

QCD Jet Measurements at ATLAS

Inclusive dijet and $W + b$ -jets cross-sections, and unbinned Z +jets in the charged-particle phase space

Hanfei Cui

on behalf of the ATLAS Collaboration


28 September 2026

QCD@LHC 2026, Oxford




Three Run 2 Jet Analyses

All three use **140 fb⁻¹ of Run 2 data**, unfolded to particle level.




Inclusive dijet cross-sections
Double-differential in
(m_{jj} , y^*) and (m_{jj} , y_{boost}).

arXiv:2512.19073 (EPJC)



W + b-jets cross-sections
Adds jet flavour tagging, and probes
heavy quarks in the initial state.

arXiv:2605.31205 (JHEP)




Z+jets OmniFold
All charged particles are unfolded, so
the observables can be chosen after
unfolding.

ATLAS-CONF-2026-009

Inclusive dijet cross-sections

arXiv:2512.19073, EPJ C

Dijet: two rapidity variables




- $y^* = |y_1 - y_2|/2$ sets the parton scattering angle:
 $\tanh y^* = \cos \theta^*$
- $y_{\text{boost}} = |y_1 + y_2|/2$ gives the boost of the pair, so it scans x_1/x_2 : $x_{1,2} = (m_{jj}/\sqrt{s}) e^{\pm y_{\text{boost}}}$
- Small $y_{\text{boost}} \rightarrow x_1 \approx x_2$; large $y_{\text{boost}} \rightarrow$ asymmetric x ; combined with m_{jj} , adds PDF sensitivity
- **First ATLAS dijet cross-section binned in y_{boost}**

Phase space


- ≥ 2 jets, $p_T > 75$ GeV, $|y| < 3.0$
- $H_{T,2} = p_T^1 + p_T^2 > 200$ GeV (fixed-order stability)
- $y^*, y_{\text{boost}} \in [0, 3]$, six bins of 0.5
- m_{jj} : 240 GeV to ~ 10 TeV; 237 (y^*) + 160 (y_{boost}) bins
- highest dijet mass observed: **9.54 TeV**

Dijet: experimental uncertainties


first y^* bin (most precise)



last y_{boost} bin (least precise)





against early Run 2, $y^* < 0.5$





- Single-jet triggers (60–420 GeV), 99.5% efficient
- JES dominates: **1172** systematic components, plus 34 for JER
- IDS unfolding: 1 iteration (y^*), 2 (y_{boost})
- New calibration (EPJC 85 (2025) 927): JES **2–3× smaller** than early Run 2
- Total: **2–4%** (best) to 22% (worst)

Dijet: PDF uncertainties





- NNLOJET full-colour NNLO, $\mu_R = \mu_F = m_{jj}$
- Five PDF sets; $\alpha_s(m_Z) = 0.118 \pm 0.001$
- Highest m_{jj} : CT18 **15–22%**, NNPDF4.0 **2%**
- Scale variation dominates NNPDF4.0; PDF uncertainty dominates CT18

Dijet: electroweak and non-perturbative corrections





- EW (Dittmaier et al., JHEP 11 (2012) 095): +10% at low y^* , high m_{jj} ; within 3% in y_{boost}
- NP: Pythia 8 / Herwig 7 spread, < 5% EW and NP corrections, all bins
- χ^2 with EW: $y^* 396 \rightarrow 329$ for 237 bins; $y_{\text{boost}} 184 \rightarrow 183$
- y_{boost} : **no sensitivity to either correction**

Dijet: measured spectra





Slices scaled by 10^{-3n} ; NNLO CT18 $\times$ EW $\times$ NP bands over the data.
Data uncertainties smaller than theory almost everywhere.

Dijet: NLO and NNLO against the data



- NLO and NNLO **15–20% above data**, after EW and NP corrections
- y^* : ratio falls with m_{jj} at low y^* , rises at high; NNLO improves shape only at low y^*
- y_{boost} : NNLO improves every m_{jj} bin for $y_{\text{boost}} > 1.5$ (low- y^* events dominate)
- At fixed m_{jj} , y_{boost} scans x_1/x_2 : new PDF constraint

Dijet: comparison with five PDF sets



Global χ^2/ndf	ATLASpdf21T3	CT18	NNPDF3.1	NNPDF4.0	MSHT20
(m_{jj}, y^*) , ndf 237	288	329	343	342	347
$(m_{jj}, y_{\text{boost}})$, ndf 160	166	183	178	185	195

- Lowest χ^2 : ATLASpdf21T3; NNPDF3.1 best global set in y_{boost}
- NNPDF4.0 closest centrally; narrow 2% bands raise χ^2 [▶ per-bin tables](#)
- **No set removes the 15–20% offset**, which remains after EW and NP corrections
- Candidates: fixed order, $\mu = m_{jj}$, missing resummation, small $R = 0.4$ ($\alpha_s \ln(1/R)$)

Take-away: Data reach 2–4%; theory 15–20% off after EW and NP. Theory now limits.

$W + b$ -jets cross-sections

arXiv:2605.31205, JHEP

$W + b$: heavy quarks in the initial state

Why measure it


- pQCD with initial-state heavy quark; dominant: $W+g(\rightarrow b\bar{b})$
- Background to $VH(b\bar{b})$, single top, BSM
- **Open since Run 1:** ATLAS at 7 TeV sat 1.5σ above NLO

Selection


Lepton	$e/\mu, p_T > 28 \text{ GeV}$
E_T^{miss}	$28 < E_T^{\text{miss}} < 60 \text{ GeV}$
m_T	$> 65 \text{ GeV}$
Jets	anti- k_t , $R = 0.4$ PFlow, $p_T > 25 \text{ GeV}$ $ \eta < 2.5$, jet-vertex tagged
b -tag	DL1r, 60% working point
Regions	1j1b and 2j1b

Rejects c by 40, light by 2500. Upper E_T^{miss} cut suppresses top; events with ≥ 2 b -jets or ≥ 3 jets vetoed.

$W + b$: flavour fit and unfolding




fitted scale factors, 1j1b



b -jet p_T after the fit, electron 2j1b


- **Flavour fit:** DL1r score, two bins (20–60%, 0–20%); p_0 for $W+b$, p_1 for $W+c/\text{light}$
- All scale factors **consistent with 1** ▶ all four channels
- Multijet and fakes from matrix method
- Unfolding: matrix inversion (total), 5 Bayesian iterations (spectrum)
- BLUE combination: $\chi^2/\text{ndf} = 0.2/2$ and $2.0/8$

W + b: cross-sections and comparison with predictions



$\sigma_{\text{fid}}(1+2 \text{ jets}) = 16.6 \pm 0.2 (\text{stat}) \pm 1.9 (\text{syst}) \text{ pb}$

[pb]	1 jet	2 jets	1+2
Data	9.1 ± 0.9	7.5 ± 1.1	16.6 ± 1.9
Sherpa 2.2.11	9.0 ± 0.9	7.8 ± 1.3	16.8 ± 2.3
MG5+Py8 FxFx	7.4 ± 0.9	6.5 ± 0.6	13.9 ± 1.3




- Systematics dominate (11.6% vs 0.9% stat); b -tagging **8.3%**, top up to 13.8% [breakdown](#)
- Sherpa agrees; MG5+FxFx 1σ low in every bin
- 2 $\times$ precision of 7 TeV; **Run 1 excess not reproduced**
- Data and 5FNS NLO both at $\sim 10\%$; needs NNLO and 4FNS, not more data

Unbinned Z+jets

fully differential in the **charged-particle** phase space


ATLAS-CONF-2026-009 (preliminary)

Z+jets: an unbinned charged-particle release



- Binned unfolding: observables and bins fixed before results seen
- Published as weighted events: **162 variations** (27 syst, 135 bootstrap)
- σ for any cut = sum of weights that pass
- Jet algorithm, radius, binning all free; **no new calibration**
- Each new observable checked on pseudodata (≥ 5000 eff. events/bin)
- **First published unbinned, fully differential collider cross-section in the charged-particle phase space.**
ATLAS-CONF-2026-009, preliminary

Z+jets: Omnifold at ~150 dimensions




Measurement

- $Z \rightarrow \mu\mu$, 140 fb^{-1} , $\sim 247\text{k}$ events
- $\mu^+ \mu^-$: $p_T > 25 \text{ GeV}$, $|\eta| < 2.4$, $81 < m_{\mu\mu} < 101 \text{ GeV}$
- $p_T^{\mu\mu} > 200 \text{ GeV}$, **no jet cut**
- Charged particles: $p_T > 0.5 \text{ GeV}$, $|\eta| < 2.5$, $c\tau_0 > 10 \text{ mm}$
- Top background 0.25%, subtracted unbinned


Method

- **~ 150 dimensions** (50 particles $\times$ 3 coords); up to 750
- Particle Transformer, 1.6M params, trained at $p_T^{\mu\mu} > 190 \text{ GeV}$
- **Pretraining on 14M events**: cuts bias and run-to-run variation
- 400k A100 GPU-hours (Perlmutter)

Z+jets: validation and uncertainties





closure test on simulated data, n_{ch} : $\chi^2 = 2.41$, $p = 0.790$



IBU, Multifold and Omnifold on $p_T^{\mu\mu}$



- Closure on **Sherpa** pseudodata, kernel-reweighted (not the network)
- **25 of 26** observables pass; m_{j1} at $p = 0.016$
- Matches IBU and Multifold (PRL 133 (2024) 261803)
- Dominant syst: tracking, muon isolation; unfolding comparable. H_T : 7, 2, 4, 5, 7, 12% [► breakdown](#)
- **150-D costs almost nothing** vs 1-D; more precise for m_{j1}, m_{j2} [► vs IBU](#)

Application 1: observables outside the unfolding



- n_{ch} and H_T not in the unfolding; computed from published events
- n_{ch} : **50% disagreement**, generators in opposite directions; constrains tuning
 → Soft QCD measurements with pp collisions in ATLAS (Pretel, Mon 14:00)
- The MG5+Py8 H_T shape is also off, by less


Application 2: three jet algorithms from one dataset



- $\langle \rho(r) \rangle$ for k_t , anti- k_t and C/A at $R = 0.4$, same events, no extra calibration. Jets $p_T \in [120, 150] \text{ GeV}$, $|y| < 1$
- Generators too narrow: excess at $r < 0.04$, deficit at $0.04-0.3$
- anti- k_t : bump near $r \approx R$ (soft edge); k_t and C/A track each other



Take-away: This measures radiation inside the jet; the unbinned release allows any R and any jet algorithm.

Application 3: energy correlator with no jet radius



- No jet-edge cut-off; six decades in $z = (1 - |\cos \theta_{jj}|)/2$ at $\leq 5\%$. **First all-charged-particle EEC in pp**
→ EEC, Track functions and related topics (Waalewijn, Tue 10:00)
→ Accounting for the Effect of Fiducial Cuts on ... (Vorona, Tue 16:20)
- **Three regimes:** wide-angle UE ($z > 10^{-2}$), pQCD ($10^{-4} - 10^{-2}$), confinement ($z < 10^{-4}$)
- Both generators miss at small z , Sherpa by the larger amount

Application 4: an observable with no response matrix



- NN intrinsic dimension vs median EMD; anti- k_t $R = 1.0$, $p_T \in [330, 370] \text{ GeV}$
→ Event isotropy in perturbative QCD (Atzori, Tue 17:00)
- **First collider measurement of intrinsic dimension**
- Defined for the dataset, not per event; binned unfolding cannot measure it
- N -particle jet spans $3N$ dimensions, but **< 10 effective degrees of freedom** at every scale
- Generators agree; MG5+Py8 deviates near minimum

What limits each measurement

Measurement	Precision	Limited by	Needs
Inclusive dijet cross-sections	2-4%	Theory	resummation, a better scale, the R dependence → Accurate predictions for LHC Processes ... (Andersen, Mon 16:20)
$W + b$ -jets cross-sections	11.6%	Detector	NNLO, and a 4FNS prediction → Matching flavour schemes via isolation of ... (Gauld, Thu 12:30)
Z +jets, unfolded charged particles	2-12%	Method	neutral particles, a goodness-of-fit test

Summary

- Measurement of inclusive dijet cross-sections in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector

arXiv:2512.19073 (EPJC)

- ▶ Full Run 2, double-differential; first ATLAS y_{boost} binning
- ▶ Data reach 2–4%
- ▶ NNLO **15–20% above** after EW and NP corrections; no PDF set removes it

- Measurement of the cross-section for the production of a W boson in association with b -jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector

arXiv:2605.31205 (JHEP)

- ▶ $\sigma_{\text{fid}} = 16.6 \pm 1.9$ pb
- ▶ Sherpa agrees; MG5+FxPx 1σ low
- ▶ Twice the 7 TeV precision; Run 1 excess not reproduced

- A fully differential charged-particle phase-space measurement of Z +jets production in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector

ATLAS-CONF-2026-009 (preliminary)

- ▶ Particle-level events with unfolded weights, instead of histograms
- ▶ Observable, jet algorithm and binning **chosen after publication**; no new calibration
- ▶ Code and data: gitlab.cern.ch/atlas-physics/sm/photons-jets/ana-stdm-2024-02

Thank you!

Questions?

Hanfei Cui

hanfei.cui@cern.ch

University College London

Backup



Backup: dijet χ^2 tables

	(m_{jj}, y^*)				
y^*	CT18	NNPDF 3.1	NNPDF 4.0	MSHT 20	ATLAS pdf21T3
< 0.5	58/58	68/58	59/58	61/58	58/58
0.5–1.0	62/54	70/54	63/54	67/54	59/54
1.0–1.5	56/45	67/45	63/45	65/45	58/45
1.5–2.0	51/36	59/36	57/36	54/36	63/36
2.0–2.5	30/28	41/28	36/28	34/28	39/28
2.5–3.0	17/16	21/16	17/16	21/16	19/16
all	329 /237	343 /237	342 /237	347 /237	288 /237

	$(m_{jj}, y_{\text{boost}})$				
y_{boost}	CT18	NNPDF 3.1	NNPDF 4.0	MSHT 20	ATLAS pdf21T3
< 0.5	33/36	48/36	37/36	43/36	44/36
0.5–1.0	51/34	55/34	53/34	55/34	54/34
1.0–1.5	41/31	40/31	39/31	47/31	47/31
1.5–2.0	26/26	27/26	26/26	31/26	35/26
2.0–2.5	24/20	24/20	21/20	28/20	31/20
2.5–3.0	15/13	15/13	14/13	20/13	16/13
all	183 /160	178 /160	185 /160	195 /160	166 /160


- ATLASpdf21T3 publishes no α_s uncertainty, so that term is left out. Including it would lower its χ^2 further
- Bin by bin the agreement is good. The global $y^* \chi^2$ gives $p \sim 10^{-6}$, which the paper attributes to **bin-to-bin fluctuations in the perturbative prediction**, from large real-real and real-virtual integrals that do not quite cancel, rather than to a physical effect
- The χ^2 is insensitive to an overall shift, because systematics correlated across bins absorb it. That is why the χ^2 passes while the 15–20% offset remains
- Drop the corrections (NNLO, CT18) and (m_{jj}, y^*) moves from 329 to 396 without EW and to 338 without non-perturbative corrections. $(m_{jj}, y_{\text{boost}})$ is unaffected

Backup: dijet systematics, remaining extreme bins



- At high m_{jj} the η -intercalibration leads every y_{boost} bin and the high- y^* bins; at low y^* the response to single hadrons leads instead
- At low m_{jj} pile-up leads at low y^* and low y_{boost} , and the η -intercalibration takes over as either grows
- Statistical errors come from 1000 bootstrap replicas. For the JES, JER, Tile and unfolding terms, bins are merged until each is worth more than 2σ , then smoothed with a Gaussian kernel


Backup: dijet electroweak and non-perturbative corrections




- $NP = \sigma^{PS+HAD+MPI}/\sigma^{PS}$. Pythia 8.210 with NNPDF2.3 and A14 gives the central value; the spread over all seven generator and tune pairs gives the uncertainty
- Pythia 8.210 (Lund string) with CTEQ6L1/C4, CTEQ6L1/AU2, CT10/AU2, NNPDF2.3/A14; Herwig 7.2.3p2 (cluster) with NNPDF2.3 and H7.1-Default, -BaryonicReconnection, -Soft
- They reach 5% at low dijet mass and fall away at high m_{jj} in every y^* bin. In the lower y_{boost} bins they stay flat

Backup: dijet theory uncertainties, remaining bins


CT18, last y^*




CT18, first y_{boost}




CT18, last y_{boost}




NNPDF4.0, last y^*









NNPDF4.0, first y_{boost}



NNPDF4.0, last y_{boost}



Backup: NLO vs NNLO, all six y^* bins









ATLAS
Particle level

$\sqrt{s} = 13 \text{ TeV}, 140 \text{ fb}^{-1}$
anti- k_t $R=0.4$
 $\mu_R = \mu_F = m_{jj}$
 $\alpha_s(m_Z) = 0.118$

pQCD
⊗ EW ⊗ NPC

Backup: NLO vs NNLO, all six y_{boost} bins





ATLAS
Particle level

$\sqrt{s} = 13 \text{ TeV}, 140 \text{ fb}^{-1}$
anti- k_t $R=0.4$

$\mu_R = \mu_F = m_{jj}$
 $\alpha_s(m_Z) = 0.118$

pQCD
 $\otimes \text{EW} \otimes \text{NPC}$

Backup: PDF comparison, all six y^* bins





Backup: $W + b$ uncertainties and correlations

Relative uncertainties [%]


Source	1 jet	2 jets	1+2
b -tagging	7.3	10.2	8.3
Top background	3.1	7.0	4.3
E_T^{miss}	2.6	4.7	3.5
Jet p_T	3.5	3.0	3.1
Multijet bkg	3.3	3.2	3.2
Muon	2.2	3.3	2.7
Luminosity	1.3	2.2	1.7
Electron	1.2	1.7	1.4
MC statistics	0.9	1.2	0.9
V +jet scale	0.6	1.5	0.8
V +jet PDF	0.6	0.5	0.4
Unfolding	0.4	0.6	<0.1
JVT	<0.1	<0.1	<0.1
Total syst.	10.1	14.7	11.6
Total stat.	0.8	1.2	0.9

Highest p_T bin (60–140 GeV, 1+2 jet): top background **13.8%**,
 b -tag 9.6%, E_T^{miss} 8.8%, jet p_T 7.0%, total syst. **21.7%**




The correlations are large because systematics dominate and most of them move together across bins.

Backup: $W + b$ post-fit distributions, all four regions




e 1j1b, e 2j1b (top); μ 1j1b, μ 2j1b (bottom)




fitted scale factors, 2j1b

- $p_0 = 1.02 / 0.99 / 0.99 / 0.97$ and $p_1 = 0.93 / 1.04 / 0.93 / 1.02$ over the four regions, with ± 0.01 from statistics and ± 0.15 – 0.26 from systematics. All are consistent with unity
- For $W + c$ and $W + \text{light}$ the 0–60% calibration is used in both tagging bins, with a 100% uncertainty added in the 0–20% one
- Fake-lepton rates come from a control region with the cuts inverted, $E_T^{\text{miss}} < 28$ GeV and $m_T < 40$ GeV, after the prompt leptons are subtracted using MC
- The fit errors come from 1000 bootstrap replicas; systematics are evaluated by shifting one source at a time

Backup: $W + b$ differential cross-sections per jet region



1-jet region




2-jet region

$p_T^{b\text{-jet}}$ [GeV]	25-30	30-40	40-60	60-140
Data, 1+2 jet [pb]	4.64 ± 0.72	5.02 ± 0.48	4.02 ± 0.44	2.84 ± 0.62
Sherpa 2.2.11	4.49 ± 0.46	5.16 ± 0.64	4.13 ± 0.62	2.67 ± 0.47
MG5+Py8 FxFx	3.75 ± 0.39	4.20 ± 0.42	3.31 ± 0.32	2.27 ± 0.20


- The spectrum stops at 140 GeV because the top-quark background dominates above it
- In Sherpa, 0.98 / 0.92 / 0.87 / 0.84 of 2j1b fiducial events hold exactly one b -jet, so the “ $W +$ one b -jet + one other jet” cross-section can be recovered

Backup: Z+jets uncertainty breakdown

H_T : 13-component breakdown




H_T : uncertainty correlation matrix




- From the detector: muon efficiency and calibration, track reconstruction, pile-up, luminosity. From theory: Drell-Yan PDF, α_s , QCD scale and tune. The diboson yield is varied by 30% and the EW Zjj yield by 20%
- From the unfolding: training noise over 10 repeats, method bias taken from the data, the hadron mix (pions, kaons, protons and the rest reweighted to Sherpa), and the hidden-variable term, obtained by repeating the whole procedure with Sherpa as the MC
- For most observables the detector terms are larger than the modelling ones

Backup: Omnifold vs dedicated IBU, uncertainties

total fractional uncertainty: n_{ch} and H_T




where Omnifold beats IBU: m_{j1} and m_{j2}




- On H_T Omnifold is more precise than IBU in the first four bins and less precise in the last two, where the statistics run out: 0.072 / 0.011 / 0.034 / 0.048 / 0.064 / 0.121 against 0.085 / 0.020 / 0.047 / 0.056 / 0.061 / 0.079
- On the two jet masses Omnifold is the more precise of the two
- n_{ch} runs the other way: Omnifold is slightly worse in the first five bins and better in the last. The penalty is small in every case
- Multifold measured neither n_{ch} nor H_T , so only IBU compares here. On the other 24 observables all three agree

Backup: Omnifold vs IBU cross-sections, and the top subtraction

cross-sections: n_{ch} and H_T



top-background subtraction weights



- Top is the only background worth counting, 0.25% of the yield. It is removed by event-by-event classifier reweighting rather than assigned as a systematic uncertainty, which is why this total is smaller than Multifold's
- The subtraction uses five variables (a dedicated $t\bar{t}$ classifier output, n_{ch} , H_T , p_T^{j1} and m_{j1}), because over the full phase space the network could not pick out a component worth 0.25%
- Powheg+Pythia gives $\langle w \rangle = 0.9973, \sigma_w = 0.0059$; Sherpa gives 0.9975 and 0.0008. They remove the same total but spread it differently over events, and **that gap is the subtraction uncertainty**