Jet reconstruction (experimental)
Part 2

Steven Schramm
University of Geneva

PREFIT School, DESY
March 12, 2020

Reminder from last lecture

UNIVERSITÉ DE GENÈVE

- Experimentally, jet reconstruction is mostly about dealing with the detector
  1. Devising robust input objects that exploit both trackers and calorimeters (PFlow)
  2. Mitigating pileup effects, both before (CHS) and after (JVT/fJVT) jet reconstruction
  3. Calibrating reconstructed jets back to the truth jet scale, thus reducing detector effects
  4. Comparing jets between data and MC, to correct for detector simulation imperfections
- On Tuesday, we went into these topics in detail for $R = 0.4$ anti-k_t jets
  - You then looked at pileup and detector response plots in the exercise

- Today, we will switch focus to large-R jets and identifying the origin of such jets

Large-R jets overview

UNIVERSITÉ DE GENÈVE

1. Boosted jets and jet definitions
   - Collimation, grooming, constituent modifiers, and more
2. Jet mass scale and resolution
   - Calibrating a more complex quantity than the energy
3. Hadronic decay tagging and efficiencies
   - Identifying jets containing the decay products of a parent particle

Boosted jets and jet definitions

Boosted objects and collimation

- Particles always decay back-to-back in their own reference frame
  - $p_T^A = p_T^B \approx m^X/2$
- However, particles are typically moving in the laboratory (detector) frame
  - $p_T^X \neq 0$, thus there is a "boost"
- Decay products A and B may be collimated in the detector frame
  - $\Delta R_{AB} \approx \frac{m^X}{p_T^X} \frac{1}{\sqrt{f_A(1-f_A)}}$
  - f_A = fraction of p_T^X carried by A
  - $\implies \Delta R_{AB} \gtrsim \frac{m^X}{p_T^X} \frac{1}{\sqrt{(1/2)\cdot(1/2)}}$
  - $\implies \Delta R_{AB} \gtrsim \frac{2m^X}{p_T^X}$

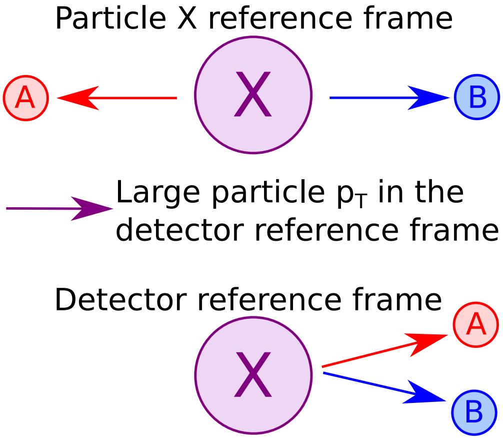

Collimation in practice

UNIVERSITÉ DE GENÈVE

- This is ΔR for $W \to qq$
  - $\Delta R_{\mathrm{qq}} \gtrsim \frac{2m^{\mathrm{W}}}{p_{\mathrm{T}}^{\mathrm{W}}}$ is pretty accurate
- Note that this is for a two-body decay
  - Complex decays like $t \to Wb \to qqb$ will not follow this as a whole
  - Individual steps ($t \to Wb$, $W \to qq$) will follow this, although the W mass complicates the process until $p_{\mathrm{T}}^{t} \gg m_{\mathrm{W}}$

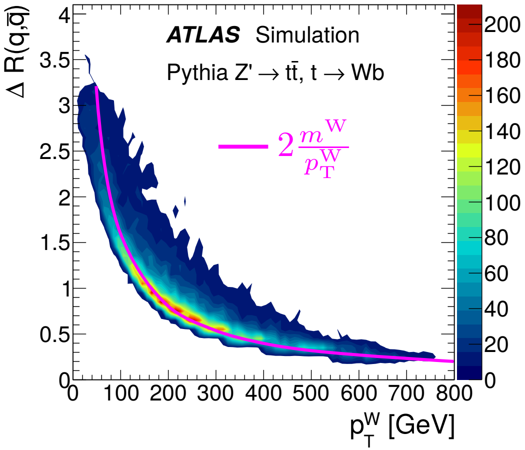

Boosted jets and collimation

UNIVERSITÉ DE GENÈVE

- What if A and B are quarks or gluons?
  - Same equations hold at parton level
  - However, the quarks/gluons hadronize
- Result: overlapping showers if $p_{\mathrm{T}}^{\mathrm{X}} \gg m_{\mathrm{X}}$
  - At some point, $R = 0.4$ jets no longer represent a single quark
- Instead, switch to a larger radius and capture all particle decay products
  - ATLAS: $R = 1.0$, CMS: $R = 0.8$
  - Then "tag" the jet to determine whether it is consistent with originating from a massive particle decay [later]

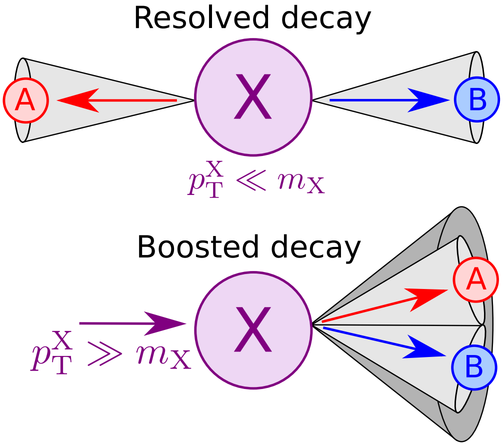

Jet mass and hadronic decays

UNIVERSITÉ DE GENÈVE

- If the entire decay is inside the jet, then we have an expectation for the jet mass
  - Mass is a boost-invariant quantity
  - Should correspond to parent particle
- If we select a sample of hadronic decays of W bosons, we should get the W mass
  - Indeed, we see the W peak!
  - Furthermore, data and MC agree
- However, large-R jets are complex
  - It took work to get here, as we will see

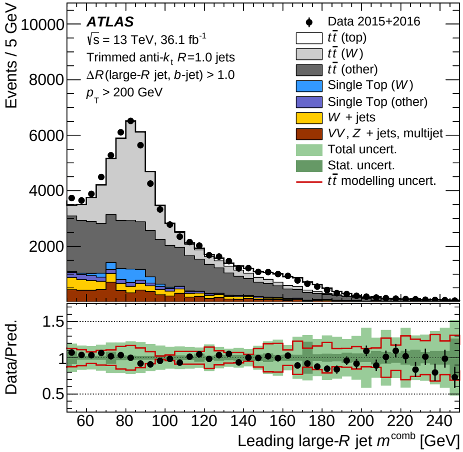

Jet mass and grooming

UNIVERSITÉ DE GENÈVE

- The jet mass is a complex quantity
  - $m^2 = \left(\sum E_i\right)^2 - \left(\sum \vec{p}_i\right)^2$
  - Involves the energy and location of every object the jet is built from
- Jet mass is very sensitive to the underlying event (UE) and pileup (N_{PV})
- Jet mass before grooming:
  - UE: m_{W} is wrong, and $m_{\mathrm{QCD}} \approx m_{\mathrm{W}}$
  - Significant pileup dependence
- Jet mass after grooming:
  - ~~UE~~: m_{W} is correct, and $m_{\mathrm{QCD}} \neq m_{\mathrm{W}}$
  - No pileup dependence

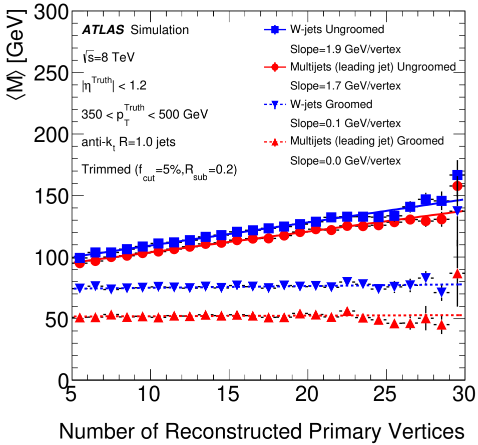

A reminder of a grooming algorithm: trimming

UNIVERSITÉ DE GENÈVE

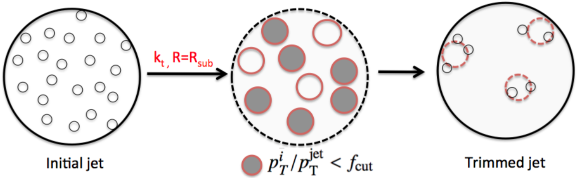

- You already discussed grooming quite a bit with Gavin, so just a brief reminder
- Trimming is currently the ATLAS default grooming algorithm
  1. Reconstruct the $R = 1.0$ anti-k_t jet as usual
  2. Build $R = 0.2$ k_t sub-jets out of the objects (the clusters) that went into a given $R = 1.0$ jet
  3. For any subjet where $p_{\mathrm{T}}^{R=0.2}/p_{\mathrm{T}}^{R=1.0} < 5\%$, remove that sub-jet and its clusters
  4. Rebuild the $R = 1.0$ jet from the clusters that survived step 3

Different grooming algorithms

UNIVERSITÉ DE GENÈVE

- There are many grooming algorithms
  - Each algorithm also has parameters
- Need to study the different options
  - This is the pileup dependence of the W mass for many grooming options
  - ATLAS default trimming is in a box
  - Trimming is very strict against pileup
- However, there are also other columns
  - "Constituent-level pileup mitigation"
  - Significantly reduces pileup dependence

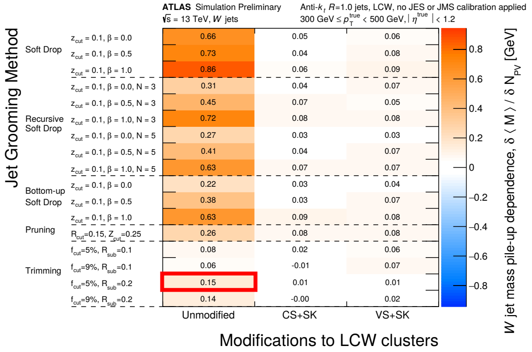

Constituent-level pileup mitigation

UNIVERSITÉ DE GENÈVE

- Recall from the first lesson that pileup has a big impact on $R = 0.4$ jets
- The impact on large-R jets is much more significant for several reasons
  - Larger jet size = larger amount of pileup in the jet, $(R = 1.0/R = 0.4)^2 = 6.25$
  - ATLAS doesn't use PFlow yet for large-R jets, so no Charged Hadron Subtraction
  - Large-R jets use variables sensitive to both energy **and location** (such as the mass)
    - Pileup is low in p_{T}, but can be at any angle, and thus can have a big impact
    - In contrast, $R = 0.4$ jets focus on p_{T} where the angle is $\sim$irrelevant
- It is thus important to have more ways to remove individual objects from large-R jets
  - This includes neutral particles, not only charged particles!
- This is the domain of constituent-level pileup mitigation

Voronoi Subtraction (VS)

UNIVERSITÉ DE GENÈVE

1. Calculate the pileup energy density, ρ
   - Same as for $R = 0.4$ jet calibration
2. Define an "area" for every object
   - Uses voronoi diagrams: split based on half-way points between objects
3. Subtract energy from each object based on its area, $p_{\mathrm{T}}^{\mathrm{corr}} = p_{\mathrm{T}}^{\mathrm{obj}} - \rho \mathcal{A}^{\mathrm{obj}}$
4. Remove resulting negative objects
5. Run jet reconstruction on what survived

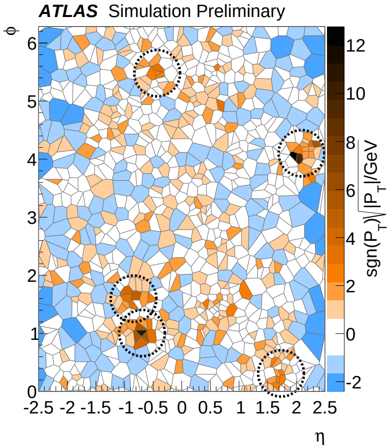

Constituent Subtraction (CS)

1. Calculate the pileup energy density, ρ
   - Same as for $R = 0.4$ jet calibration
2. Fill the detector with fake "ghosts"
   - Uniform distribution over η and ϕ
   - $p_{\mathrm{T}}^{\text{ghost}} = \rho \mathcal{A}^{\text{ghost}}$, $\mathcal{A}^{\text{ghost}}$ is the inverse of the number density of the ghosts
3. Build jets as usual, including the ghosts
4. Subtract the ghosts from nearby (ΔR) clusters/etc in each resulting jet
   - If a given cluster has a p_{T} below the nearby ghost p_{T} density, it is removed
5. Remove any remaining ghosts

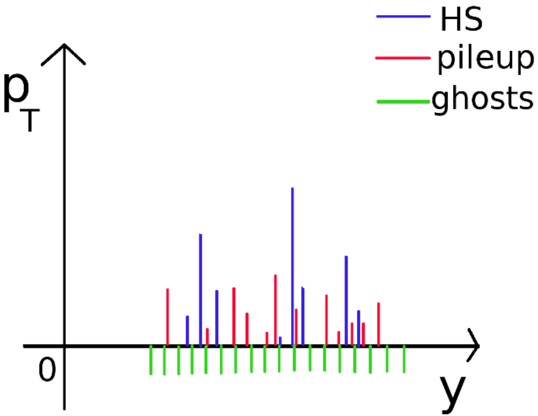

SoftKiller (SK)

UNIVERSITÉ DE GENÈVE

- Another approach is to define a dynamic p_T cut using SoftKiller
  1. Group energy deposits in the calorimeter into blocks, such as 0.6×0.6 in $\eta \times \phi$
  2. Determine the object p_T cut that will make half of the grid spaces empty
  3. Apply that cut, removing everything in any grid space which was below that threshold
- Benefits of SK alone are generally small, but pairs very well with CS and VS

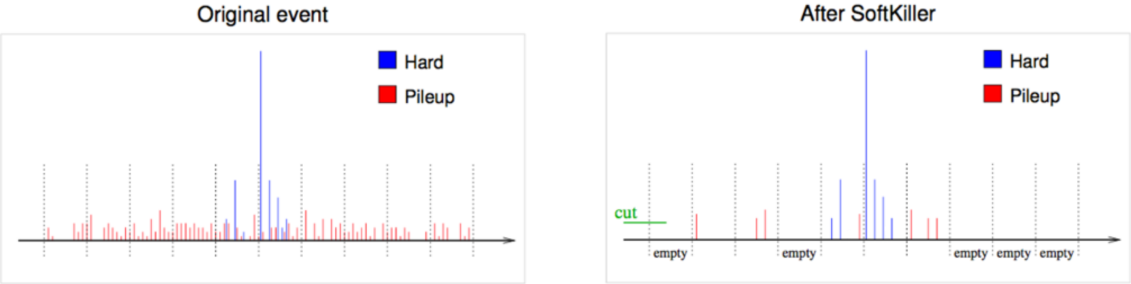

The impact of constituent-level pileup mitigation

UNIVERSITÉ DE GENÈVE

Original event

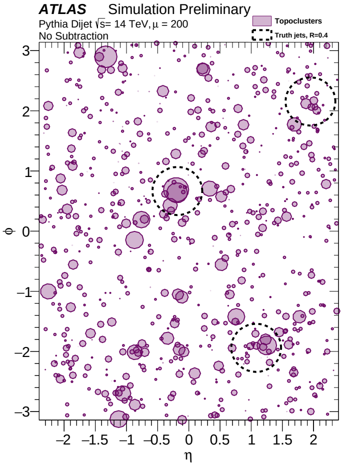

After VS

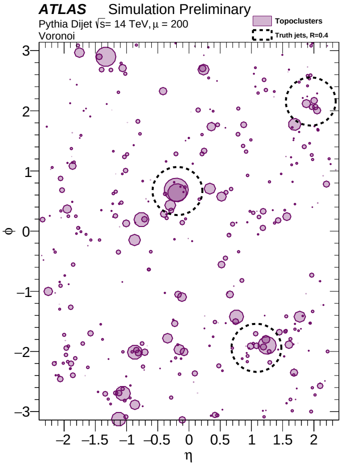

After VS+SK

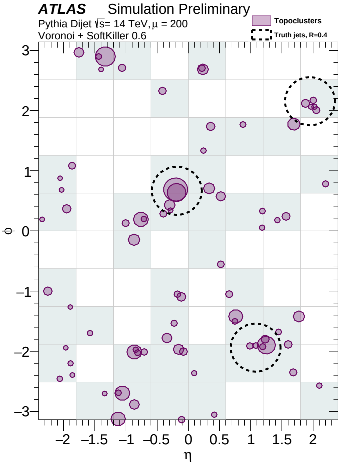

Revisiting the mass pileup dependence

UNIVERSITÉ DE GENÈVE

- Now you know the different columns
- CS+SK and VS+SK both help a lot
  - Recall that grooming suppresses both pileup and underlying event
  - By reducing pileup, grooming can focus on UE, the original intent for grooming
- Constituent mitigation techniques support the use of "loose" groomers
  - None of these algorithms is perfect
  - All choices will result in some loss of desired hard scatter energy
  - Looser combinations can mitigate this

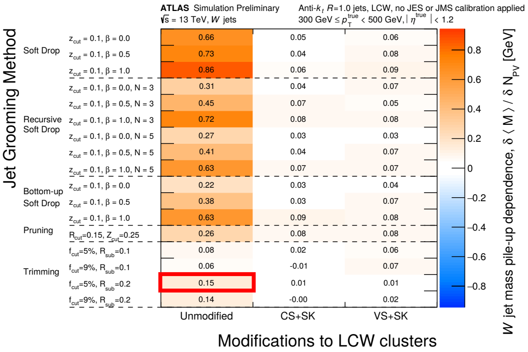

Why would we want looser groomers?

UNIVERSITÉ DE GENÈVE

- We will talk about taggers later
  - For now: the larger the number, the better we can identify W bosons
- Default ATLAS trimming is very strict
  - Result is pileup robust (last slide)
  - However, also reduces the W-boson identification potential
- Combining constituent mitigation with looser grooming can be beneficial

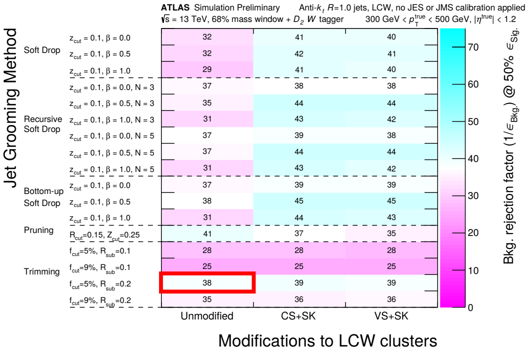

Another usage of tracks for jet reconstruction

UNIVERSITÉ DE GENÈVE

- As we discussed, tracks have a bad energy scale at high p_T and are thus not typically used
  - Related to it being hard to measure the curvature of a nearly-straight line
  - In contrast, the calorimeter works better at high p_T, so we use it more
- However, there is one very important reason to use tracks at high p_T
  - At very high p_T, the entire boosted jet may be a single cluster!
  - Tracks have a **very precise** spatial resolution and can identify the number of charged particles
  - Tracks can thus be used to "split" calorimeter clusters; ATLAS calls this "TCC"

Topo-clusters (EM or LC)
[Most ATLAS results]

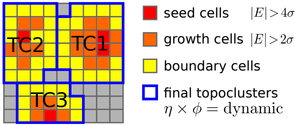

Particle Flow Objects (PFOs)
[Not yet used for R=1.0 jets]

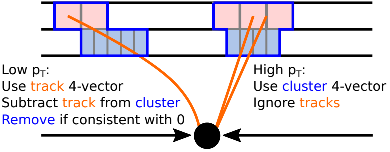

Track-CaloClusters (TCCs)
[Used in some high p_T R=1.0 results]

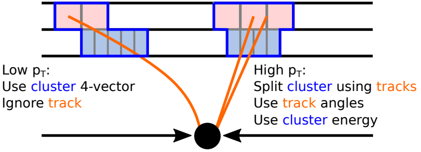

Tracks for high p_T large-R jets

- The ability to "split" high p_T calorimeter clusters using tracks is very important
  - Provides enormous improvements when trying to identify W bosons at very high p_T
- Not used by most groups in ATLAS as it's only needed for the highest p_T regime

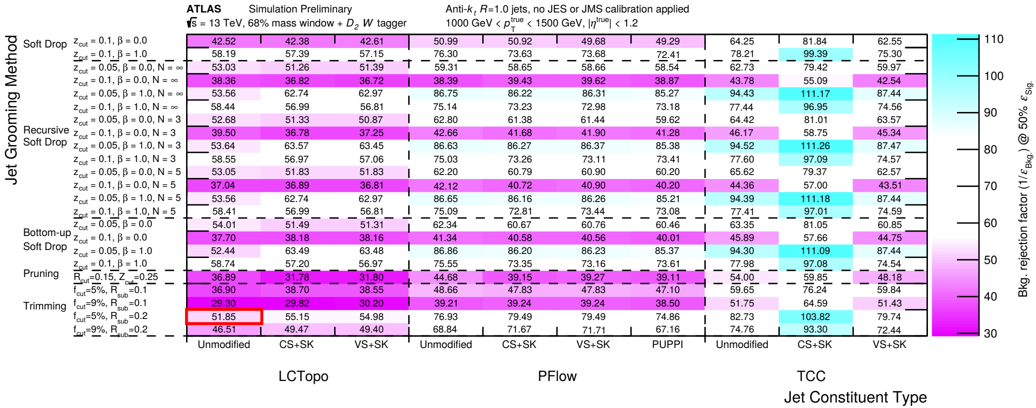

Large-R jet definitions

UNIVERSITÉ DE GENÈVE

- Large-R jets are almost exclusively used for boosted hadronic decays
  - Allows for reconstructing all of the decay products in a single object
  - The resulting jet has a well-defined mass expectation
- Large-R jets are very sensitive to pileup due to their size and use of individual angles
  - Grooming does reduce dependence on pileup and also the underlying event
  - Constituent-level pileup mitigation is another powerful approach
  - A combination of the two approaches is usually the best strategy
- Tracks can be useful at very high p_{T}, but not for their energy measurement
  - Instead, the spatial resolution of tracks can help to "split" calorimeter energy deposits

Jet mass scale and resolution

Calibrating the jet mass

UNIVERSITÉ DE GENÈVE

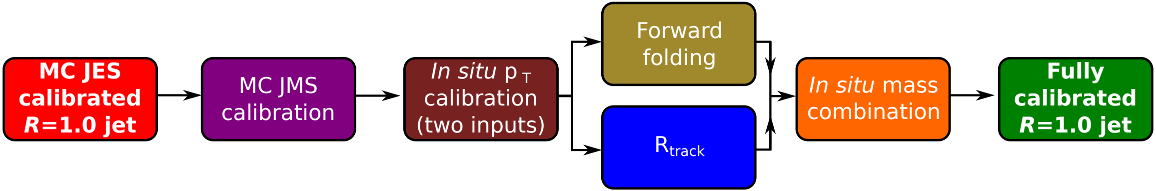

- Large-R jet calibrations generally follow what was done for $R = 0.4$ jets
  - The MC JES and *in situ* calibrations are derived in nearly the same way
- However, there are additional specific calibrations for the jet mass
  - They are derived and applied after the respective JES calibrations
  - MC JMS is MC only, forward folding and $\mathcal{R}_{\text{trk}}$ are data/MC comparisons

Mass scale and resolution

UNIVERSITÉ DE GENÈVE

- Exactly the same as for the JES/JER
  - Calculate the mass response
  - JMS = the mean of the Gaussian
  - JMR = the width of the Gaussian
- Mass response can be less Gaussian
  - Treating it as a Gaussian is usually still a good approximation for the scale
  - However, the resolution is increasingly being handled as a non-Gaussian

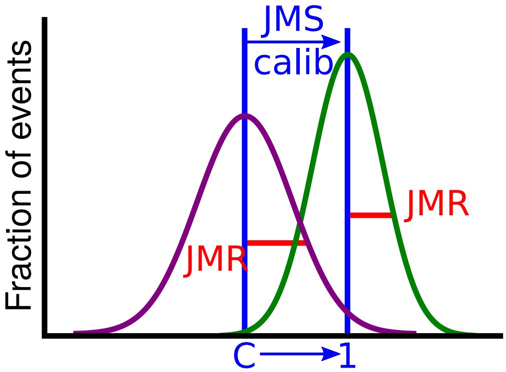

$$\text{Response} = \frac{m^{\text{reco}}}{m^{\text{true}}}$$

MC JES and JMS

UNIVERSITÉ DE GENÈVE

- The JES and JMS are linked: they both are related to the overall scale of the jet
- However, there are important differences related to detector geometry
  - JES is defined by the centroid of the jet and the surrounding detectors
  - JMS is defined by the width and angular profile of the shower
- Example: $\eta = 1.3$ results in a large JES
  - Quark jet: concentrated at $\eta \approx 1.3$
  - Two quarks from a W jet: likely on either side, at $\eta = 1.0$ and $\eta = 1.6$
  - QCD and W jets JMS can be different

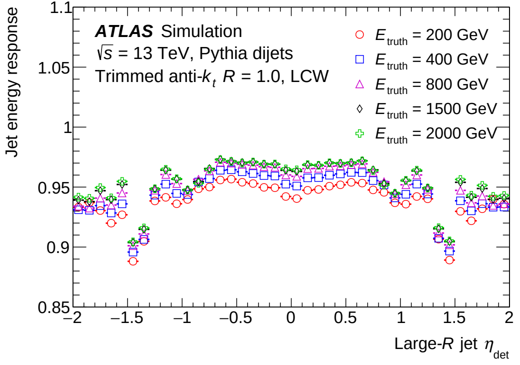

MC JMS

- The same JES is applied to all jets, then the JMS differs based on jet angular properties
  - Shown for W-mass jets (left) and top-mass jets (right), clear difference in required JMS
- ATLAS actually uses QCD jets for all of this and checks that it applies to real W/top later
  - Calorimeter responds the same to QCD and W/top, it just cares about the shower “width”

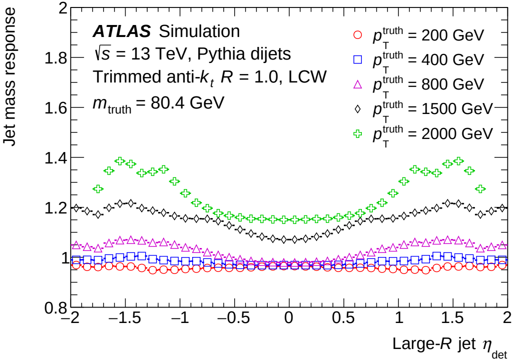

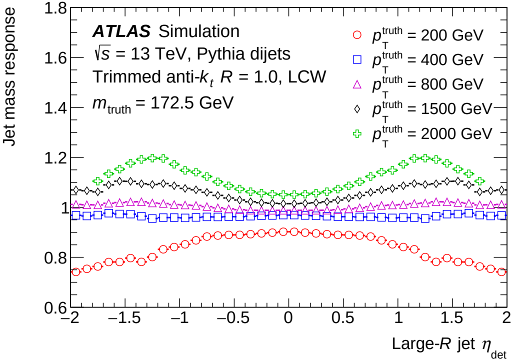

Forward folding

UNIVERSITÉ DE GENÈVE

1. Obtain a pure sample of high-mass jets
   - Here, we have used semi-leptonic $t\bar{t}$
2. Fit the MC to the data using:
   - $m^{\text{fold}} = s \cdot m^{\text{reco}} + (m^{\text{reco}} - m^{\text{truth}} \cdot \mathcal{R}_m)(r - s)$
   - s is the JMS, r is the JMR
   - $\mathcal{R}_m$ is the MC JMS

- This specifies how to modify MC to match data, both JMS and JMR
  - No assumption of a Gaussian JMR!
  - Can also anti-smear MC if needed

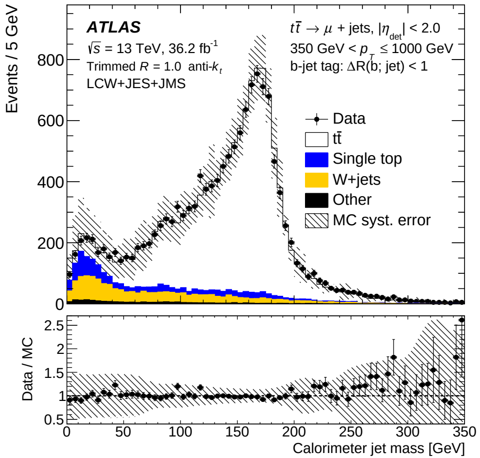

$\mathcal{R}_{\text{trk}}$

UNIVERSITÉ DE GENÈVE

- Forward folding requires a peak to fit
  - What about masses other than W/top?
  - Not enough stats to fit Higgs yet, but we do look for $H \to bb$ in a single jet
  - Need another method for such masses
- $\mathcal{R}_{\text{trk}}$ provides a rough fix
  - Use all tracks matched to the jet as an independent measurement of the scale
  - $r_{\text{trk}} = m^{\text{calo}}/m^{\text{track}}$, $\mathcal{R}_{\text{trk}} = r_{\text{trk}}^{\text{MC}}/r_{\text{trk}}^{\text{data}}$
  - Propagate track uncertainties to m^{calo}
- Only works if the jet mass is built using only the calorimeter
  - Otherwise not an independent measure

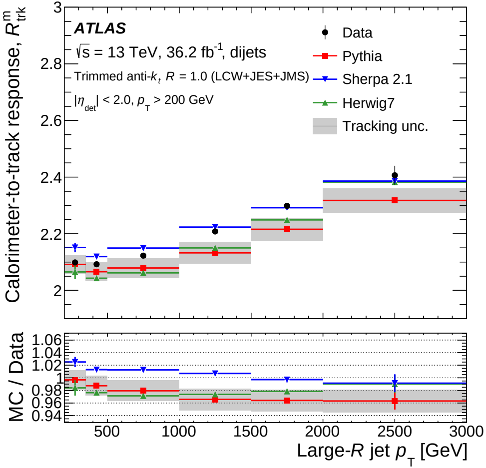

In situ JMS combination

- The forward folding and $\mathcal{R}_{\text{trk}}$ JMS are combined similar to what is done for the JES
  - Excellent control where forward folding is possible, larger uncertainties when only $\mathcal{R}_{\text{trk}}$
- Derived after the MC JMS and *in situ* JES: the *in situ* JMS is consistent with 1!

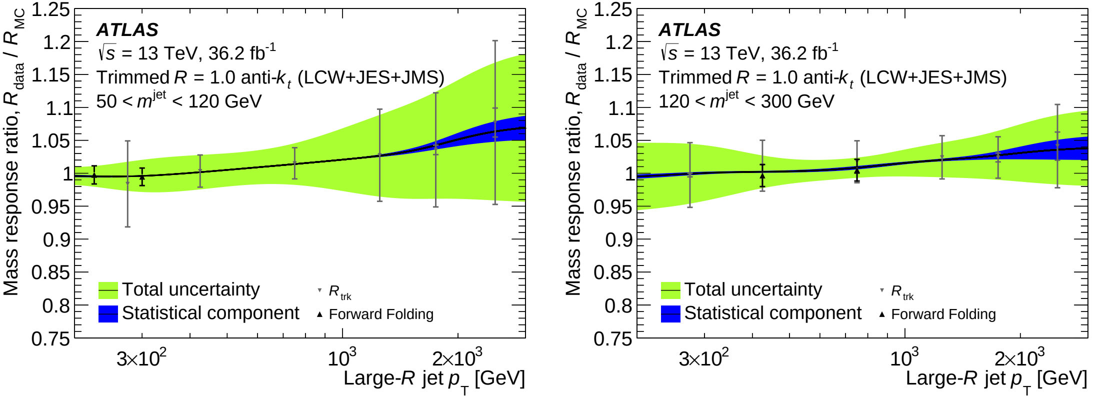

Large-R JMS calibrations

UNIVERSITÉ DE GENÈVE

- Large-R jets have similar JES calibrations as what we discussed for $R = 0.4$ jets
- They follow this up with dedicated mass calibrations
  - Important as the shower angular profile in the detector influences the calorimeter response

- Both MC-based and data/MC JMS calibrations are performed
  - Forward folding extracts the JMS and JMR, but is limited in range
  - $\mathcal{R}_{\text{trk}}$ is used to extend coverage to other mass ranges, but with larger uncertainties

Hadronic decay tagging

Hadronic decay tagging

UNIVERSITÉ DE GENÈVE

- Hadronic decays of massive particles can come from many sources
  - Main examples: $W \to qq'$, $Z \to qq$, $H \to b\bar{b}$, $top \to bW \to bqq'$
- We already saw the jet mass
- There are other *jet substructure* variables which quantify the energy and angular correlations within the jet
  - Together with the mass, these can be used to identify jets of interest

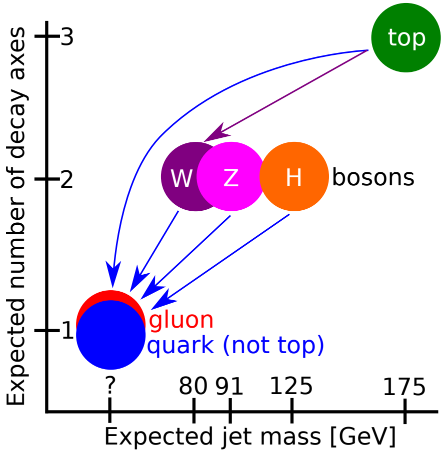

Designing W/Z/top taggers

- To first order, designing a tagger is as straightforward as the previous slide
  - Cut on the jet mass and other substructure variable(s) correlated to the number of decay axes
  - Common variables: $D_2^{(\beta=1)}$ for two-body decays (W/Z), τ_{32}^{WTA} for three-body decays (top)
  - You will use fastjet to calculate these variables in today's exercise
- Such a simple two-variable cut-based tagger provides a powerful starting reference

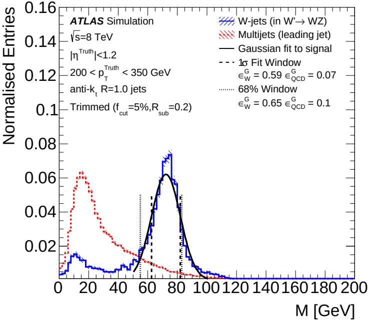

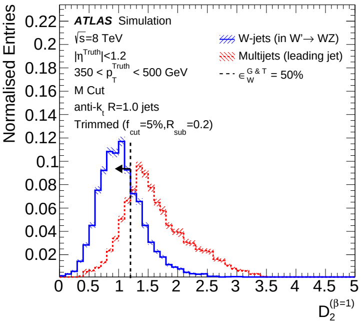

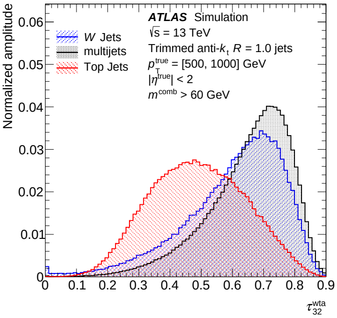

A brief aside: defining the tagger target

UNIVERSITÉ DE GENÈVE

- Just like for jet calibration, we need to define our target when designing a tagger
  - W/Z jet: jet is matched to a truth W/Z boson and both quark decay products ($q_1\&q_2$)
  - Top jet: jet is matched to a truth top quark and all three quark decay products ($b\&q_1\&q_2$)
- We typically optimize with respect to these "contained" truth labels, vs QCD background

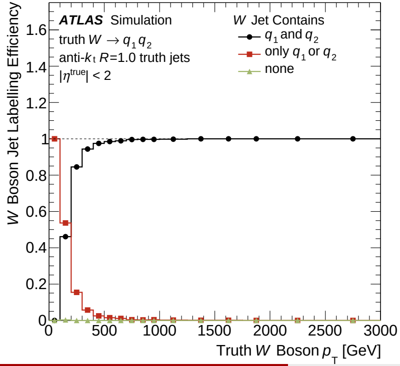

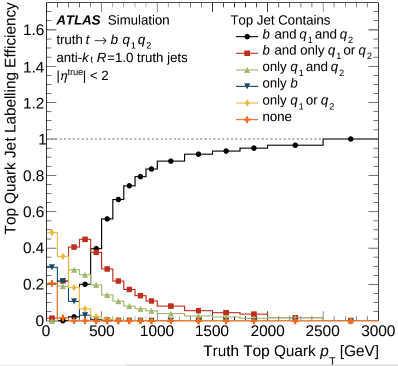

Tagging terminology

UNIVERSITÉ DE GENÈVE

- Jet tagging has a few key words that are important to know
- Tagging: identifying a jet as likely originating from a given type of particle
- Signal: the type of jet you want to keep, usually W/Z/H/top
- Background: the type of jet you want to get rid of, typically QCD
- Rejection of X: 1/X events survives the tagger (typically used for background)
- Efficiency of Y%: Y% of events pass the tagger (used for signal, sometimes background)

Simple two-variable taggers

UNIVERSITÉ DE GENÈVE

- As mentioned earlier, two-variable taggers are already quite powerful
  - W-tagging: mass+D_2 has $\sim 60\times$ background rejection $(1/60 = 1.7\%)$, 50% signal efficiency
  - Top-tagging: mass+τ_{32} (among others) has $\sim 5\times$ background rejection for 80% signal eff.
- However, we can definitely improve upon these simple taggers

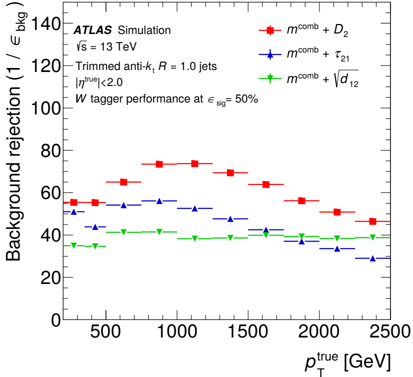

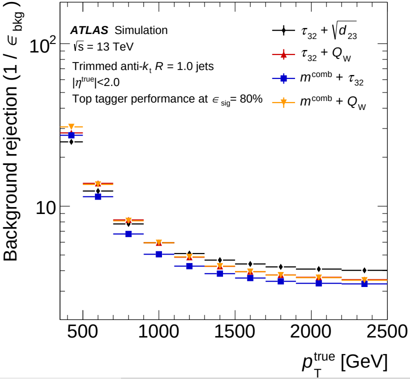

Moving towards machine learning

- Jet tagging is a very active domain for machine learning, which you saw earlier from Kyle
- Trained BDTs and DNNs for W/top-tagging with $\mathcal{O}(10)$ jet substructure variables
  - Small gain for W-tagging: $60\times \to 80\times$ QCD rejection for 50% signal eff ($\sim$ 30% gain)
  - Large gain for top-tagging: $5\times \to 12\times$ QCD rejection for 80% signal eff ($\sim$ 240% gain)
- Performance limited by amount of unique info: substructure variables are highly correlated

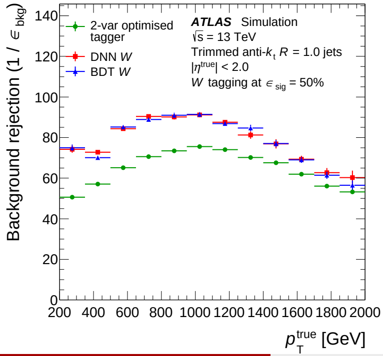

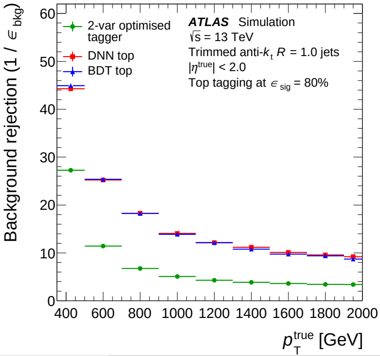

Going deeper into machine learning

UNIVERSITÉ DE GENÈVE

- Studied DNN with topocluster inputs, not substructure variables (TopoDNN)
  - Trained specifically for high p_T
  - Better than BDT,DNN in this regime
- Still, could be improved
  - Tagging power depends on jet inputs
  - Use of more precise inputs to DNN should provide further gains
  - Note: no b-tagging used so far

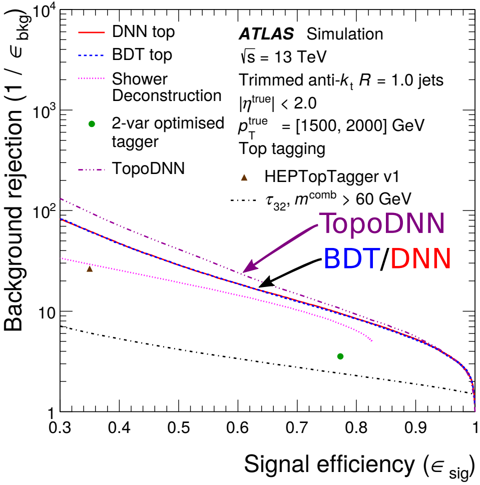

Controlling tagger performance in data and MC

UNIVERSITÉ DE GENÈVE

- Everything so far has been done based on MC optimization
  - No guarantee that it works the same in data, especially for complex ML taggers
  - We need a way to compare data and MC, like we had for *in situ* JES and JMS calibrations
- For taggers, this is done by considering tagging efficiency
  - Study the fraction of events that pass or fail a tagger, in both data and MC
  - This has to be done separately for both signal and background jets
- Studying background efficiency is pretty straightforward, as QCD events are numerous
- Studying signal efficiency is harder, as we only have a few relevant processes
  - We will often need to study one kinematic regime and extrapolate to another regime

Deriving tagging efficiency signal scale factors

- Semi-leptonic $t\bar{t}$ events are very pure in hadronic W and top jets
  - We can use these events to derive the tagging efficiency in data and compare to simulation

$200\ \text{GeV} < p_{\text{T}}^{\text{W}} \lesssim 500\ \text{GeV}$

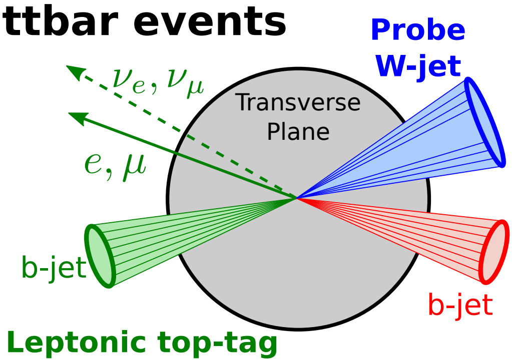

$350\ \text{GeV} < p_{\text{T}}^{\text{top}} \lesssim 1000\ \text{GeV}$

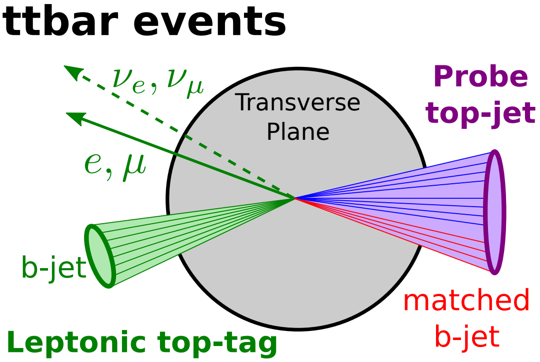

Scale factor extraction regions

UNIVERSITÉ DE GENÈVE

- The selected W- and top-candidate pre-tag mass distributions show high signal purity
  - Fit three templates: $t\bar{t}$ signal process, $t\bar{t}$ background processes, and non-$t\bar{t}$ backgrounds

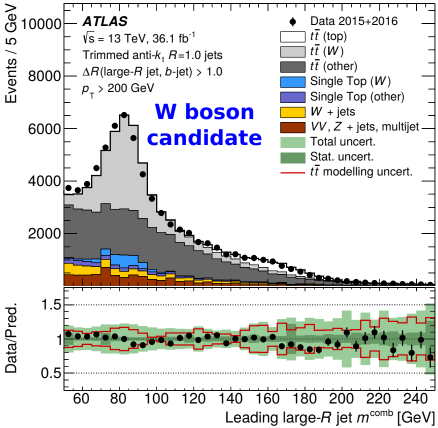

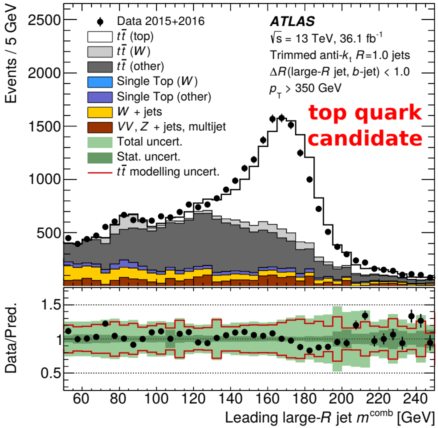

Extracting the efficiency

- Define templates from the shapes in MC, but allow their normalizations to float
- Fit them to the pass and fail distributions to extract the number of signal events in data
- $\epsilon^{\text{signal}}_{\text{data}} = \frac{N^{\text{tagged}}_{\text{fit,signal}}}{N^{\text{tagged}}_{\text{fit,signal}} + N^{\text{not tagged}}_{\text{fit,signal}}}$, $\epsilon^{\text{signal}}_{\text{MC}} = \frac{N^{\text{tagged}}_{\text{signal}}}{N^{\text{tagged}}_{\text{signal}} + N^{\text{not tagged}}_{\text{signal}}}$, scale factor $= \epsilon^{\text{signal}}_{\text{data}} / \epsilon^{\text{signal}}_{\text{MC}}$

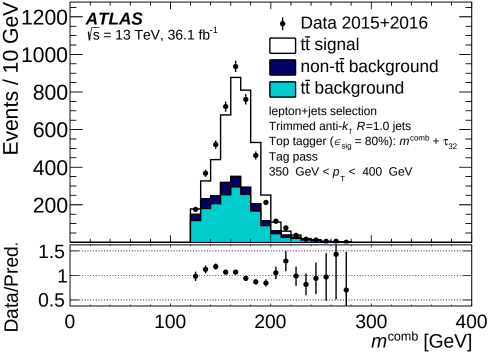

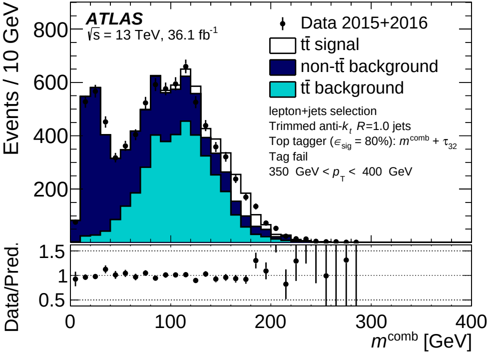

Resulting signal scale factors (p_T sliced)

UNIVERSITÉ DE GENÈVE

- Data and simulation agree within data statistical uncertainties, even for complex taggers
  - Shown for the W-tagging DNN and the top-tagging topocluster-based DNN
  - Agreement also stable against large range of pileup (not shown here, but important to check)

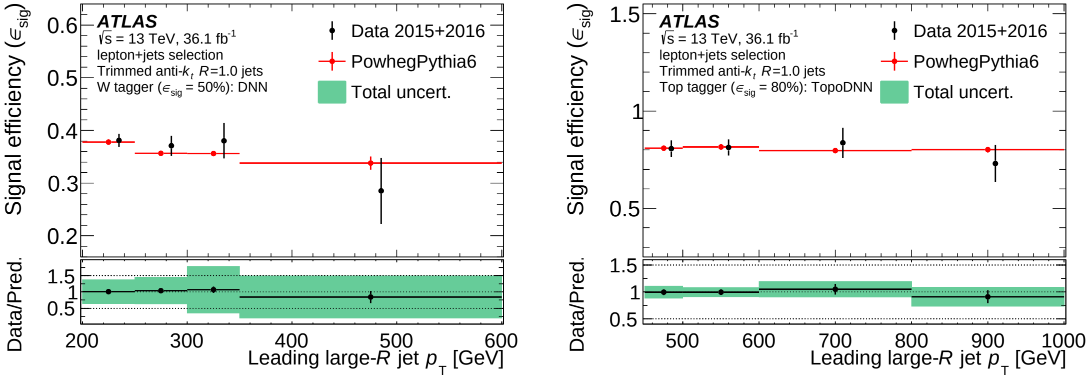

A different way to extract W/Z signal efficiency

- V+jets events can also be used to evaluate the tagger efficiency
- Extends measurement to higher p_T, but much harder as background dominates signal
  - Need to be careful: breaks down if signal and background peak at the same mass

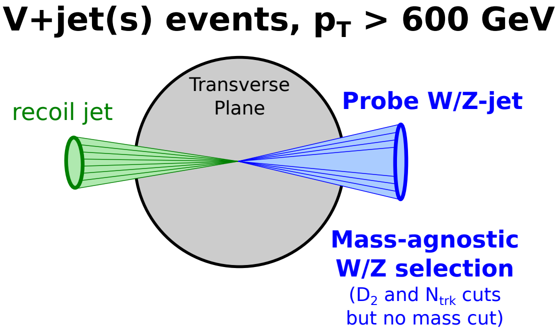

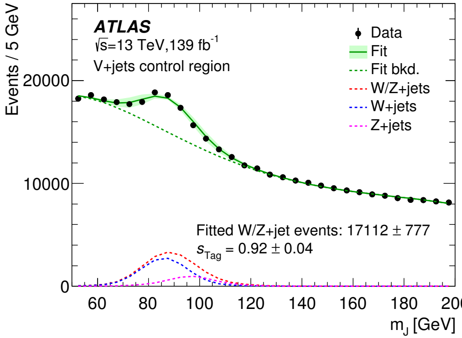

A challenge of ML taggers

- Machine learning is great, but it can learn the expected jet mass shape
  - Improves tagging performance
  - However, this results in QCD and signal have the same mass spectrum
- This breaks the V+jet signal efficiency extraction approach
  - QCD is much more common and has the same shape, so can't fit W/Z peak
  - $t\bar{t}$ approach still works (signal-pure)
- It is possible to get around this by telling ML not to learn the mass shape
  - Mass decorrelated taggers, not covered

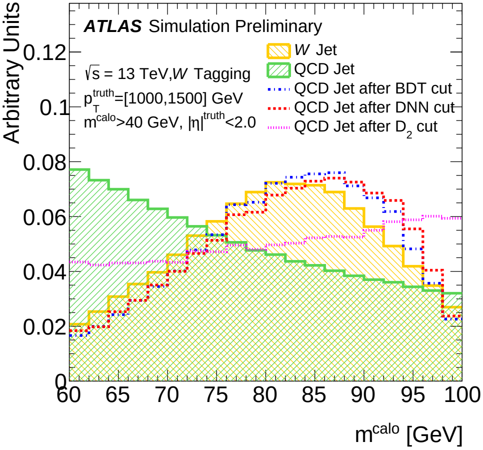

Extracting background scale factors

UNIVERSITÉ DE GENÈVE

- Much simpler than for the signal efficiency, as the sample is $\sim 99\%$ pure in background
  - $\epsilon^{\text{background}}_{\text{data}} = \frac{N^{\text{tagged}}_{\text{background}}}{N^{\text{total}}_{\text{background}}}, \quad \epsilon^{\text{signal}}_{\text{MC}} = \frac{N^{\text{tagged}}_{\text{background}}}{N^{\text{total}}_{\text{background}}}, \quad \text{scale factor} = \epsilon^{\text{background}}_{\text{data}} / \epsilon^{\text{background}}_{\text{MC}}$
- Study performance of taggers in background samples: QCD multijets and γ+jet
  - Allows for studying modelling differences between gluons (QCD) and light quarks (γ+jet)

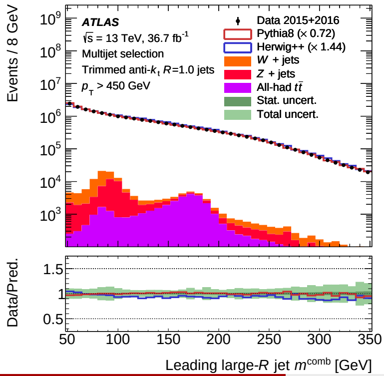

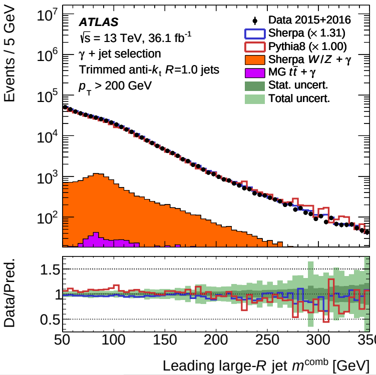

Resulting background scale factors

UNIVERSITÉ DE GENÈVE

- As before, these are for the W-tagging DNN and the top-tagging topocluster-based DNN
  - Multijet (below): W- and top-tagging both agree with Pythia8, disagree with Herwig++
  - γ+jet (backup): W-tagging agrees with Sherpa not Pythia8, top-tagging agrees with both
- The observed differences between generators are taken as an uncertainty

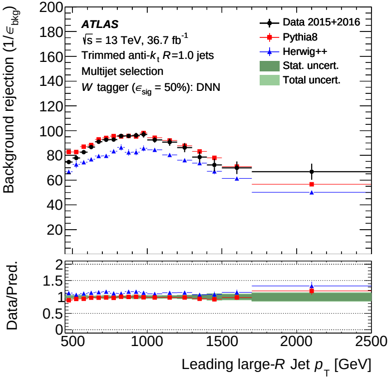

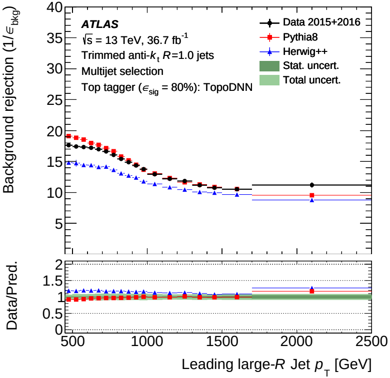

Designing taggers and extracting scale factors

UNIVERSITÉ DE GENÈVE

- Jet substructure is a powerful tool for identifying hadronically decaying massive particles
- Jet taggers continue to improve dramatically
  - Started out as simple but powerful two-variable taggers (mass and another variable)
  - Increasingly the playground of advanced machine learning techniques
- Tagging efficiency can be compared between data and MC in well-identified scenarios
  - Works well for some kinematic regimes for W, top, and Z
  - Other particles and kinematic regimes require extrapolations and larger uncertainties
- Jet taggers can be very sensitive to the jet definition
  - Grooming strategy, constituent-level pileup mitigation, and jet inputs all matter
  - Important to keep the full picture in mind when deciding on a jet definition and tagger

Summary

Summary

UNIVERSITÉ DE GENÈVE

- Large-R jets are an increasingly important topic and use case
  - At the LHC, W/Z/H bosons and top quarks often have enough p_{T} to end up in a single jet
- There are three main experimental aspects to working with large-R jets
  1. Identifying an optimal jet definition: grooming, constituent-modifiers, and jet inputs
  2. Calibrating the mass scale and resolution of the large-R jet
  3. Designing advanced taggers to identify W/Z/H/top decays, and then evaluating the tagger efficiency for both signal and background in data and MC
- Today's exercise will look at large-R jets in more detail
  - You will be building, grooming, and calculating substructure variables for large-R jets
  - This will allow you to understand how jet definitions impact key large-R jet observables

- Large-R jet reconstruction and tagging is rapidly evolving, and is increasingly using ML
  - This is only the start - there are lots of open opportunities to develop new ideas!