

Recent developments in shower matching

status and outlook


James Whitehead

Jagiellonian University, Kraków

QCD@LHC 2026 · Oxford, 30 September 2026



JAGIELLONIAN UNIVERSITY
IN KRAKÓW



- **Exclusive, hadron-level events**
detector simulation, fiducial cuts, unfolding, ML training
- **The best available perturbative precision**
fixed order for rates and hard tails; resummation where logs are large; uncertainty variations as event weights
- **The jet multiplicities analyses use**
 V +jets, $t\bar{t}$ +jets, H +2 jets, ...
- **Large samples at acceptable cost**
 ~ 10 billion events per year, per experiment, through full detector simulation

No single calculation provides all four: matching and merging combine them.

Two complementary approximations

	m	$m+1$	$m+2$	$m+3$	$\dots$
tree	B	R	RR	RRR	$\dots$
1-loop	V	RV	RRV		
2-loop	VV				

Red: exact at fixed order (here NLO)

Blue: filled *approximately* by the shower

Fixed order: exact coefficients at fixed multiplicity; inclusive, signed weights; breaks down where $\alpha_s L^2 \sim 1$.

Parton shower: iterate the singular limits,

$$R \xrightarrow{\text{soft/collinear}} BK, \quad \Delta = \exp \left[- \int K \right].$$

Exclusive and unitary, at all multiplicities. Hard, wide-angle emissions are only approximate.

Matching: NLO initial condition for shower evolution

	m	$m+1$	$m+2$	$m+3$	$\dots$
tree	B	R	RR	RRR	$\dots$
1-loop	V	RV	RRV		
2-loop	VV				

Red: exact NLO $\{B, V, R\}$

With $\bar{B} = B + V + \int (R - D)$:

POWHEG (multiplicative)

$$\bar{B} \left[\Delta(t_0) + \frac{R}{B} \Delta(t) d\Phi_r \right],$$

$$\Delta(t) = \exp \left[- \int d\Phi_r \frac{R}{B} \theta(t_r - t) \right]$$

MC@NLO (subtractive)

$$\underbrace{\bar{B}_{\text{MC}} \otimes \text{PS}}_{\mathbb{S}} + \underbrace{(R - D_{\text{MC}}) \otimes \text{PS}}_{\mathbb{H}, \text{ either sign}}$$

D_{MC} : the shower's own first-emission approximation.

Reproduce the NLO expansion without double counting or spoiling the shower.

Merging: several multiplicities in one sample

	m	$m+1$	$m+2$	$m+3$	$\dots$
tree	B	R	RR	RRR	$\dots$
1-loop	V	RV	RRV		
2-loop	VV				

Red: NLO X ; violet: NLO $X + j$

A direct sum of $X+PS$ and $X+j+PS$ double counts.

A **merging scale** Q_{cut} assigns each configuration:


- above: matrix element of that multiplicity;
- below: shower, with Sudakov factors and vetoes.

LO: CKKW(-L), MLM

NLO: MEPS@NLO, FxFx, UNLOPS

Accuracy is set per multiplicity: NNLO, NLO, LO or shower.

Three goals for matching



3: usable samples ($w \geq 0$, cost/event)

1. Formal accuracy

NNLO+PS beyond colour singlets; logarithmic accuracy that survives matching

2. High multiplicity



the jet multiplicities analyses use, at the best available order

3. Usable large samples

positive weights and throughput for detector simulation

Goal 1: formal accuracy


Established for colour singlets: **GENEVA**, **MINNLO_{PS}**,
UN2LOPS.




[Two-loop 2 $\rightarrow$ 3 amplitudes since 2021; C. Signorile-Signorile, QCD@LHC 2026]

Which processes are possible depends largely on which two-loop amplitudes are known.


NNLO+PS: from colour singlets to coloured Borns



Since 2020: diboson, heavy quarks (approximate two-loop MEs where necessary).




[$gg \rightarrow HH$: $\text{MiNNLO}_{\text{PS}}$ vs GENEVA]



[$t\bar{t}H$ in $\text{MiNNLO}_{\text{PS}}$ (validation, $\mathbb{H}^{(2)} = 0$)]

- Loop-induced HH : top-mass effects included; the unknown multi-loop amplitudes are approximated.
- $t\bar{t}H$: a coloured Born. The two-loop finite remainder is unknown: approximated from soft-Higgs and high-energy limits, with its own uncertainty (set to zero in the validation plot).

Resolution variable



[GENEVA $b\bar{b} \rightarrow H$: q_T before/after Pythia8]

MiNNLO_{PS}: q_T or $\mathcal{T}_0$; $\mathcal{T}_1$ formulated for singlet + jet.

Or local counterterms

MC@NNLO requirements on the counterterms

Prerequisites for MC@NLO-type matching at NNLO without loss of logarithmic accuracy:

- 1. **Coincide with the shower splitting kernels**
 - known $1 \rightarrow 3$ splittings defined only in the collinear limit
 - ~ **exact spin-blind radiation**: scalar radiators + pure remainders
- 2. **NLL-safe momentum mapping**
 - ~ ALARIC recoil, applied **iteratively**
- 3. **Collinear limits within one eikonal disentangled**
 - ~ partial fractioning in angles, **no sectorisation**

The scheme presented here is designed to satisfy all three.


[M. Knobbe, QCD@LHC 2026]

Both seek to preserve the shower's logarithmic accuracy

NLO + NLL

ESME, PanScales

Towards NNLO + NNLL




[NNDL tests: unmatched shower fails (red); matched pass (green)]

Two coloured Born legs: e^+e^- , Drell-Yan, Higgs, DIS.


Also positive weights (Goal 3).

Matching must agree with the shower in the unresolved limits.



[NNLL tests with the matched double-real; N. Schalch]

Full $\mathcal{O}(\alpha_s^2)$ double-real, $\gamma^* \rightarrow 4$ partons, leading colour; spread over all shower histories.




[Reconstructed m_t , $e^+e^- \rightarrow W^+W^-b\bar{b}$ at 365 GeV: Alaric S-MC@NLO, default vs resonance- and width-aware]

- Radiation and recoil can shift a reconstructed resonance mass.
- **Recoil must stay inside the decaying system:** resonance-aware maps in shower, subtraction and matching.
- **Width-aware evolution:** the finite Γ_t screens soft wide-angle radiation.

Goal 1, open: NNLO+PS for coloured and loop-induced Borns; NNLL-preserving NNLO matching; decays.

Goal 2: high multiplicity




[$gg \rightarrow H$ in the VBF fiducial $H + 2j$ region]

- $gg \rightarrow H + 2j$ (the VBF background): inclusive NNLO+PS samples are only **LO** here.
- The setups used by ATLAS and CMS differ by more than 20%.
- $VH \rightarrow b\bar{b}/c\bar{c}$: V+HF backgrounds from LO ($\leq 4j$) and FxFx NLO ($\leq 3j$) samples; negative weights dilute the NLO one.
- Cost per event grows steeply with multiplicity.

We need the right perturbative order in each jet bin.

Merging: how it works, and what to specify



[$pp \rightarrow h + \text{jets}$: 0,1,2 jets (exclusive) at NLO, 3 jets at LO, and their sum]

Common structure


1. Resolve: matrix elements above Q_{cut}
2. Cluster: shower-like branching history
3. Reweight: α_s and Sudakov factors
4. Shower: veto double-counted emissions

To specify

- Per-multiplicity accuracy
- Inclusive rates: unitarised (UNLOPS) or not
- Highest multiplicity: hard tail filled by shower
- Histories and starting scales

Fixed-order accuracy and unitarisation are separate properties.


Merging introduces an additional scale dependence



[Alaric MEPS@NLO (5 jets), e^+e^- at 91.2 GeV; Durham y_{23} for $Q_{\text{cut}} = 5, 9, 16$ GeV; coloured: 2- to 5-jet sub-samples]

- Q_{cut} is an unphysical parameter: sub-samples hand over there, and the sum should not depend on it beyond the claimed accuracy.
- Like $\mu_{R,F}$ and shower scales, it must be varied to build the uncertainty band; it probes the ME/shower interface.
- Here: $Q_{\text{cut}} = 5\text{-}16$ GeV changes the total by $<10\%$.

A small Q_{cut} dependence does not prove logarithmic accuracy.




[ALEPH thrust at the Z pole: MEPS@NLO (5 jets) vs MEPS (LO, 6 jets), μ_R bands]

Alaric, e^+e^- at the Z pole

- NLO through five jets; LO through six.
- Subtraction and matching adapted to the shower's recoil.
- NLO merging halves the scale uncertainty in the bulk of thrust vs LO merging.

Open: is NLL preserved by merging (histories, vetoes)?

Combining NNLO+PS with multijet merging



UN²LOPS

Unitary NNLO for the inclusive Born, merged higher multiplicities

MiNNLO_{PS}

NNLO X , NLO $X+j$, LO $X+2j$ in one sample

GENEVA

NNLO + resummed resolution variable, NLO for the next multiplicity


Beyond one extra jet at NLO, NNLO+PS and NLO multijet merging are not yet routinely combined.

Goal 2, open: log accuracy after merging; NNLO + multijet NLO; cost at 5–6 jets.

Goal 3: usable large samples

*“The **limited size** of the simulated samples that can be produced under resource constraints is a source of major uncertainty in many analyses [...] and may get significantly worse at HL-LHC, where the projected computing needs of the experiments **exceed the resources that are expected to be available.**”*

— HSF Physics Event Generator WG, input to the LHCC review of HL-LHC computing



[F. Giuli, QCD@LHC 2026]

Prevent (matching), reduce (generation), repair (post hoc).

- ATLAS, CMS: $\mathcal{O}(10^{10})$ fully simulated events per year each, $\sim 3\times$ the data
- Negative S and H events; more with NLO merging

$$N_{\text{eff}} = N(1 - 2f_-)^2$$


Prevent I: change the factorisation scheme (KrkNLO)

- PDFs in the **Krk scheme**: no explicit subtraction.
- Shower LO events, then multiplicative weights:


$$w_R = \frac{R}{PB}, \quad w_V = 1 + \frac{\alpha_s}{2\pi} \left(\frac{V+I}{B} + \Delta_0^{\text{FS}} \right)$$

- Colour singlets: W , Z , $Z\gamma$, WW , ZZ .

A very small fraction of negative weights remains in the setups studied.



Prevent II: positive weights, shower intact



ESME: exponentiated subtraction as a counting process; integer weights $n_B \geq 0$. Positive and NLL-preserving (Goal 1); two coloured Born legs.

MAcNLOPS: drop negative $\mathbb{H}$ events, compensate in the $\mathbb{S}$ events' first emission,

$$P_{\text{acc}} = \min\left(1, \frac{R}{K}\right).$$

- K : shower approximation to R
- Negative $\mathbb{S}$ weights remain
- ZZ in MG5_aMC + Pythia8

Reduce and repair: after the matching

“The presence of negative densities and negative weights, pose an array of challenges to traditional neural likelihood ratio estimation methods.”

— Drnevich, Jiggins, Katzy, Cranmer (2024)


Signed-weight losses make training possible; weight variance and negative-density regions remain hard.

Reduce during generation

- ARCANE: move weight between hidden histories; observables unchanged.
- MC@NLO: $\mathbb{S}$ events from reweighted LO events, heavy folding.
- Neural control variates: a trainable signed subtraction.

Repair post hoc


- Cell resampling: average nearby weights; also merged samples.
- Learned positive reweighting with uncertainties.



[Cell resampling: weight distribution]

Choosing a method: benchmark first

Preliminary




[Generation + shower, one CPU core; DY and Z+jet]

- The cost that matters is CPU per *effective* event in the analysis region, not f_- alone.
- With detector simulation, N/N_{eff} dominates overall cost.
- This can be benchmarked and strategised around.

Pending a 'Goldilocks' method: where resources are constrained we should optimise smartly.

Throughput: matching on GPUs



[NLO+PS $pp \rightarrow Z$: GPU, CPU, Herwig]

GAPS: NLO-matched initial- and final-state dipole shower.

Identical code for CPU and GPU.

$\mathcal{O}(10^6)$ NLO+PS events per minute on one GPU
(weighted, parton level)

Upcoming: GPU KrkNLO (proof-of-concept).

Also: Pepper (portable LO), MG5_aMC multi-event NLO interface, vectorised folding (MadVfold).

Open: positive weights with merging and NNLO; throughput at high multiplicity.

Accuracy

- NNLO+PS for complicated processes
- ...preserving logarithmic accuracy

Multiplicity

- NLO merging with NLL showers
- NNLO inclusive + NLO merging
- more jets at lower cost

Efficiency

- Matched samples with maximal statistical power
- ...with merging
- GPU architecture to help with throughput


Medium-term: need to accept trade-offs for different applications.

1. Matching and merging combine complementary approximations for **exclusive events** at the **best available precision**.
2. **Accuracy**: NNLO+PS **systematically expanding** process coverage; NLO matching that preserves NLL is achievable; NNLL showers emerging.
3. **Multiplicity**: NLO merging with modern showers works; its logarithmic accuracy is open.
4. **Samples**: positive-weight matching methods emerging; GPU generation being prototyped (currently for simple processes); post-hoc being studied.


Thank you.

Questions?

Backup: NLO+PS matching methods, 1996–2026



Backup: NLO+PS multijet merging and resummation



Backup: off-shell top pairs and resonance histories

$e^+e^- \rightarrow t\bar{t} \rightarrow$ **semileptonic final states**

- Full off-shell NLO QCD calculation from MoCaNLO.
- Matching to the Pythia8 final-state shower with resonance histories supplied event by event.
- Dedicated MC@NLO counterterms based on Catani–Seymour subtraction.

Complementary to resonance- and width-aware evolution: compare line shapes only after aligning the decay channel, histories and recoil settings.

$$pp \longrightarrow Z/\gamma^* \longrightarrow \ell^+ \ell^- \quad + \quad \text{QCD and QED radiation}$$

One interleaved shower

- NLO QCD+EW matrix elements matched in MC@NLO.
- Resonance-aware modifications for QCD/QED-singlet resonances.
- Radiation, recoil and matching counterterms must be consistent.

Validation in lepton-pair production

- Compare to fixed-order QCD+EW.
- Compare to QCD MC@NLO supplemented with YFS resummation.
- Keep photon definitions and fiducial cuts aligned.

Backup: QED counterparts and recent YFS developments


QED shower matching

Sequential emissions + fixed-order EW corrections;
recent $e^+e^- \rightarrow ZH$ NLO EW matching.

YFS exponentiation

Exponentiate universal soft-photon factors and
improve the finite residuals.

- Automated full NLO EW.
- Process-specific complete NNLO examples.
- General two-loop providers remain a bottleneck.



$[e^+e^- \rightarrow \mu^+\mu^- \text{ at } 161 \text{ GeV}]$

Backup: Subtractive and multiplicative matching

MC@NLO: subtract the shower overlap

$$d\sigma \sim \underbrace{\bar{B}_{\text{MC}} S_B}_{\mathbb{S} \text{ events}} + \underbrace{(R - D_{\text{MC}}) S_R}_{\mathbb{H} \text{ events}}.$$

- D_{MC} reproduces the shower's first-emission approximation.
- The remainder can have either sign.

POWHEG: generate the hardest emission

$$d\sigma \sim \bar{B} \left[\Delta(t_0) + \frac{R}{\bar{B}} \Delta(t) d\Phi_r \right].$$

- The shower must respect the generated emission and its veto.
- Positivity requires a suitable $\bar{B}$ and treatment of remnants.

Common requirement: recover the NLO expansion without double counting.

Backup: MAcNLOPS: removing negative hard events

Born-like $\mathbb{S}$ events

first shower step $\longrightarrow \begin{cases} \text{keep} & \text{no emission,} \\ P_{\text{acc}} & \text{emission.} \end{cases}$

Hard $\mathbb{H}$ events

$R - K > 0 \longrightarrow \text{keep,}$

$R - K < 0 \longrightarrow \text{discard.}$


[Branching logic adapted from the matching equations]

Remove negative $\mathbb{H}$ events and compensate through the first emission of $\mathbb{S}$ events:

$$P_{\text{acc}} = \min\left(1, \frac{R}{K}\right).$$

- K : shower approximation to R .
- Negative $\mathbb{S}$ weights remain.
- Implemented for ZZ in MG5_aMC + Pythia8.

Backup: Higgs pairs: comparing complete generator setups




$gg \rightarrow HH$ at NNLO+PS

- GENEVA and MiNNLO_{PS} extend process coverage with finite- m_t information.
- The full finite-mass NNLO amplitude information is unavailable.
- The pair- p_T comparison probes different complete setups.

Mass approximations, resolution variables and shower treatment must also be controlled.

Backup: Coloured Born systems: top-pair-Higgs production



[Validation: $\mathbb{H}^{(2)} = 0$ in both NNLO calculations]

$pp \rightarrow t\bar{t}H$ in $\text{MiNNLO}_{\text{PS}}$

- A heavy coloured Born system with an additional colour singlet.
- Two-loop finite remainder approximated by combining soft-Higgs and high-energy information.
- Approximation uncertainty is distinct from scale variation.