

Machine Learning for Jet Physics

Monday 11 December 2017 - Wednesday 13 December 2017

Lawrence Berkeley National Laboratory

Book of Abstracts

Contents

Linear jet tagging with the energy flow basis (15'+5')	1
A complete, linear basis for (machine) learning jet substructure (15'+5')	1
Machine learning in the Lund plane (15'+5')	1
Jet Response Prediction Using Jet Images (15'+5')	2
Probing heavy ion collisions using quark and gluon jet substructure with machine learning (15'+5')	2
Deep(Boosted)Jet: Boosted jet identification using particle-level convolutional neural networks (15'+5')	3
Top tagging with jet constituents and Long Short-Term Memory (LSTM) networks. (15'+5')	3
Recursive Neural Networks in quark/gluon tagging (15'+5')	4
Identifying QCD transition using convolution neural network (15'+5')	4
Learning the Physics of Jet Evolution with a Recurrent Neural Network Part I (15')	4
Deep Neural Networks for whole, multi-jet event classification and generation (15'+5')	5
Registration	5
Welcome and Logistics	5
Jets and ML in Theory	5
Jets and ML in CMS (30'+15')	5
Jets and ML in ATLAS (30'+15')	5
Discussion and Closeout	6
Gluon and ONNX	6
Introduction and Overview (15'+5')	6
"Planing" to expose what the machine is learning (15'+5')	6
Building an anti-QCD tagger (15'+5')	6
Weak Supervision in High Dimensions (15'+5')	6

Adversarial Approaches (15'+5')	7
Deep-Learned Top Taggers from Images & Lorentz Invariance (15'+5')	7
Introduction and Overview (15'+5')	7
Learning the Physics of Jet Evolution with a Recurrent Neural Network Part II (15'+5')	7
Introduction and Overview (15'+5')	7
Advanced Light Source Tour	8
Introduction and Overview (15'+5')	8
Introduction and Overview (15'+5')	8
The Latest in GANs for Jet/Calo Simulation (15'+5')	8
Machine Learning and Tracking inside Jets in ATLAS (15'+5')	8
Workshop Dinner	8
Jets as graphs: W tagging with neural message passing	8
Jets as graphs: W tagging with neural message passing (15'+5')	9
Visualization Intro	9
Gluon and ONNX	9
Meta learning	9
NIPS overview (TBC)	9

Learning more about QCD / 1**Linear jet tagging with the energy flow basis (15'+5')****Authors:** Eric Metodiev¹; Jesse Thaler¹; Patrick Komiske¹¹ MIT**Corresponding Author:** pkomiske@mit.edu

In this talk, I will demonstrate the linear power of Energy Flow Polynomials (EFPs) by applying linear classification methods to quark/gluon discrimination, boosted W tagging, and boosted top tagging, achieving performance that compares favorably to other jet representations and modern machine learning approaches. I will briefly describe novel algorithms that make use of the graph-theoretic interpretation of EFPs to improve their computational complexity over that of an arbitrary N-particle correlator, making the computation of a large number of EFPs highly feasible. I will discuss how this linear energy flow basis provides an alternative to “black-box” machine learning techniques for fully combining the (IRC-safe) information in jet observables, replacing complex models by convex linear methods with few or no hyperparameters.

Learning more about QCD / 2**A complete, linear basis for (machine) learning jet substructure (15'+5')****Authors:** Eric Metodiev¹; Jesse Thaler¹; Patrick Komiske¹¹ MIT**Corresponding Author:** metodiev@mit.edu

In this talk, I will present Energy Flow Polynomials (EFPs), a novel class of jet substructure observables that form a discrete, linear basis of all infrared- and collinear-safe information in a jet. The EFPs are multiparticle energy correlators with a powerful graph-theoretic interpretation which encompass and generalize the analytic structures present in many existing classes of jet substructure observables. I will show that many common jet substructure observables are exact linear combinations of EFPs. Further, I will demonstrate the linear, IRC-safe spanning nature of EFPs by performing linear regression with EFPs on a collection of IRC-safe and unsafe observables in a variety of jet contexts.

Learning more about QCD / 3**Machine learning in the Lund plane (15'+5')****Author:** Frederic Dreyer¹¹ MIT**Corresponding Author:** fdreyer@mit.edu

We introduce a novel representation for emission patterns inside a jet, by declustering a Cambridge-Aachen jet and using the primary-emission Lund plane coordinates. We present several possible variations of this method, and show how it can be used to construct either an n by n pixel image or a graph, which can be used as inputs for neural networks.

Using W tagging as an example, we show how these jet representations can be used as inputs for convolutional neural networks or recurrent neural networks, performing on par or better than other state-of-the-art methods. We illustrate in particular how networks trained on Lund coordinates result in excellent discrimination at high p_t .

Experimental/Practical aspects of learning with jets / 4

Jet Response Prediction Using Jet Images (15'+5')

Author: Alexx Perloff¹

Co-authors: Andrew Whitbeck²; Jean-Roch Vilmant³; Pierini Maurizio⁴; Raghav Kunnawalkam Elayavalli⁵

¹ TAMU

² FNAL

³ Caltech

⁴ CERN

⁵ Rutgers University

Corresponding Author: kunnawalkamraghav@gmail.com

Understanding and appropriately correcting for the detector response on any observable of interest is an important chore for experimentalists. Such a procedure is ultimately necessary to remove the impact of the finite detector and to facilitate direct comparisons with theoretical predictions. All current experiments take on this major task by generating Monte Carlo samples and running them through a detector simulation in GEANT. In the case of reconstructed jets, one often ends up with a parametrized extraction of the jet energy scale and resolution as function of the jet's transverse momenta and rapidity. With the recent push towards new jet observables that involve the jet structure and fragmentation, a better representation of detector driven correction is paramount to better parameterize the energy resolution and hence the inherent smearing. Since the jet image contains the full jet fragmentation and energy distribution on the incident detector, we train a deep convolutional neural network to extract the jet energy response from a given jet image. This method is shown to effectively reproduce the parametrized input and as an additional feature, capture the dependence on the energy scale on the jet's internal structure observables. We show comparisons of our model with standard multi-variable machine learning techniques and highlight the importance of such an unbiased extraction on jets in data, with the near future goal of reduced jet energy resolution uncertainties on a jet-by-jet basis.

Heavy Ions / 5

Probing heavy ion collisions using quark and gluon jet substructure with machine learning (15'+5')

Author: Raghav Kunnawalkam Elayavalli¹

Co-authors: Eric Metodiev²; Patrick Komiske²; Yang-Ting Chien²

¹ Wayne State University

² MIT

Corresponding Author: kunnawalkamraghav@gmail.com

We study the classification of quark-initiated jets and gluon-initiated jets in proton-proton and heavy ion collisions using modern machine learning techniques. We train the deep convolutional neural network on discretized jet images. The classification performance is compared with the multivariate analysis of several physically-constructed jet observables including the jet mass, the p_T^D , the multiplicity and the radial moments. We also compare with the systematic N -subjett expansion in telescoping deconstruction to exploit the information carried by the subjets. The quark and gluon jet samples generated from JEWEL are used as an example to demonstrate this general method. We find that the classification performance gradually worsens in central or high multiplicity PbPb events at 2.76 TeV in JEWEL w/recoils. The information carried by the subleading subjets can be washed out by the possible subjet thermalization or randomization due to the soft event activities. Our method provides a systematically improvable framework for analyzing and comparing all jet simulations and measurements in heavy ion collisions.

Jet tagging / 6

Deep(Boosted)Jet: Boosted jet identification using particle-level convolutional neural networks (15'+5')

Authors: Markus Stoye¹; Qu Huilin²

¹ CERN

² UCSB

Corresponding Author: markus.stoye@cern.ch

Identification of boosted top quarks from their hadronic decays can play an important role in searches for new physics at the LHC. We present DeepBoostedJet, a new approach for boosted jet identification using particle-flow jets at CMS. One dimensional convolutional neural networks are utilized to classify a jet directly from its reconstructed constituent particles. The new method shows significant improvement in performance compared to alternative multivariate methods using jet-level observables.

Jet tagging / 7

Top tagging with jet constituents and Long Short-Term Memory (LSTM) networks. (15'+5')

Authors: Wojciech Fedorko¹; Wojciech Fedorko²

¹ University of British Columbia

² UBC

Corresponding Author: wfedorko@phas.ubc.ca

Multivariate techniques based on engineered features have found wide adoption in the identification of jets resulting from hadronic top decays at the Large Hadron Collider (LHC). Recent Deep Learning developments in this area include the treatment of the calorimeter activation as an image or supplying a list of jet constituent momenta to a fully connected network. This latter approach lends itself well to the use of Recurrent Neural Networks. We study the applicability of architectures incorporating Long Short-Term Memory (LSTM) networks. We explore several network architectures, methods of ordering of jet constituents, and input pre-processing. The best performing LSTM-based network achieves a background rejection of 100 for 50% signal efficiency in the jet transverse momentum range of 600 to 2500 GeV. This represents more than a factor of two improvement over a fully connected Deep Neural Network (DNN) trained on similar types of inputs.

Jet tagging / 8**Recursive Neural Networks in quark/gluon tagging (15'+5')****Author:** Taoli Cheng¹¹ *University of Chinese Academy of Sciences***Corresponding Author:** chengtaoli.1990@gmail.com

I am writing to propose a talk based on the recent paper <https://arxiv.org/abs/1711.02633>. The main topic is exploring the performance of Recursive Neural Networks in quark/gluon tagging.

Heavy Ions / 9**Identifying QCD transition using convolution neural network (15'+5')****Author:** Long Pang^{None}

The initial state fluctuations in relativistic heavy ion collisions are converted to the final state correlations of soft particles in momentum space, through strong collective expansion of the quark gluon plasma (QGP) and the QCD transition from QGP to hadrons. The patterns (equations of state) encoded in the relativistic hydrodynamic evolution are extracted from the final particle spectra $\rho(p_T, \phi)$ using supervised learning with a deep convolution neural network (DCNN). Comparisons with traditional machine learning methods (such as support vector machine, decision trees and random forests...) show that the DCNN is very good at decoding physical information from complex and dynamical evolving systems.

Learning more about QCD / 10**Learning the Physics of Jet Evolution with a Recurrent Neural Network Part I (15')****Author:** Anders ANDREASSEN¹¹ *Harvard*

Many early applications of Machine Learning in jet physics are classifiers that use Convolutional Neural Networks trained on jet images. We will present a work-in-progress custom probabilistic model, tailored to learning the physics of jet production in an unsupervised way. Our model is built on a Recurrent Neural Network suited to modeling the approximate sequential splitting of a tree, which can be explicitly defined through a clustering algorithm. The model also contains fully-connected sub-networks modeling physical quantities like the QCD splitting functions.

We train our network on Pythia jets as a proof-of-principle, but our framework importantly admits training on LHC data, including the potential to be jet-algorithm independent. Given the general structure, our model can be used as a generative model for jets, though we do not anticipate that to be its primary use. Instead, we will investigate the extraction of splitting functions in various environments and their sensitivity to global jet structure using unsupervised machine learning. Further possible physics applications will be explored.

Experimental/Practical aspects of learning with jets / 11**Deep Neural Networks for whole, multi-jet event classification and generation (15'+5')****Author:** Wahid Bhimji^{None}

Several studies have had success applying deep convolutional neural nets (CNNs) to a subset of the calorimeter for individual jet classification / tagging. We explore approaches that use the entire calorimeter, combined with track information, for directly conducting multi-jet physics analyses, without the need for any jet reconstruction. We use an existing RPV-Susy analysis as a case study and 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 discuss work in progress, and possible directions, using GraphCNNs on this data and GAN approaches for generating new events of this type.

Networks are applied on GPU and multi-node CPU architectures (including Knights Landing (KNL) Xeon Phi nodes) on the Cori supercomputer at NERSC, so we also provide time-to-solution performance of CPU (scaling to multiple KNL nodes) and GPU implementations.

12

Registration

13

Welcome and Logistics**Corresponding Authors:** bpnachman@lbl.gov, naroe@lbl.gov

14

Jets and ML in Theory**Corresponding Author:** schwartz@g.harvard.edu

15

Jets and ML in CMS (30'+15')**Corresponding Author:** markus.stoye@cern.ch**Summary:**

On overview of recently introduced methods of machine learning for jet physics at the CMS experiment is given. Both, the actual machine learning tools and their inputs are discussed. Briefly also an outlook on the next challenges in these recent developments are presented.

16

Jets and ML in ATLAS (30'+15')**Corresponding Author:** francesco.rubbo@cern.ch

17

Discussion and Closeout

18

Gluon and ONNX**Heavy Ions / 19****Introduction and Overview (15'+5')****Corresponding Author:** pmjacobs@lbl.gov**Learning from data / 20****”Planing” to expose what the machine is learning (15'+5')****Author:** Bryan Ostdiek¹¹ *University of Oregon*

Applications of machine learning tools to problems of physical interest are often criticized for producing sensitivity at the expense of transparency. In this talk, I explore a procedure for identifying combinations of variables – aided by physical intuition – that can discriminate signal from background. Weights are introduced to smooth away the features in a given variable(s). New networks are then trained on this modified data. Observed decreases in sensitivity diagnose the variable’s discriminating power. Planing also allows the investigation of the linear versus non-linear nature of the boundaries between signal and background. I will demonstrate these features in both an easy to understand toy model and an idealized LHC resonance scenario.

Learning from data / 21**Building an anti-QCD tagger (15'+5')****Corresponding Author:** jhc296@umd.edu**Learning from data / 22**

Weak Supervision in High Dimensions (15'+5')

Corresponding Author: metodiev@mit.edu

Learning from data / 23

Adversarial Approaches (15'+5')

Corresponding Author: kyle.cranmer@nyu.edu

We use an adversarial neural network to train a jet classifier that remains largely uncorrelated with the jet mass — a nuisance parameter that is highly correlated with the observed features. This adversarial training strategy balances the dual objectives of classification accuracy and decorrelation, reducing the deleterious effect of systematic uncertainties in the background modeling. The result is a robust classifier with improved discovery significance relative to existing jet classification strategies.

Jet tagging / 24

Deep-Learned Top Taggers from Images & Lorentz Invariance (15'+5')

Author: Gregor Kasieczka¹

¹ *Uni Hamburg*

Distinguishing hadronic top quark decays from light quark and gluon jets (top tagging) is an important tool for new physics searches at the LHC and allows the comparison of different machine learning approaches. We present results on using convolutional neural networks as well as recent studies employing a physics motivated network architecture based on Lorentz Invariance (and not much else) for top tagging. We also discuss further generalisations of this approach.

Learning from data / 25

Introduction and Overview (15'+5')

Corresponding Author: freytsis@uoregon.edu

Learning more about QCD / 26

Learning the Physics of Jet Evolution with a Recurrent Neural Network Part II (15'+5')

Corresponding Author: christopher.robert.frye@gmail.com

Learning more about QCD / 27

Introduction and Overview (15'+5')

Corresponding Author: ytchien@mit.edu

28

Advanced Light Source Tour

Jet tagging / 29

Introduction and Overview (15'+5')

Corresponding Author: matthew.dolan@unimelb.edu.au

Experimental/Practical aspects of learning with jets / 30

Introduction and Overview (15'+5')

Corresponding Author: jchilders@anl.gov

Experimental/Practical aspects of learning with jets / 31

The Latest in GANs for Jet/Calo Simulation (15'+5')

Corresponding Author: mpaganini@lbl.gov

Experimental/Practical aspects of learning with jets / 32

Machine Learning and Tracking inside Jets in ATLAS (15'+5')

Corresponding Author: wpmccormack@lbl.gov

33

Workshop Dinner

34

Jets as graphs: W tagging with neural message passing

Corresponding Author: henrion@nyu.edu

Jet tagging / 35

Jets as graphs: W tagging with neural message passing (15'+5')

Corresponding Author: henrion@nyu.edu

Representing jets / 36

Visualization Intro

Corresponding Author: fillmein@lbl.gov

37

Gluon and ONNX

Recent results in ML / 38

Meta learning

Corresponding Author: lukedeoliveira@lbl.gov

39

NIPS overview (TBC)