiers

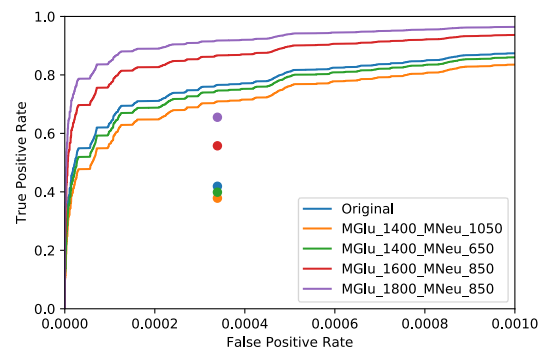


Figure 3: ROC curves for application of CNN to different signal mass samples

signal and background class probabilities. We use softmax with cross-entropy as the loss function and the ADAM optimizer [9]. The network employed is shown in figure 1.

2.3. Computational Resources

The CPU timing and scaling results reported here used the NERSC Cori system: a Cray XC40 with 2388 Intel (E5-2698 v3) Haswell and 9688 Intel XeonPhi 7250 Knight's Landing (KNL) compute nodes. Each KNL processor includes 68 1.4GHz cores with 4 HyperThreads and peak performance of 6 TeraFLOP/s (single-precision). Each Haswell node has dual 16-core 2.3 GHz processors with peak performance of 2.4 TeraFLOP/s. We compare performance to a standalone machine containing Titan X (Pascal) GPUs each capable of 10.2 TeraFLOP/s peak.

3. Results

3.1. CNN Performance

In order to obtain good significance to new physics it is necessary to combine high signal efficiency (True Positive Rate (TPR)) with very high background rejection (low False Positive Rate (FPR)). In figure 2 we show the ROC curve of TPR vs FPR for the CNN architecture described in 2.2 with a single channel for calorimeter energy (labeled 'CNN'). We compare to the physics selections described in 2.1. We achieve increased signal efficiency at same background rejection without using physics-based jet variables. We also compare the AMS (approximate median significance) [10] which gives a single measure accounting for initial pre-selection efficiency and relative signal/background cross-sections, and achieve a 1.8x improvement in AMS at the same background rejection as our implementation of the selections described in 2.1 and Appendix A.

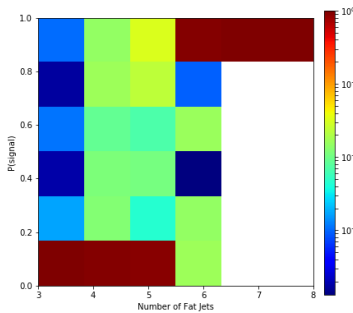


Figure 4: Comparison of CNN output with Number of Fat Jets

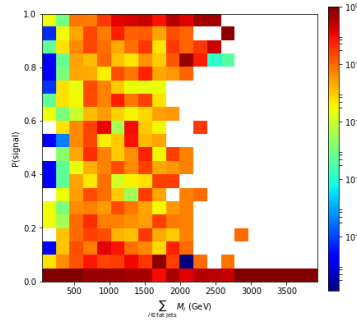


Figure 5: Comparison of CNN output with sum of jet mass

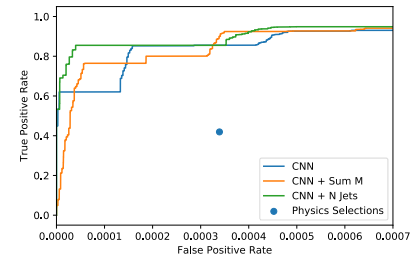


Figure 6: Change in performance on combining CNN output with jet variables

Comparison to shallow classifiers: We also compare to two shallow machine-learning approaches using high-level physics variables as inputs as detailed in Section 2.1 implemented with the scikit-learn package [11]. The first is a gradient boosted decision tree (GBDT) classifier with 100 estimators and a max depth of 3. The second is a single-hidden-layer fully-connected NN (MLP) with 100 hidden units. As shown in figure 2, these models outperform the traditional selections but under-perform relative to the CNN.

Further improving performance: In figure 2 we show further improvements to performance from modifications to our network or training. Firstly we alter the cross-section based weights applied to the training loss (described in section 2.2). The background cross-section weights are in some cases are seven orders of magnitude larger than that of the signal sample. This range of values causes some instability for training in some of the deep learning frameworks we use, so we also train the network with applying the natural logarithm of these normalized positive weights. This slightly improves the performance in the false-positive region of interest, though it performs worse than the simple cross-section for higher false-positive rates.

Next, we explore using three input channels in our CNN: namely to separate the energy deposited in the electromagnetic and hadronic calorimeters, and to add the number of tracks reconstructed in the same η/ϕ bin. Figure 2, shows the further performance improvements.

Finally a slight further improvement can be seen in figure 2 from creating an ensemble of the networks trained with original weights and the log weights (with 3 channels). This is done by taking the mean of the two network outputs as a prediction.

Robustness to different signal samples and pileup: The model described in preceding sections was trained on a specific cascade decay with gluino mass (M_{Glu}) of 1400 GeV, and neutralino mass (M_{Neu}) of 850 GeV. However, in figure 3 we show that applying this network to other signal samples with different input masses, without retraining, still offers good performance that exceeds that of the benchmark selections for each sample.

The above studies used simulated samples without a contribution from additional interactions per beam crossing (pile-up). We also generate samples using the default Delphes ATLAS detector configuration ‘card’ with added pile-up ($\mu = 20$). The physics selections benchmark has a higher false-positive rate for this sample (0.001) with a TPR of 0.6 at that point. The CNN still performs well with a TPR of around 0.9 at that FPR.

3.2. Relation between CNN and physics variables

In figures 4 and 5 we compare the predicted class probability of the CNN with two of the jet variables used in the physics analysis and a clear correlation can be seen. We also train a one-layer network combining the CNN output with each of these physics variables in turn.

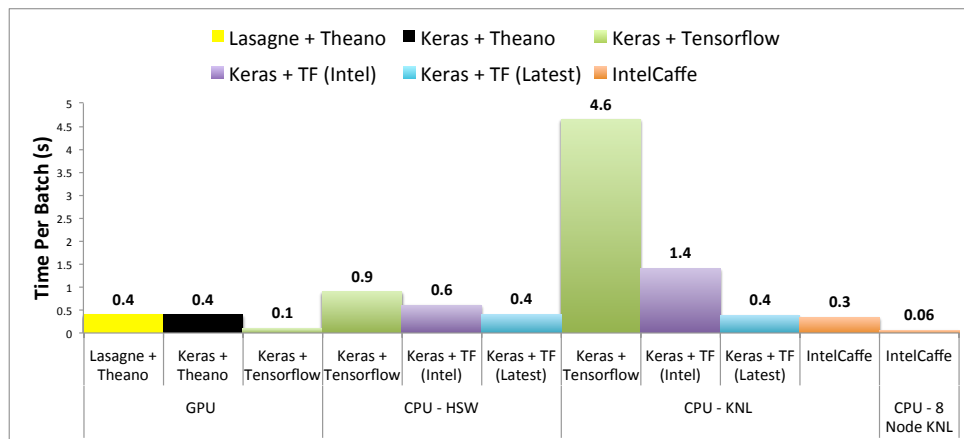


Figure 7: Comparison of time per batch (batch size=512) for different deep-learning frameworks on GPU, Haswell (HSW) and Knights Landing (KNL) CPU

The performance is shown in figure 6 and there is little or no further improvement within the statistics of the sample used, suggesting that the CNN is capturing much of the discriminating power of these jet variables, despite not being explicitly trained on them.

3.3. Training Time Performance on CPU and GPU architectures

In figure 7 we compare the training time for implementations of our network in different popular deep-learning frameworks and on different architectures. We report the time per batch of 512 examples excluding I/O. Not all implementations have been optimized so the aim is not to provide detailed comparisons of frameworks but rather to drive improved CPU performance. It can clearly be seen that the default Tensorflow (v1.2) [12] implementation performed very poorly on the KNL architecture. Several optimizations were recently introduced by Intel for Tensorflow on CPU [13] building on the Intel Math Kernel Library (MKL) [14], such as multi-threaded convolutional layers, vectorizing over channels or filters and cache blocking. These offer substantial improvements as shown in the times labeled ‘TF (Intel)’ in figure 7. Further optimizations have been carried out in a collaboration between NERSC and Intel using the architecture and data from this paper, as well as other use-cases, resulting in the ‘TF (Latest)’ measurement in figure 7. These include MKL implementation of element-wise operations to avoid MKL to Eigen conversions. This version is currently available for use on NERSC systems and will be included into future releases of Tensorflow. Code, datasets and recipes are provided in Appendix B.

Performance improvements for CPU were also made in IntelCaffe [15]. Figure 7 shows this implementation also performs well. In addition a multi-node version of IntelCaffe using the Intel ML SL library [16] is available at NERSC. Figure 7 shows that using 8 nodes (with a constant batch size per node) one can train this network approximately 6x faster than a single node. With other collaborators, we have scaled this Caffe implementation on the same problem, but with large 224x224 images, to 9600 KNL nodes on Cori. That study is reported in detail in [17].

4. Conclusions

We have implemented deep convolutional networks on large whole-detector images directly for physics analysis. We find this offers improved sensitivity than physics-variable based selections, and shallow classifiers using those variables, without the need for jet reconstruction. Further improvements come, in particular, from adding multiple sub-detectors as channels. This network is robust to pileup and can be applied to other masses of the signal sample without retraining.

Furthermore, we have used this implementation and data to benchmark and improve popular deep learning libraries on CPU architectures including XeonPhi/KNL on the NERSC Cori supercomputer as well as to demonstrate distributed training up to 9600 KNL nodes.

References

- [1] ATLAS Collaboration 2016 Search for massive supersymmetric particles in multi-jet final states produced in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC Tech. Rep. ATLAS-CONF-2016-057 CERN Geneva URL <http://cds.cern.ch/record/2206149>
- [2] Sjöstrand T, Mrenna S and Skands P 2008 *Computer Physics Communications* **178** 852 – 867
- [3] de Favereau J, Delaere C *et al.* 2014 *JHEP* **2014** 57
- [4] Cacciari M, Salam G P and Soyez G 2012 *The European Physical Journal C* **72** 1896
- [5] de Oliveira L, Kagan M, Mackey L, Nachman B and Schwartzman A 2016 *JHEP* **07** 069
- [6] Komiske P T, Metodiev E M and Schwartz M D 2017 *JHEP* **01** 110
- [7] LeCun Y, Bottou L, Bengio Y and Haffner P 1998 *Proceedings of the IEEE* **86** 2278–2324
- [8] Nair V and Hinton G E 2010 *Proceedings of ICML-10* pp 807–814
- [9] Kingma D P and Ba J 2014 *CoRR* **abs/1412.6980** URL <http://arxiv.org/abs/1412.6980>
- [10] Adam-Bourdarios C, Cowan G, Germain C, Guyon I, Kegl B and Rousseau D 2014 URL <http://higgsml.lal.in2p3.fr/documentation>
- [11] Pedregosa F *et al.* 2011 *Journal of Machine Learning Research* **12** 2825–2830
- [12] 2017 TensorFlow version 1.2.0 <https://github.com/tensorflow/tensorflow/releases/tag/v1.2.0>
- [13] 2017 TensorFlow Optimizations on Modern Intel Architecture <https://software.intel.com/en-us/articles/tensorflow-optimizations-on-modern-intel-architecture>
- [14] 2017 Introducing DNN primitives in Intel Math Kernel Library <https://software.intel.com/en-us/articles/introducing-dnn-primitives-in-intelr-mkl>
- [15] 2017 Intel distribution of Caffe* <https://github.com/intel/caffe>
- [16] 2017 Intel Machine Learning Scaling Library for Linux OS <https://github.com/01org/MLSL>
- [17] Kurth T *et al.* 2017 *Accepted for SC17, arXiv:1708.05256* URL <http://arxiv.org/abs/1708.05256>

Acknowledgements

We thank Ben Nachman and Brian Amadio for many valuable discussions and physics input, and Mustafa Mustafa and the Intel Tensorflow team for CPU optimizations. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Appendix A. Benchmark analysis and pre-selection

The selections applied for the benchmark analysis were (following [1]):

Fat-jet object selection:

AntiKt R=1.0 trimmed ($R_{trim} = 0.2$, $P_{Tfrac} = 0.05$)

$P_T > 200$ GeV , $|\eta| < 2.0$

Preselection:

Leading Fat-jet $P_T > 440$ GeV

$N_{Fat-Jet} > 2$

Analysis Selection:

$|\Delta\eta_{12}|$ between leading two Fat-jets < 1.4

($N_{Fat-Jet} \geq 4$ AND $\Sigma M_{Fat-jet} > 800$ GeV)

OR ($N_{Fat-Jet} \geq 5$ AND $\Sigma M_{Fat-jet} > 600$ GeV)

Appendix B. Datasets, Code and Recipes

Example code for implementing the networks used in this study, together with datasets and recipes for running at NERSC are provided at <http://www.nersc.gov/users/data-analytics/data-analytics-2/deep-learning/deep-networks-for-hep/>