Jets and substructure in proton-proton collisions

Heavy-ion Jet Substructure Workshop, University of Bergen, 14 May 2019

Frédéric Dreyer

UNIVERSITY OF OXFORD

Physics at the high energy frontier

- LHC has been colliding protons at $13\,\mathrm{TeV}$ center-of-mass energy.
- Particle physics entering precision phase in study of EW symmetry breaking.
- Searching for new physics at the highest energy ever attained.

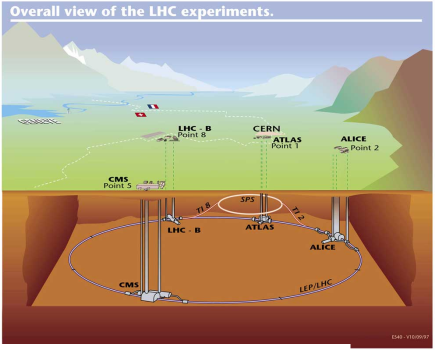

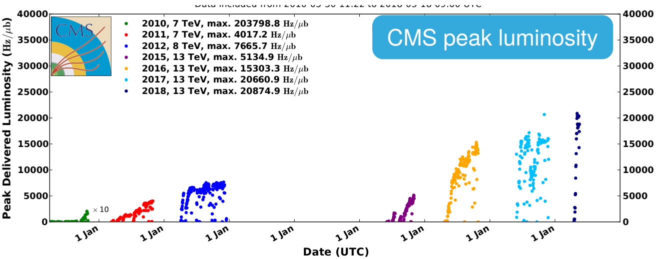

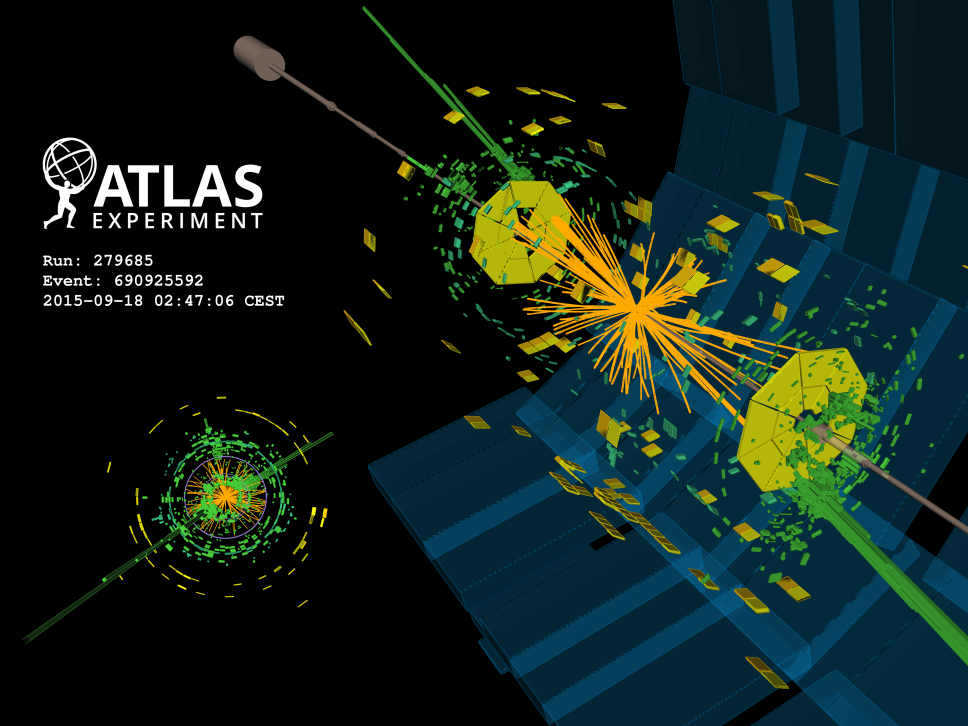
ATLAS
EXPERIMENT
Run: 279685
Event: 690925592
2015-09-18 02:47:06 CEST

JET SUBSTRUCTURE AND MACHINE LEARNING

Jets as proxies for partons

Because of color confinement, quarks and gluons shower and hadronise immediately into collimated bunches of particles.

Hadronic **jets** can emerge from a number of processes

- scattering of partons inside colliding protons,
- hadronic decay of heavy particles,
- radiative gluon emission from partons, ...

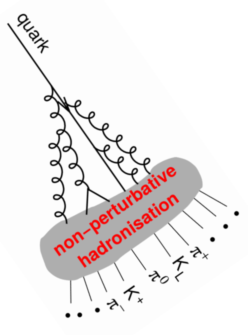

Jets are prevalent at hadron colliders

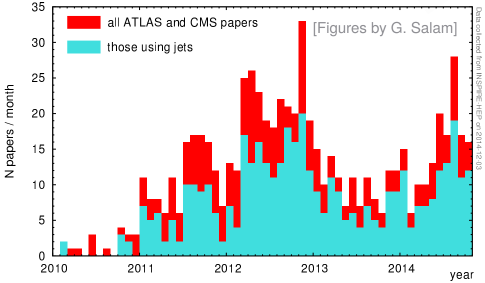

Jet algorithms

A jet algorithm maps final state particle momenta to jet momenta.

$$\underbrace{\{p_i\}}_{\text{particles}} \Longrightarrow \underbrace{\{j_k\}}_{\text{jets}}$$

This requires an external parameter, the jet radius R, specifying up to which angle separate partons are recombined into a single jet.

Basic idea of jet algorithm is to invert QCD branching process, clustering pairs which are closest in metric defined by the divergence structure of the theory.

$$d_{ij} = \min(k_{t,i}^{2p}, k_{t,j}^{2p}) \frac{\Delta_{ij}^2}{R^2}$$

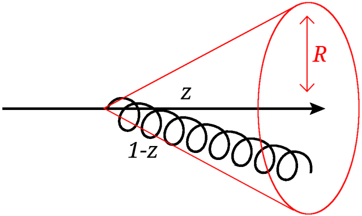

Choice of jet radius

The jet radius parameter roughly controls the size of the jet

- Standard choice for small-R jets: $R = 0.4$ (ATLAS and CMS)
  - Used for QCD jets by most experimental analyses
  - Aim is to contain most of the decay of light quarks and gluons
- Typical choice for large-R jets: $R = 0.8$ (CMS) or $R = 1$ (ATLAS)
  - Used for boosted jets by most experimental analyses
  - Aim is to contain hadronic decay of decaying particle such as W, Z, top, H, $\ldots$

Boosted objects at the LHC

- At LHC energies, EW-scale particles (W/Z/t...) are often produced with $p_t \gg m$, leading to collimated decays.
- Hadronic decay products are thus often reconstructed into single jets.

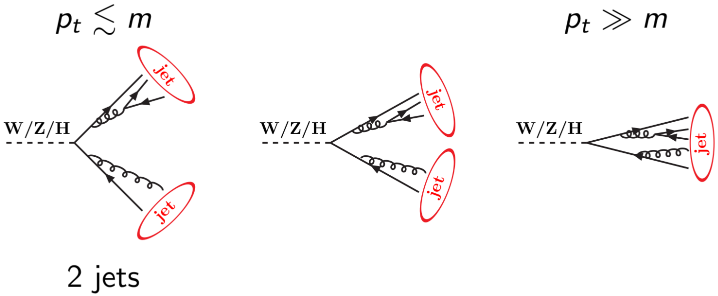

[Figure by G. Soyez]

Boosted objects at the LHC

- At LHC energies, EW-scale particles (W/Z/t. . .) are often produced with $p_t \gg m$, leading to collimated decays.
- Hadronic decay products are thus often reconstructed into single jets.

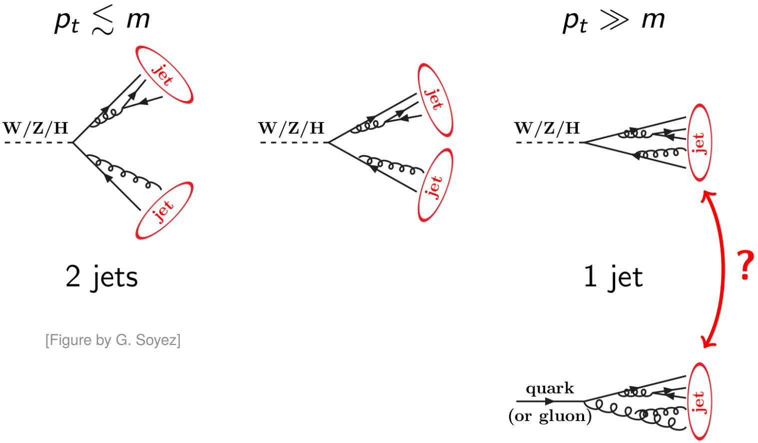

[Figure by G. Soyez]

Boosted objects at the LHC

- Many techniques developed to identify hard structure of a jet based on radiation patterns.
- In principle, simplest way to identify these boosted objects is by looking at the mass of the jet.

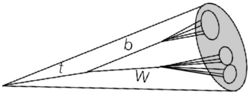

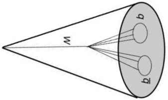

Boosted objects at the LHC

- Many techniques developed to identify hard structure of a jet based on radiation patterns.
- In principle, simplest way to identify these boosted objects is by looking at the mass of the jet.
- But jet mass distribution is highly distorted by QCD radiation and pileup.

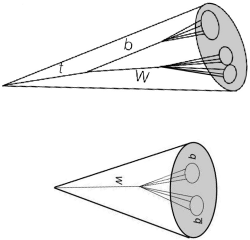

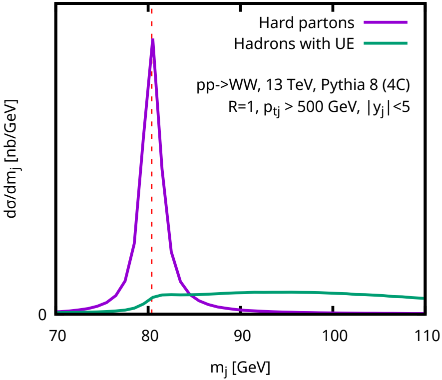

Identifying boosted objects

Two main approaches to identify boosted decays:

1. Manually constructing substructure observables that help distinguish between different origins of jets.
2. Apply machine learning models trained on large input images or observable basis.

Later in this talk: new approaches bridging some of the gap between these two techniques.

Jet grooming: (Recursive) Soft Drop / mMDT

- Mass peak can be partly reconstructed by removing unassociated soft wide-angle radiation (grooming).
- Recurse through clustering tree and remove soft branch if

$$\frac{\min(p_{t,1}, p_{t,2})}{p_{t,1} + p_{t,2}} > z_{\text{cut}} \left(\frac{\Delta R_{12}}{R_0}\right)^{\beta}$$

W jet

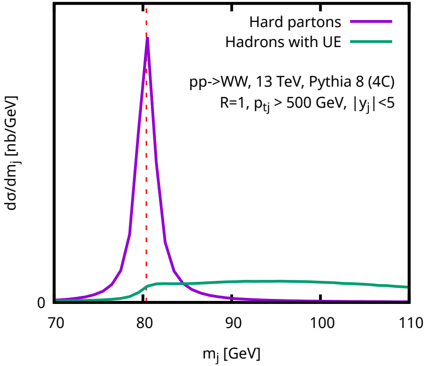

[Dasgupta, Fregoso, Marzani, Salam JHEP 1309 (2013) 029]
[Larkoski, Marzani, Soyez, Thaler JHEP 1405 (2014) 146]
[FD, Necib, Soyez, Thaler JHEP 1806 (2018) 093]

Jet grooming: (Recursive) Soft Drop / mMDT

- Mass peak can be partly reconstructed by removing unassociated soft wide-angle radiation (grooming).
- Recurse through clustering tree and remove soft branch if

$$\frac{\min(p_{t,1}, p_{t,2})}{p_{t,1} + p_{t,2}} > z_{\text{cut}} \left(\frac{\Delta R_{12}}{R_0}\right)^{\beta}$$

W jet

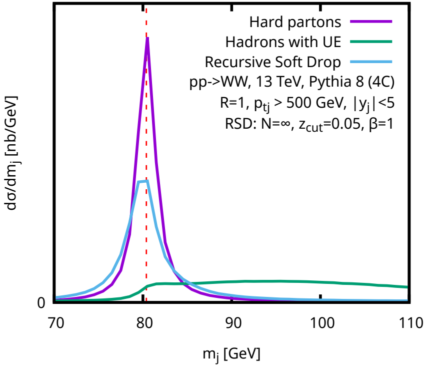

[Dasgupta, Fregoso, Marzani, Salam JHEP 1309 (2013) 029]
[Larkoski, Marzani, Soyez, Thaler JHEP 1405 (2014) 146]
[FD, Necib, Soyez, Thaler JHEP 1806 (2018) 093]

Jet grooming: common tools

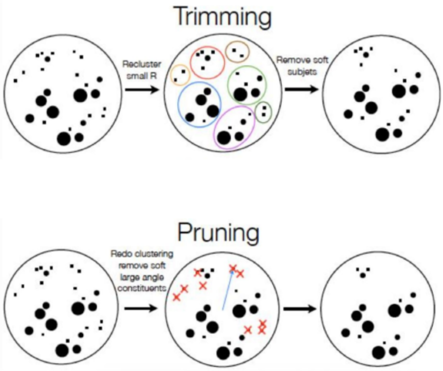

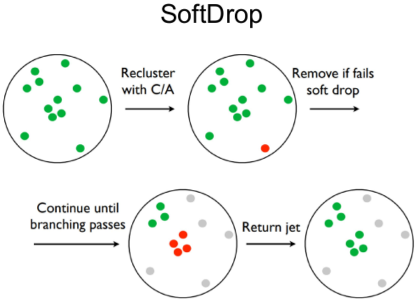

Groomed jet mass

- The jet mass is one of the simplest observables.
- Provides a unique connection between measurements and theoretical calculations.
- Grooming mitigates some non-perturbative effects such as underlying event.
- Can be calculated to high precision by resumming large logarithms and matching to fixed order

Resummation of large logs $\longleftarrow$ $\underbrace{N^kLL}_{\text{small }\rho} + \underbrace{N^mLO}_{\text{large }\rho}$ $\longrightarrow$ **Fixed order** at ~$O(\alpha_s)$

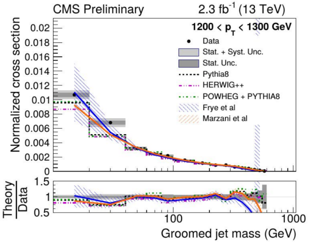

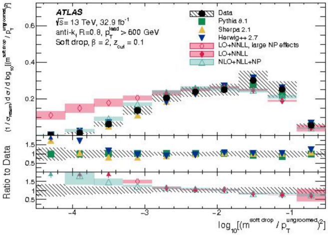

Substructure observables

- Variety of observables have been constructed to probe the hard substructure of a jet ($V/H/t$ decay lead to jets with multiple hard cores).
- Radiation patterns of colourless objects ($W/Z/H$) differs from quark or gluon jets.
- Efficient discriminators can be obtained e.g. from ratio of N-subjettiness or energy correlation functions.

[Thaler, Van Tilburg JHEP 1103 (2011) 015]
[Larkoski, Salam, Thaler JHEP 1306 (2013) 108]
[Larkoski, Moult, Neill JHEP 1412 (2014) 009]

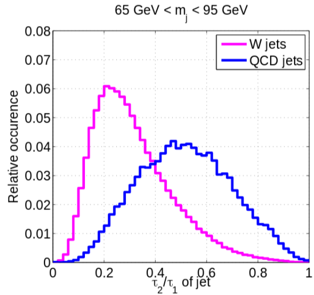

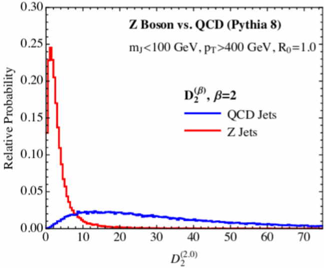

- Measures radiation around N axes that align with the dominant radiation directions

$$\tau_N^{(\beta)} = \frac{1}{p_t R^\beta} \sum_{i \in \text{jet}} p_{t,i} \min_{a_1,\dots,a_n} (\theta^\beta_{ia_1}, \dots, \theta^\beta_{ia_N})$$

- Use $\tau_{21}^{(\beta)} = \tau_2^{(\beta)}/\tau_1^{(\beta)}$ for 2-pronged jets and $\tau_{32}^{(\beta)} = \tau_3^{(\beta)}/\tau_2^{(\beta)}$ for 3-pronged jets

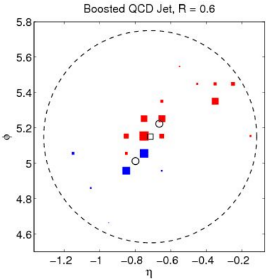

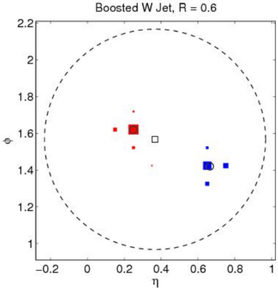

Jet shapes: Energy correlation functions

- Measures dispersion through N-point correlation functions, which are sensitive to $(N-1)$-prong substructure

$$e_2^{(\beta)} = \sum_{1\leq i<j\leq N} z_i z_j \theta_{ij}^{\beta}\,, \qquad e_3^{(\beta)} = \sum_{1\leq i<j<k\leq N} z_i z_j z_k \theta_{ij}^{\beta}\theta_{ik}^{\beta}\theta_{jk}^{\beta}$$

- Advantage: doesn't need subjet/axes finding procedure
- Efficient 2-prong discriminants can be constructed through ratio

$$D_2^{(\beta)} = \frac{e_3^{(\beta)}}{\left(e_2^{(\beta)}\right)^3}$$

- While for 3-pronged jets

$$C_3^{(\beta)} = \frac{e_4^{(\beta)} e_2^{(\beta)}}{\left(e_3^{(\beta)}\right)^2}$$

Applying Machine Learning in Jet Physics

Recent wave of results in applications of ML algorithms to jet physics.

Classification problems have been tackled through several orthogonal approaches

- Convolutional Neural Networks used on representation of jet as image
- Recurrent Neural Networks used on jet clustering tree.
- Linear combination or dense network applied to an observable basis (e.g. N-subjettiness ratios, energy flow polynomials)

Convolutational Neural Networks and Jet Images

- Project a jet onto a fixed $n \times n$ pixel image in rapidity-azimuth, where each pixel intensity corresponds to the momentum of particles in that cell.
- Can be used as input for classification methods used in computer vision, such as deep convolutional neural networks.

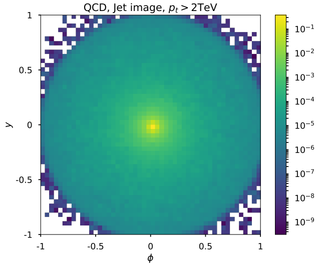

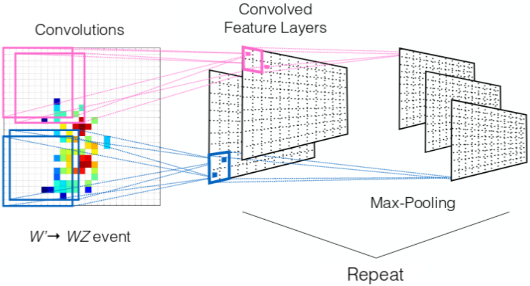

[Cogan, Kagan, Strauss, Schwartzman JHEP 1502 (2015) 118]

[de Oliveira, Kagan, Mackey, Nachman, Schwartzman JHEP 1607 (2016) 069]

Recurrent Neural Networks and clustering trees

- Train a recurrent/recursive neural network on kinematic information of successive declusterings of a jet.
- Techniques inspired from Natural Language Processing with powerful applications in handwriting and speech recognition.

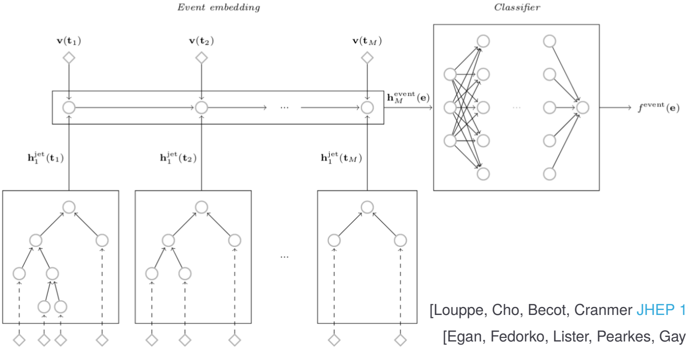

[Louppe, Cho, Becot, Cranmer JHEP 1901 (2019) 057]

[Egan, Fedorko, Lister, Pearkes, Gay 1711.09059]

Observable basis as low-dimensional representation

- Construct an observable basis that encodes the main physical properties of a jet (e.g. set of N-subjettiness ratios, energy flow polynomials, $\ldots$).
- Train a dense neural network or use linear methods to build a classifier from these inputs.

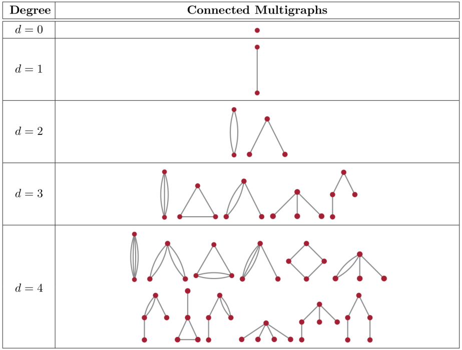

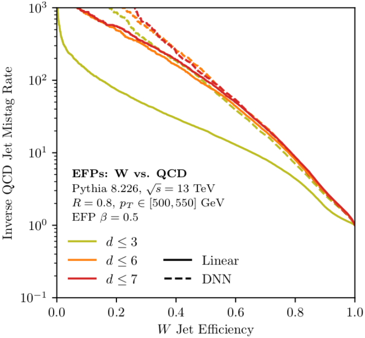

[Komiske, Metodiev, Thaler JHEP 1804 (2018) 013]

[Datta, Larkoski JHEP 1706 (2017) 073]

Beyond classification problems

- Classification problems are one of the easiest application of ML, but by far not the only one!
- Many promising applications of ML methods for:
  - fast simulations using unsupervised generative models
    [Paganini, de Oliveira, Nachman PRL 120 (2018) 042003]
  - regression tasks such as pile-up subtraction
    [Komiske, Metodiev, Nachman, Schwartz JHEP 1712 (2017) 051]
  - anomaly detection for new physics
    [Collins, Howe, Nachman PRL 121 (2018) 241803]
  - distance metric of collider events
    [Komiske, Metodiev, Thaler arXiv:1902.02346]
  - etc ...

THE LUND PLANE
(arXiv:1807.04758)

Lund diagrams

- Lund diagrams in the $(\ln z\theta, \ln\theta)$ plane are a very useful way of representing emissions.
- Different kinematic regimes are clearly separated, used to illustrate branching phase space in parton shower Monte Carlo simulations and in perturbative QCD resummations.
- Soft-collinear emissions are emitted uniformly in the Lund plane

$$dw^2 \propto \alpha_s \frac{dz}{z}\frac{d\theta}{\theta}$$

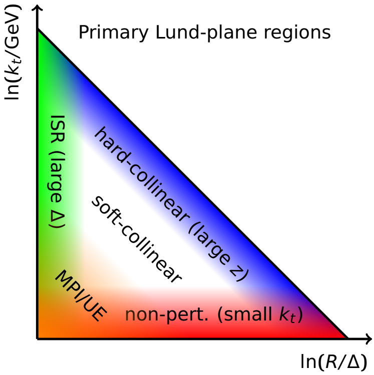

Lund diagrams

Features such as mass, angle and momentum can easily be read from a Lund diagram.

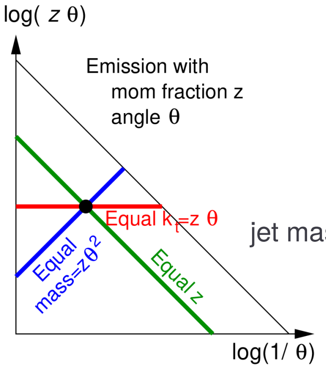

jet mass $\equiv \frac{m^2}{p_t^2 R^2} \approx z_1 \theta_1^2$

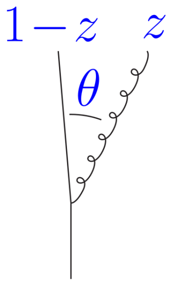

Lund diagrams for substructure

Substructure algorithms can often also be interpreted as cuts in the Lund plane.

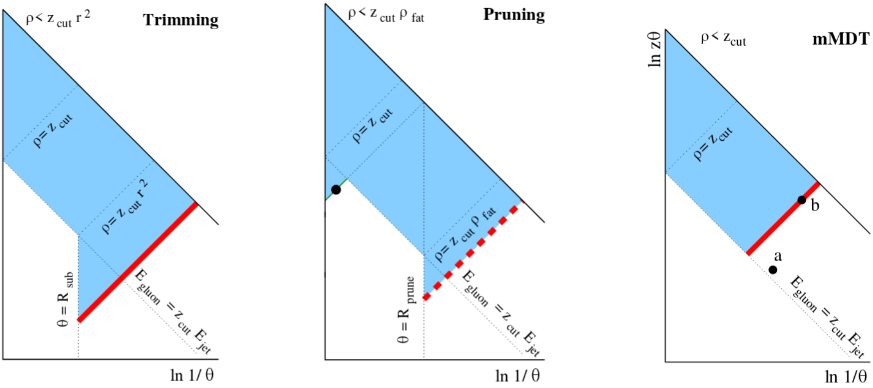

[Dasgupta, Fregoso, Marzani, Salam JHEP 1309 (2013) 029]

Studying jets in the Lund plane

Lund diagrams can provide a useful approach to study a range of jet-related questions

- First-principle calculations of Lund-plane variables.
- Constrain MC generators, in the perturbative and non-perturbative regions.
- Brings many soft-drop related observables into a single framework.
- Impact of medium interactions in heavy-ion collisions.
- Boosted object tagging using Machine Learning methods.

We will use this representation as a novel way to characterise radiation patterns in a jet, and study the application of recent ML tools to this picture.

Lund plane representation

To create a Lund plane representation of a jet, recluster a jet j with the Cambridge/Aachen algorithm then decluster the jet following the hardest branch.

1. Undo the last clustering step, defining two subjets j_1, j_2 ordered in p_t.
2. Save the kinematics of the current declustering
$$\Delta \equiv (y_1 - y_2)^2 + (\phi_1 - \phi_2)^2, \quad k_t \equiv p_{t2}\Delta,$$
$$m^2 \equiv (p_1 + p_2)^2, \quad z \equiv \frac{p_{t2}}{p_{t1}+p_{t2}}, \quad \psi \equiv \tan^{-1}\frac{y_2 - y_1}{\phi_2 - \phi_1}.$$
3. Define $j = j_1$ and iterate until j is a single particle.

Lund plane representation

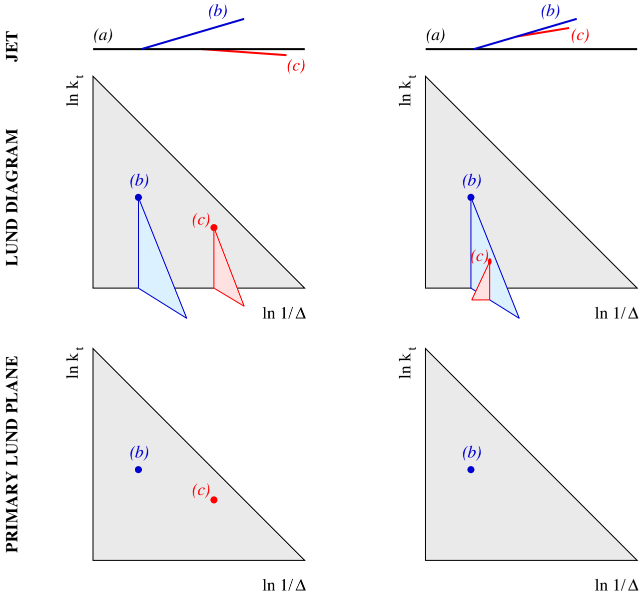

Lund representation of a jet

- Each jet has an image associated with its primary declustering.
- For a C/A jet, Lund plane is filled left to right as we progress through declusterings of hardest branch.
- Additional information such as azimuthal angle ψ can be attached to each point.

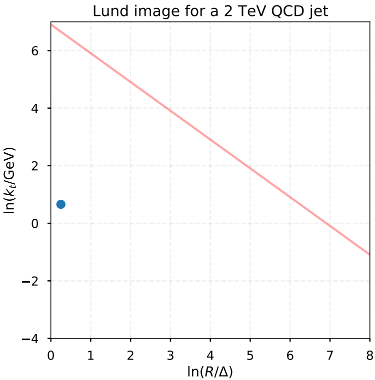

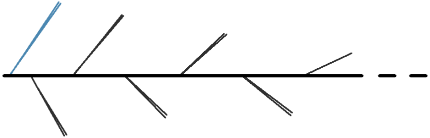

Lund representation of a jet

- Each jet has an image associated with its primary declustering.
- For a C/A jet, Lund plane is filled left to right as we progress through declusterings of hardest branch.
- Additional information such as azimuthal angle ψ can be attached to each point.

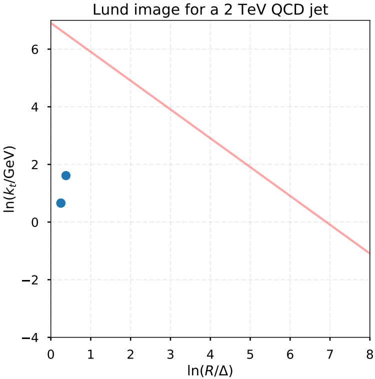

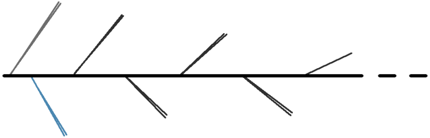

Lund representation of a jet

- Each jet has an image associated with its primary declustering.
- For a C/A jet, Lund plane is filled left to right as we progress through declusterings of hardest branch.
- Additional information such as azimuthal angle ψ can be attached to each point.

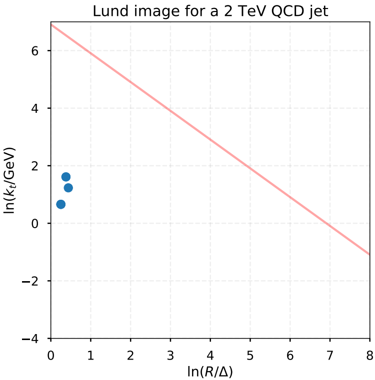

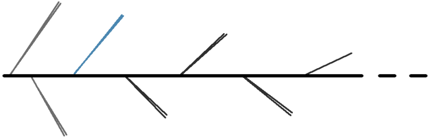

Lund representation of a jet

- Each jet has an image associated with its primary declustering.
- For a C/A jet, Lund plane is filled left to right as we progress through declusterings of hardest branch.
- Additional information such as azimuthal angle ψ can be attached to each point.

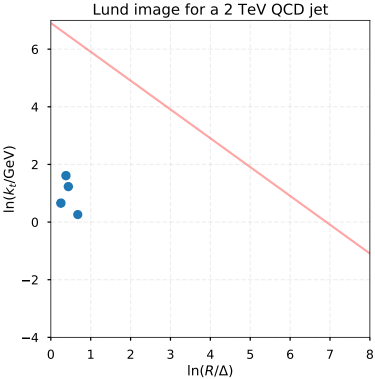

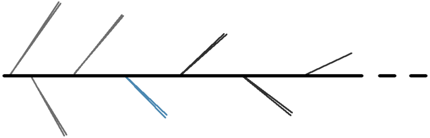

Lund representation of a jet

- Each jet has an image associated with its primary declustering.
- For a C/A jet, Lund plane is filled left to right as we progress through declusterings of hardest branch.
- Additional information such as azimuthal angle ψ can be attached to each point.

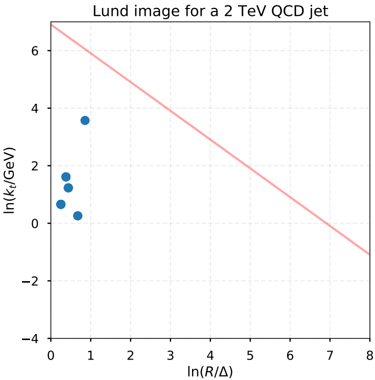

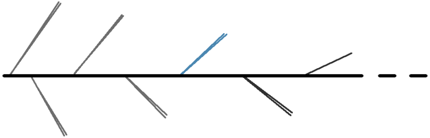

Lund representation of a jet

- Each jet has an image associated with its primary declustering.
- For a C/A jet, Lund plane is filled left to right as we progress through declusterings of hardest branch.
- Additional information such as azimuthal angle ψ can be attached to each point.

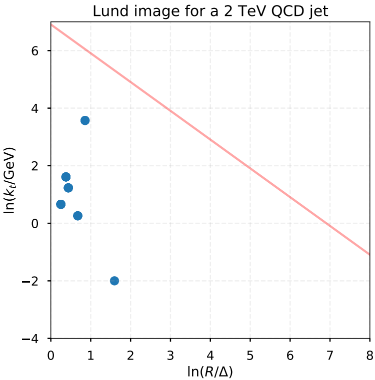

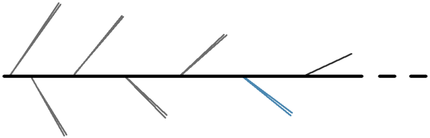

Lund representation of a jet

- Each jet has an image associated with its primary declustering.
- For a C/A jet, Lund plane is filled left to right as we progress through declusterings of hardest branch.
- Additional information such as azimuthal angle ψ can be attached to each point.

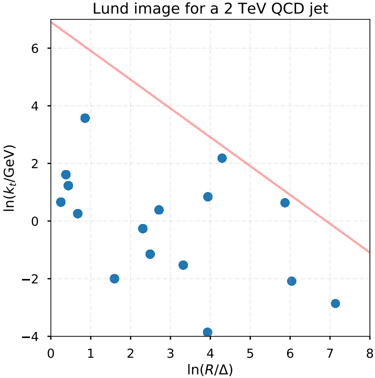

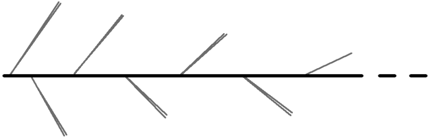

Jets as Lund images

Average over declusterings of hardest branch for 2 TeV QCD jets.

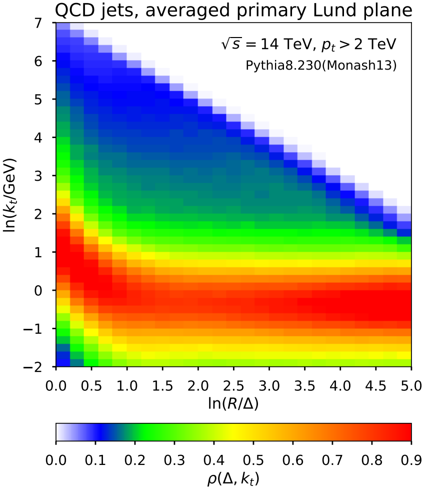

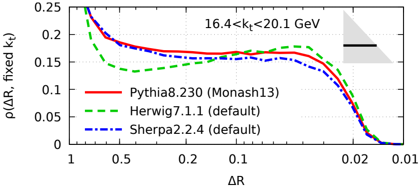

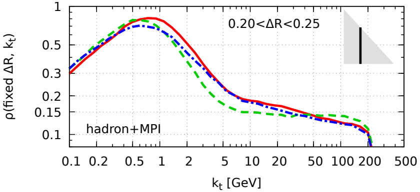

$$\rho \sim 2C\frac{\alpha_s(k_t)}{\pi}$$

Jets as Lund images

Average over declusterings of hardest branch for 2 TeV QCD jets.

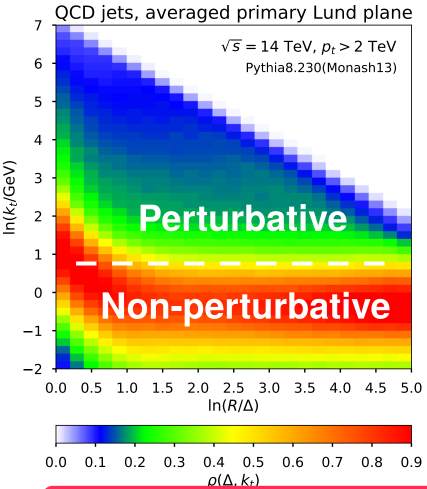

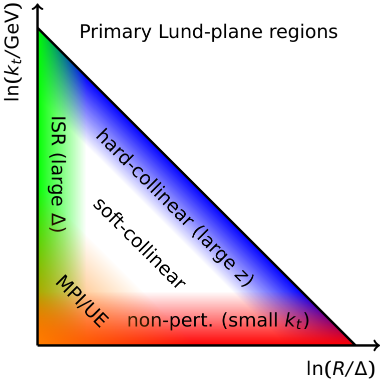

Non-perturbative region clearly separated from perturbative one.

Jets as Lund images

Average over declusterings of hardest branch for 2 TeV QCD jets.

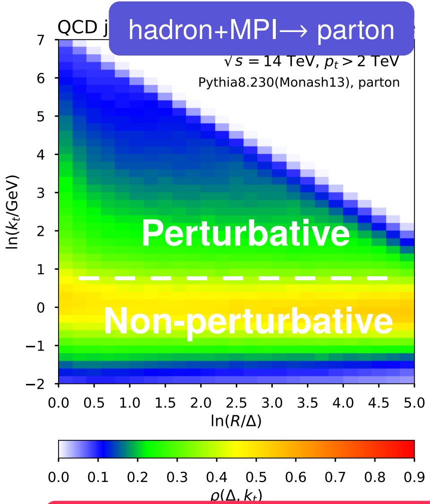

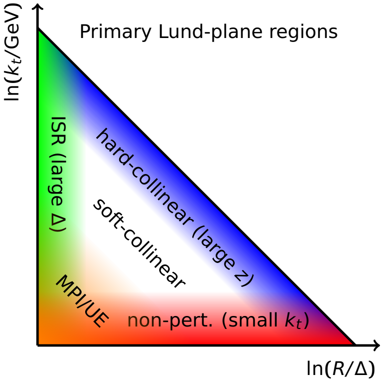

Non-perturbative region clearly separated from perturbative one.

Analytic study of the Lund plane

To leading order in perturbative QCD and for $\Delta \ll 1$, one expects for a quark initiated jet

$$\rho \simeq \frac{\alpha_s(k_t) C_F}{\pi} \bar{z} \left(p_{gq}(\bar{z}) + p_{gq}(1-\bar{z})\right), \quad \bar{z} = \frac{k_t}{p_{t,\text{jet}}\Delta}$$

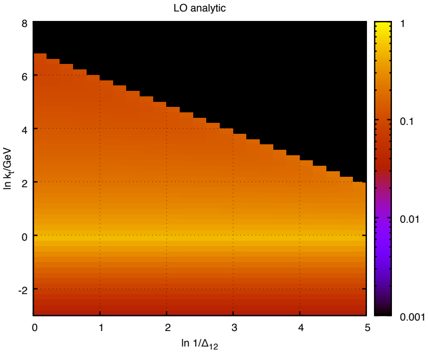

- Lund plane can be calculated analytically.
- Calculation is systematically improvable.

Analytic study of the Lund plane

To leading order in perturbative QCD and for $\Delta \ll 1$, one expects for a quark initiated jet

$$\rho \simeq \frac{\alpha_s(k_t)C_F}{\pi}\bar{z}\left(p_{gq}(\bar{z}) + p_{gq}(1-\bar{z})\right), \quad \bar{z} = \frac{k_t}{p_{t,\text{jet}}\Delta}$$

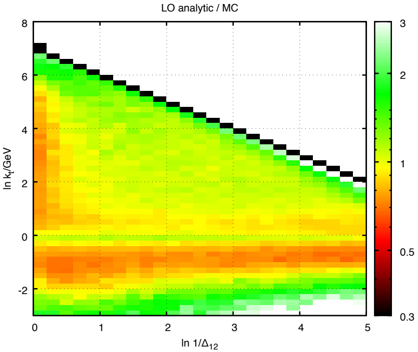

- Lund plane can be calculated analytically.
- Calculation is systematically improvable.

Declustering other jet-algorithm sequences

- Choice of C/A algorithm to create clustering sequence related to physical properties and associated to higher-order perturbative structures
- anti-k_t or k_t algorithms result in double logarithmic enhancements

$$\bar{\rho}_2^{(\text{anti-}k_t)}(\Delta,\kappa) \simeq +8 C_F\, C_A \ln^2 \frac{\Delta}{\kappa}$$

$$\bar{\rho}_2^{(k_t)}(\Delta,\kappa) \simeq -4 C_F^2 \ln^2 \frac{\Delta}{\kappa}$$

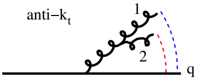

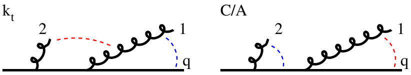

- Choice of C/A algorithm to create clustering sequence related to physical properties and associated to higher-order perturbative structures
- anti-k_t or k_t algorithms result in double logarithmic enhancements

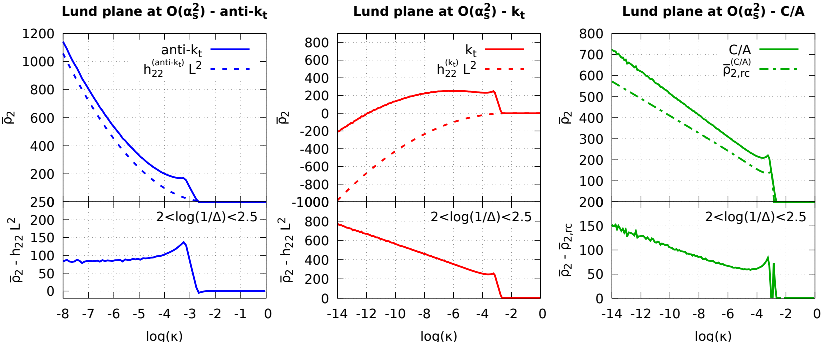

Declustering other jet-algorithm sequences

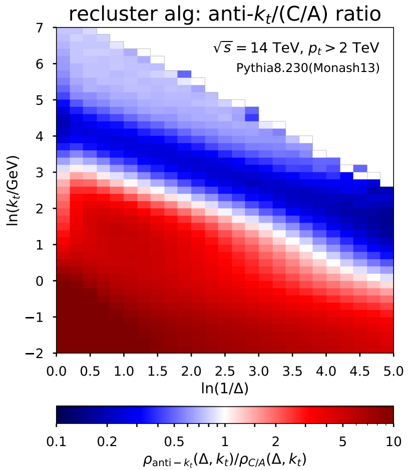

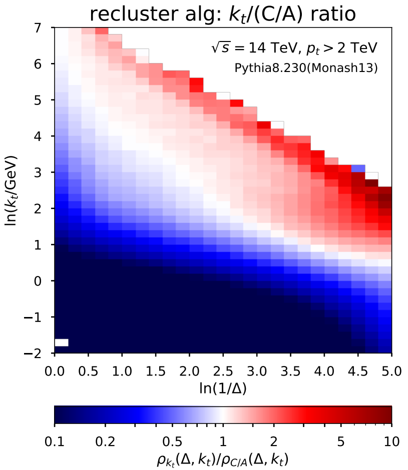

Lund images for QCD and W jets

- Hard splittings clearly visible, along the diagonal line with jet mass $m = m_W$.

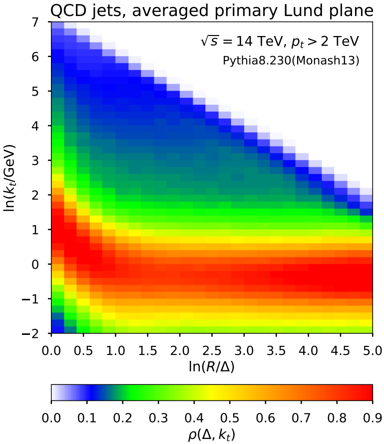

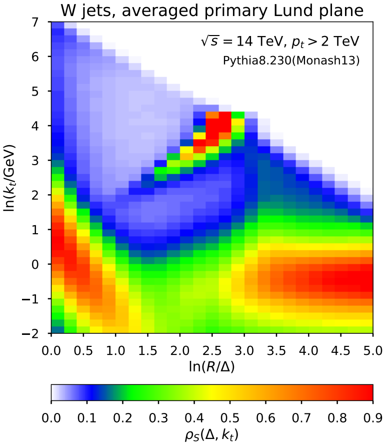

Discriminating features in the Lund plane

- Can identify discriminating features by considering log ratio of averaged images.
- W peak is clearly visible – but after cuts, depletion of emissions at relatively large angles remains distinctive signature.

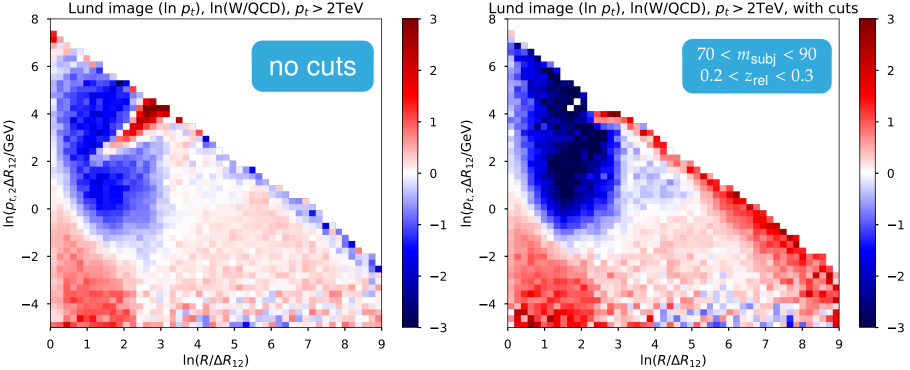

Secondary Lund plane

- Secondary Lund planes are ignored: some information is therefore lost, but still achieves good performance.
- This limitation can be overcome by extending the methods we will discuss to include secondary planes as inputs.

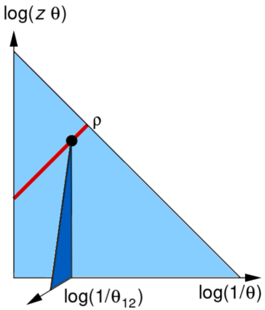

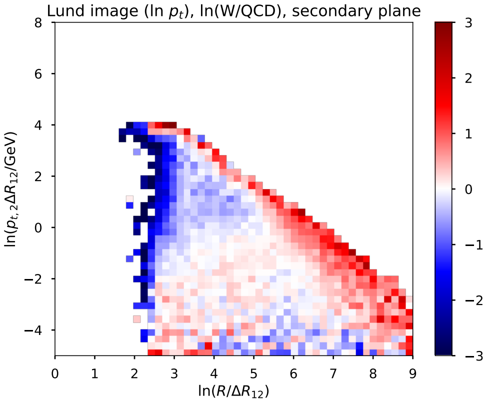

APPLICATION TO BOOSTED W TAGGING

We will now investigate the potential of the Lund plane for boosted-object identification.

Two different approaches:

- A log-likelihood function constructed from a leading emission and non-leading emissions in the primary plane.
- Use the Lund plane as input for a variety of Machine Learning methods.

As a concrete example, we will take dijet and WW events, looking at CA jets with $p_t > 2$ TeV.

Log-likelihood use of Lund Plane: leading emission

Log-likelihood approach takes two inputs:

- First one obtained from the "leading" emission.
- The second one which brings sensitivity to non-leading emissions.

Leading emission is determined to be the first emission in the Lund declustering sequence that satisfies $z > 0.025$ ($\sim$ mMDT tagger)

Define a $\mathcal{L}_\ell$ log likelihood function

$$\boxed{\mathcal{L}_\ell(m,z) = \ln\left(\frac{1}{N_S}\frac{dN_S}{dmdz}\bigg/\frac{1}{N_B}\frac{dN_B}{dmdz}\right)}$$

where the ratio of $\frac{dN_{S/B}}{dmdz}$ is the differential distribution in m and z of the leading emission for signal sample (background) with $N_S(N_B)$ jets.

Log-likelihood use of Lund Plane: non-leading emissions

Non-leading $(n\ell)$ emissions within the primary Lund plane are incorporated using a function

$$\mathcal{L}_{n\ell}(\Delta, k_t; \Delta^{(\ell)}) = \ln\left(\rho_S^{(n\ell)} / \rho_B^{(n\ell)}\right)$$

where $\rho^{(n\ell)}$ is determined just over the non-leading emissions,

$$\rho^{(n\ell)}(\Delta, k_t; \Delta^{(\ell)}) = \frac{dn^{(n\ell)}_{\text{emission}}}{d\ln k_t \, d\ln 1/\Delta \, d\Delta^{(\ell)}} \bigg/ \frac{dN_{\text{jet}}}{d\Delta^{(\ell)}}$$

as a function of the angle $\Delta^{(\ell)}$ of the leading emission.

Log-likelihood use of Lund Plane: non-leading emissions

$\mathcal{L}_{n\ell}$ log-likelihood function in a specific bin.

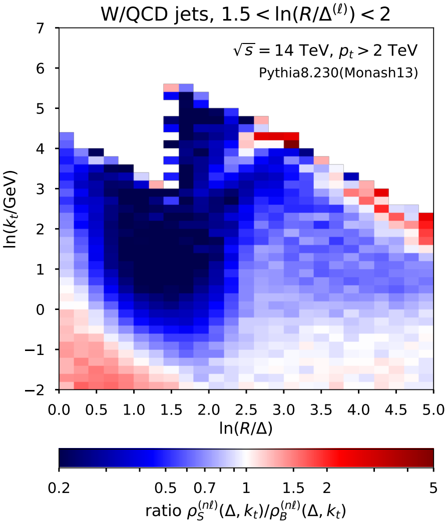

Log-likelihood use of Lund Plane: full discriminator

Overall log-likelihood signal-background discriminator for a given jet is then given by

$$\mathcal{L}_{\text{tot}} = \mathcal{L}_{\ell}(m^{(\ell)}, z^{(\ell)}) + \sum_{i \neq \ell} \mathcal{L}_{n\ell}(\Delta^{(i)}, k_t^{(i)}; \Delta^{(\ell)}) + \mathcal{N}(\Delta^{(\ell)})$$

where $\mathcal{N} = -\int d\ln\Delta\, d\ln k_t \left(\rho_S^{(\ell)} - \rho_B^{(\ell)}\right)$.

Each subjet i in the sum brings information about whether it is in a more background-like or signal-like part of the Lund plane.

Optimal discriminator if:

- Leading emission correctly associated with W's two-prong structure.
- Non-leading emissions are independent from each other.
- Emission patterns for those emissions depend only on $\Delta^{(\ell)}$.

- Compare the LL approach in specific mass-bin with equivalent results from the Les Houches 2017 report (arXiv:1803.07977).
- Substantial improvement over best-performing substructure observable.

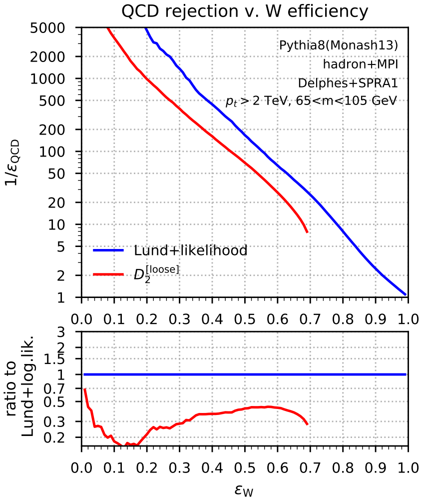

ML applied to the Lund plane

A variety of ML methods can be applied to the Lund plane in order to construct efficient taggers.

We will investigate three approaches:

- Convolutional Neural Networks (CNN) applied on 2D Lund images.
- Deep Neural Networks (DNN) applied on the sequence of declusterings.
- Long Short-Term Memory (LSTM) networks applied on the sequence of declusterings.

Recurrent networks with a Lund plane

- Jets generally associated with a clustering trees, where each node contains similar type of information.
- Particularly well-adapted for recurrent networks, which loop over inputs and use the same weights.
- LSTMs are a widely used variant designed to have memory over longer separations.
- For each declustering node, we consider the inputs

$$\left\{ \ln(R/\Delta R_{12}),\, \ln(k_t/\,\mathrm{GeV}) \right\}$$

- Inputs are IRC safe as long as there is a cutoff in transverse momentum.

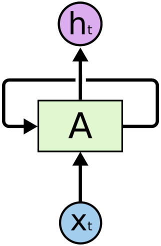

Figure from
http://colah.github.io/posts/2015-08-Understanding-LSTMs/

LSTMs for jet tagging

- LSTM network substantially improves on results obtained with other methods.
- Large gain in performance, particularly at higher efficiencies.

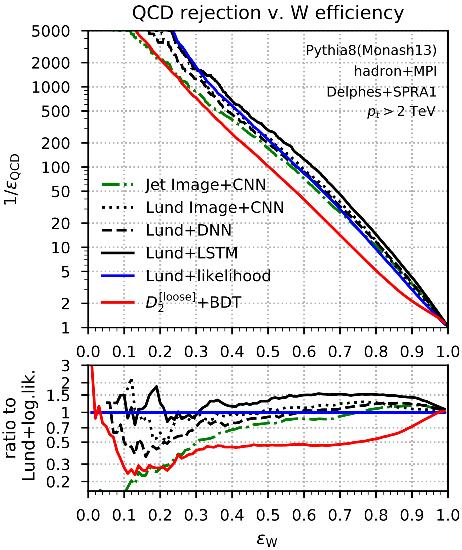

Sensitivity to non-perturbative effects

- Performance compared to resilience to MPI and hadronisation corrections.
- Vary cut on k_t, which reduces sensitivity to the non-perturbative region.

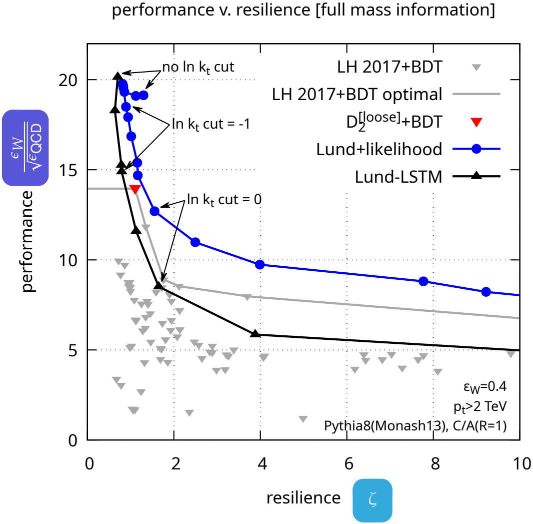

$$\Delta\epsilon = \epsilon - \epsilon'$$

$$\zeta = \left(\frac{\Delta\epsilon_S^2}{\langle\epsilon\rangle_S^2} + \frac{\Delta\epsilon_B^2}{\langle\epsilon\rangle_B^2} \right)^{-\frac{1}{2}}$$

(c.f. arXiv:1803.07977)

$$\langle\epsilon\rangle = \frac{1}{2}(\epsilon + \epsilon')$$

- Lund-likelihood performs well even at high resilience.
- ML approach reaches very good performance but is not particularly resilient to NP effects.

REINFORCED JET GROOMING
(arXiv:1903.09644)

Grooming a jet tree

- Cast jet as clustering tree where state of each node $\mathcal{T}^{(i)}$ is a tuple with kinematic information on splitting

$$s_t = \{z, \Delta_{ab}, \psi, m, k_t\}$$

- Grooming algorithm defined as a function π_g observing a state and returning an action $\{0, 1\}$ on the removal of the softer branch, e.g.

$$\pi_{\text{RSD}}(s_t) = \begin{cases} 0 & \text{if} \quad z > z_{\text{cut}} \left(\frac{\Delta_{ab}}{R_0}\right)^{\beta} \\ 1 & \text{else} \end{cases}$$

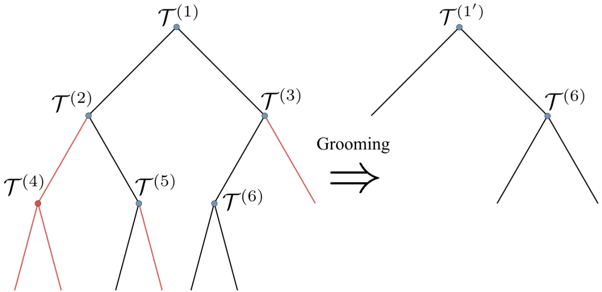

Reinforcement learning with Deep-Q-Networks

Reinforcement learning are usually built from two elements:

- an agent deciding which actions to take in order to maximize reward
- an environment, observed by the agent and affected by the action

Deep Q-Network is a RL algorithm which uses a table of Q-values $Q(s,a)$, determining the next action as the one that maximizes Q.

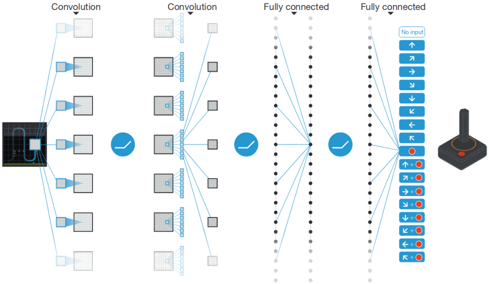

A neural network is used to approximate the optimal action-value function

$$Q^*(s,a) = \max_{\pi} \mathbb{E}[r_t + \gamma r_{t+1} + \ldots | s_t = s, a_t = a, \pi]$$

[Mnih et al, Nature 2015]

Defining a grooming environment

To find optimal grooming policy π_g, define an environment and a reward function so that problem can be solved with RL.

- Initialize list of all trees used for training.
- Each episode starts by randomly selecting a tree and adding its root to a priority queue (ordered in Δ_{ab}).
- Each step removes first node from priority queue, then takes action on removal of soft branch based on state s_t of node.
- After action, update kinematics of parent nodes, add current children to priority queue, and evaluate reward function.
- Episode terminates once priority queue is empty.

Defining the reward function

- Key ingredient for optimization of grooming policy is reward function used at each training step.
- We construct a reward with two components
  - First piece R_M evaluated on the full jet tree, comparing the jet mass to a target value.
  - Second component R_{SD} looks at kinematics of current node.
- Total reward is then given by

$$R(m, a_t, \Delta, z) = R_M(m) + \frac{1}{N_{\mathrm{SD}}} R_{\mathrm{SD}}(a_t, \Delta, z)$$

- where mass reward is defined using a Cauchy distribution

$$R_M(m) = \frac{\Gamma^2}{\pi(|m - m_{\mathrm{target}}|^2 + \Gamma^2)}$$

- To provide baseline behaviour for the groomer, we include a “Soft-Drop” reward R_{SD} evaluated on the current node
- Calculated on the current node state, gives positive reward for removal of wide-angle soft radiation and for keeping hard-collinear emissions.

$$R_{\mathrm{SD}}(a_t,\Delta,z) = a_t \min\left(1, e^{-\alpha_1 \ln(1/\Delta)+\beta_1 \ln(z_1/z)}\right)$$
$$+ (1-a_t)\max\left(0, 1-e^{-\alpha_2 \ln(1/\Delta)+\beta_2 \ln(z_2/z)}\right)$$

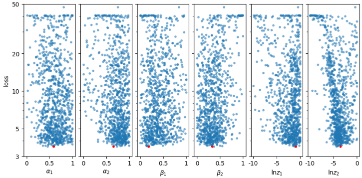

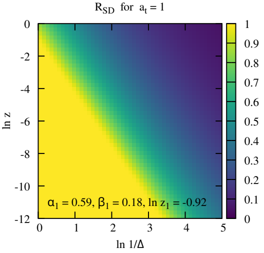

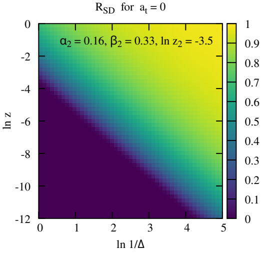

Implementation and multi-level training

- Train RL agent with multi-level approach using both signal and bkg into account. Sample consists of 500k W/QCD or Top/QCD `Pythia` 8 jets.
- At the beginning of each episode, randomly select a signal or background jet with probability $1 - p_{\text{bkg}}$.
- In the background case, mass reward function is changed to

$$R_M^{\text{bkg}}(m) = \frac{m}{\Gamma_{\text{bkg}}} \exp\left(-\frac{m}{\Gamma_{\text{bkg}}}\right).$$

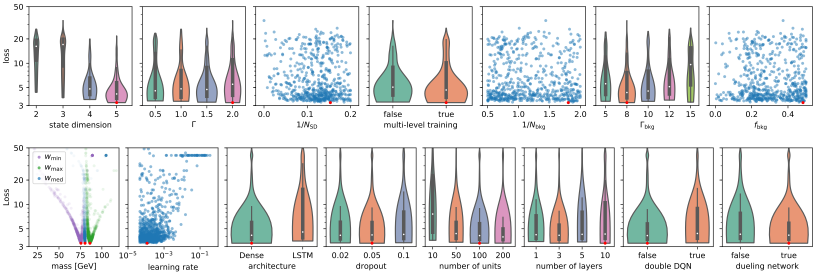

Groomed jet mass spectrum

- To test the grooming algorithm derived from the DQN agent, we apply our groomer to three test samples: QCD, W and Top jets.
- Improvement in jet mass resolution compared to RSD.
- Algorithm performs well on data beyond its training range.

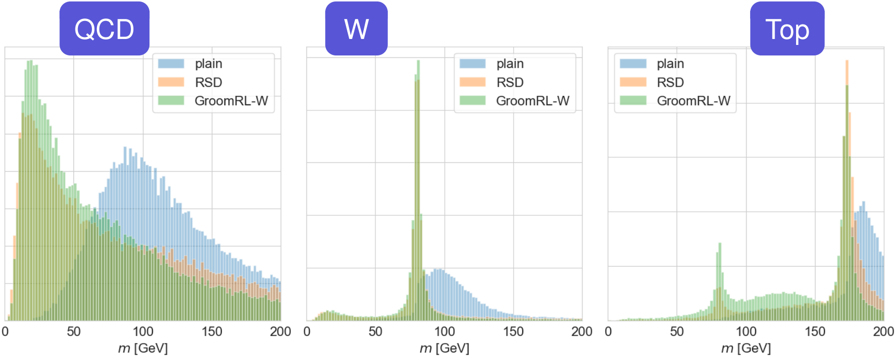

code available at **github.com/JetsGame/GroomRL**

Groomed jet mass spectrum

- To test the grooming algorithm derived from the DQN agent, we apply our groomer to three test samples: QCD, W and Top jets.
- Improvement in jet mass resolution compared to RSD.
- Algorithm performs well on data beyond its training range.

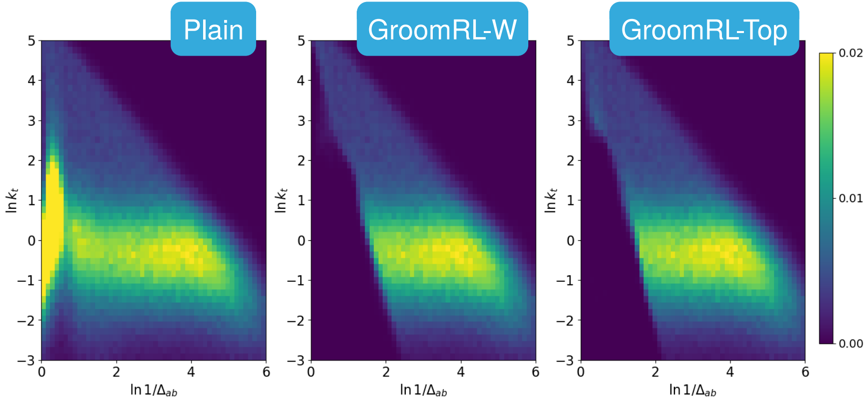

code available at **github.com/JetsGame/GroomRL**

Robustness to non-perturbative effects

- Resilience to hadronisation and underlying event corrections is a key feature of modern grooming algorithms
- Strategy derived from reinforcement learning shows similar behaviour to heuristic method
- No parton or hadron-level data was used in the training!

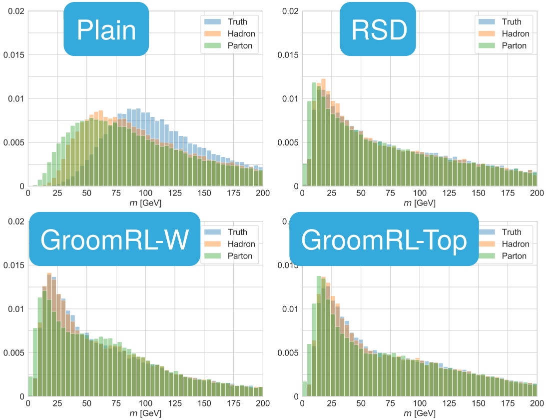

CONCLUSIONS

Conclusions: the view from pp physics

- Jet substructure is a very active subfield providing a wide range of tools that can be readily applied in heavy ion physics.
- Cross-talk with machine learning community has lead to many new advances and insights.
- Many yet to come, e.g. in tackling more complicated regression tasks or unsupervised learning approaches.
- Discussed a way to study and exploit radiation patterns in a jet using the Lund plane.
- Introduced a framework for promising application of reinforcement learning to jet grooming

  $\Longrightarrow$ easily extendable to other choices of reward function.