



Precision event generation NNLO subtraction, and robust calibration of non-perturbative models

Max Knobbe

Fermi National Accelerator Laboratory

Michigan State University, HEP Seminar

September 2026


Based on:

J. Campbell, S. Höche, MK, C. Preuss, D. Reichelt arXiv:2505.10408, Phys. Rev. D 113, 054031 (2026)

J. Campbell, S. Höche, MK arXiv:2606.26247

A. Iskauskas, MK, F. Krauss, S. Schumann arXiv:2602.22324, JHEP 08 (2026) 066

A small signal on a simulated background




- $H \rightarrow b\bar{b}$ in VH : 4.9σ observed, 4.3σ expected, rate $1.16^{+0.27}_{-0.25} \times$ Standard Model
 - **Per-cent-level signal:** 128 ± 28 on 7850 ± 90 background events
 - **Background shapes from simulation**
 - SHERPA 2.2.1, POWHEG+PYTHIA 8
 - normalisations float in the fit
- $\Rightarrow$ background *and* signal template: generator predictions

ATLAS, 79.8 fb^{-1} : $m_{b\bar{b}}$ weighted by
 S/B , signal $\times 5$.

[ATLAS, Phys. Lett. B 786 (2018) 59]

Anatomy of a simulated collision



What we should be aiming for

Part A Perturbative accuracy

- Today: NLO+PS, NNLO only as a parton-level K -factor
- Aim: **NNLO+PS matching** — process-independent, fully differential

⇒ **a requirement on the subtraction scheme**

Part B Non-perturbative uncertainties

Expansion in $\alpha_s(m_Z) \approx 0.118$:

$$\sigma = \sigma^{(0)} + \alpha_s \sigma^{(1)} + \alpha_s^2 \sigma^{(2)} + \dots$$

LO	NLO	NNLO
$\mathcal{O}(1)$	$\approx 10\%$	$\approx 1\%$

Rough estimates; process and observable dependent

NNLO calculations exist for a growing list of processes.


What is missing is a way to put them into the events.

What we should be aiming for

Part A Perturbative accuracy

- Today: a **two-point difference**, string against cluster
- The *same* cluster model, retuned to LEP, moves the response by $\approx 1\%$
 - calibration matters as much as model choice
- Aim: the **parameter region the data allow**, propagated like PDFs

Part B Non-perturbative uncertainties



Jet energy response relative to PYTHIA 8, $|\eta| < 0.7$.

[ATLAS, ATL-PHYS-PUB-2022-021]

A.I


Splitting Functions

Infrared subtraction, and the scalar-radiator decomposition


Campbell, Höche, MK, Preuss, Reichelt, arXiv:2505.10408, Phys. Rev. D 113, 054031 (2026)

Infrared divergences at higher orders

- Two classes at each order
 - **virtual**: extra loops, same final state
 - **real**: extra radiated partons
 - Each **separately divergent** in $d = 4$; radiation **soft** or **collinear**
 - **Cancel** in the sum, for infrared-safe observables
[Bloch, Nordsieck (1937); Kinoshita, Lee, Nauenberg (1962/64)]
 - Different phase spaces: $n+1$ against n partons
- ↪ **no common variable, no bin-by-bin combination**



real: $n+1$ partons, tree



virtual: n partons, one loop

$$\sigma^{\text{NLO}} = \underbrace{\int^{n+1} d\sigma_{n+1}^{(0)}}_{\text{div.}} + \underbrace{\int^n d\sigma_n^{(1)}}_{\text{div.}}$$

Subtraction, and the two universal singular limits

- Add and subtract $\mathcal{S}$, reproducing the singular behaviour:

$$\sigma^{\text{NLO}} = \underbrace{\int d\phi_{n+1} (|\mathcal{M}_{n+1}|^2 - \mathcal{S})}_{\text{finite in } d=4} + \int d\phi_n \left[\underbrace{|\mathcal{M}_n^{(1)}|^2}_{\text{poles in } 1/\epsilon} + \underbrace{\int d\phi_1 \mathcal{S}}_{\text{poles in } 1/\epsilon} \right]$$

- **both terms separately finite**, the result fully differential

- $\mathcal{S}$ from the two universal limits of one emission:

$$s_{ij}(q) = \frac{2s_{ij}}{s_{iq}s_{jq}} \quad |\mathcal{M}_{n+1}|^2 \xrightarrow{p_1 \parallel p_2} \frac{8\pi\mu^{2\epsilon}\alpha_s}{s_{12}} P_{a_1 a_2}^{ss'} \otimes |\mathcal{M}_n|^2$$

- **soft**: squared eikonal current, pairs (i, j)
- **collinear**: Altarelli-Parisi kernel of the splitting leg
- the two **overlap**


The soft limit resolves the **colour but not the spin**, the collinear limit the **spin but not the colour**.

Matching at NLO: the counterterm is the shower kernel

- **Fixed order:** exact in the **hard** region, fixes the **normalisation**
- **Parton shower:** resums **soft and collinear**, fixes the **shapes**, gives a full final state
- Same first emission in both $\leadsto$ matching **subtracts what the shower already generates:**

$$d\sigma = \underbrace{\left[B + V + \int d\Phi_1 S \right]}_{\text{shower off the Born}} + \underbrace{\left[R - S \right]}_{\text{hard remainder}}$$

- S : regulator of R and shower approximation to it



The overlap is the first shower branching, expanded to $\mathcal{O}(\alpha_s)$.


MC@NLO [[Frixione, Webber \(2002\)](#)], POWHEG

[[Nason \(2004\)](#)].

At NLO the subtraction counterterm and the shower kernel are the **same object**:
one emission, one $1 \rightarrow 2$ branching.

At second order: what a fixed-order scheme does not supply

- **Not** the singular functions: $1 \rightarrow 3$ splittings and double-soft currents known for decades
[Campbell, Glover (1997); Catani, Grazzini (1999)]
- NNLO+PS exists **only for colour singlets**, with no generalisation
- Soft and collinear **overlap**: **no unique emitter**
 - schemes differ here; several **sectorise** phase space
- Built to deliver a fixed-order *number*, not to feed a shower



The **sector boundaries** are internal to the subtraction and sit *away* from the singular limits. A **shower** over the whole phase space crosses them.

A sectorised counterterm can only drive a shower with the *same* partition.
Missing: counterterms that are globally valid and one-to-one with a shower branching.

Requirements on the counterterms at NNLO

MC@NLO-style matching at NNLO, without degrading the logarithmic accuracy of the shower, imposes three conditions on the counterterms [Dasgupta, et al. (2018); Herren, et al. (2022)]:

❶ Coincide with the shower splitting kernels

- known $1 \rightarrow 3$ splittings defined *in* the collinear limit
- ↪ **exact computation of the spin-blind radiation** (scalar radiators)

❷ NLL-safe momentum mapping

- ↪ the ALARIC recoil, applied **iteratively** (Sec. A.II)

❸ Collinear limits within one eikonal disentangled

- ↪ **partial fractioning** in angles, **without sectorisation** (Sec. A.II)

The scheme presented here is built to satisfy all three.

Motivation for scalar QCD radiators

- Two distinct objects at second order, each defined only in its own limit
 - **double-soft**: $2 \rightarrow 4$ off a colour **dipole**
 - **triple-collinear**: $1 \rightarrow 3$ off a single leg
- They overlap, and the overlap is the spin-blind radiation
- **Scalar QCD computes exactly that**: radiation off a spinless charge
 - a full amplitude, no kinematic limit taken


$$P_{1 \rightarrow 3} = \underbrace{\text{scalar radiator}}_{\text{no kinematic limits}} + \underbrace{\text{pure splitting remainder}}_{\text{spin-dependent}}$$

The double-unresolved counterterm has **two** pieces:

$$S_2 = \underbrace{\sum_{(i,k)} \hat{\mathbf{T}}_i \hat{\mathbf{T}}_k \mathcal{S}_{i;k}^{(2 \rightarrow 4)}(q_1, q_2)}_{\text{scalar radiators, one per dipole}} + \underbrace{\sum_{\text{splittings}} \langle P_{1 \rightarrow 3}^{(p)} \rangle}_{\text{pure splitting remainders}}$$

Only the radiator is new; the sum is the known $1 \rightarrow 3$ function.

Scalar QCD from the Gordon decomposition



- **Exact rewriting** of the quark-gluon vertex [Gordon (1928)]

$$\frac{\not{p} + \not{q}}{(p+q)^2} T_{ij}^a \gamma^\mu = T_{ij}^a \left[S^\mu(p, q) + \frac{i\sigma^{\nu\mu} q_\nu}{(p+q)^2} - \frac{\gamma^\mu \not{p}}{(p+q)^2} \right], \quad S^\mu = \frac{(2p+q)^\mu}{(p+q)^2}$$


$\rightsquigarrow$ **scalar current**: a **spinless charge**; **magnetic-moment term**: the spin; $\gamma^\mu \not{p}$: vanishes by the equations of motion

- Same decomposition for the **triple-gluon vertex**

$$\frac{d_\sigma^\mu(p_{12})}{p_{12}^2} f^{abc} \Gamma^{\sigma\nu\rho}(p_1, p_2) = i d_\sigma^\mu(p_{12}) \text{Tr}(T^c[T^b, T^a]) \left[S^\nu(p_1, p_2) g^{\rho\sigma} - \frac{(p_1 - p_2)^\rho}{2p_{12}^2} g^{\sigma\nu} \right] + \text{swap}$$

- Manifest in a physical **axial gauge**; the subtraction must use that *same* gauge
[Ellis, Georgi, Machacek, Politzer, Ross (1979)]

NLO splitting functions from scalar QCD





- Single-gluon emission, $\mathcal{S}_g = \sum_{i,k} \hat{\mathbf{T}}_i \hat{\mathbf{T}}_k \mathcal{S}_{i;k}$:

$$\mathcal{S}_{i;k}(q_1; \bar{n}) = \frac{2z_i}{z_1 p_{i1}^2} + \frac{2z_k}{z_1 p_{k1}^2} - \frac{4p_i p_k}{p_{i1}^2 p_{k1}^2} \xrightarrow[\text{Bose symmetry}]{\text{on shell}} -\frac{4p_i p_k}{p_{i1}^2 p_{k1}^2}$$

- Remainder $\langle P_{q \rightarrow q}^{(f)} \rangle / s_{12} = C_F(1 - \epsilon) z_2 / (z_{12} s_{12})$: finite for $z_2 \rightarrow 0$
- One emission: scalar and eikonal currents give **identical matrix elements**
 - they differ only at two emissions, where $s_{i12} \neq s_{i1} + s_{i2}$

Scalar radiators: $q\bar{q}$ emission



- Soft starting point: eikonal currents
 $J^\mu(q) = \sum_i \hat{\mathbf{T}}_i p_i^\mu / (p_i q)$ on the fermionic vacuum polarisation
- Replace throughout by the **scalar** current,
 $p_i^\mu / (p_i \cdot q) \rightarrow S^\mu(p_i, q)$

$$\mathcal{I}_{i;k} = 4T_R \frac{N_{ik}}{s_{12}^2 (s_{i1} + s_{i2})(s_{k1} + s_{k2})} \quad \text{double-soft limit}$$

$$\mathcal{S}_{i;k}^{(q\bar{q})} = 4T_R \frac{N_{ik}}{s_{12}^2 s_{i12} s_{k12}} \quad \text{scalar radiator}$$

$$N_{ik} = s_{i1}s_{k2} + s_{i2}s_{k1} - s_{ik}s_{12}$$

$\Rightarrow$ identical numerators; only the **physical propagator** $s_{i12} = (p_i + q_1 + q_2)^2$ replaces its soft-limit approximation [Campbell, Glover (1997); Catani, Grazzini (1999)]

- valid over the **entire** phase space; covariant gauge here, axial gauge adds the z -dependent terms of the previous slide

One object: the double-soft current *and* the spin-averaged collinear splitting.

Triple-collinear decomposition: distinct and same flavour

- **Distinct flavour** $q \rightarrow \bar{q}' q' q$: one diagram, so it factorises exactly,

$$\langle P_{q \rightarrow \bar{q}' q' q} \rangle = \frac{s_{123}}{s_{12}} \left[\frac{C_F}{2} s_{123} \mathcal{S}_{3;3}^{\mu\nu}(p_{12}) + \langle P_{q \rightarrow q}^{(f)\mu\nu}(p_3, p_{12}) \rangle \right] P_{g \rightarrow q, \mu\nu}(p_1, p_2)$$

- **radiator** + **magnetic remainder**, the latter soft-finite
 $\rightsquigarrow$ **no new $1 \rightarrow 3$ structure**


- **Same flavour** $q \rightarrow \bar{q} q q$: two diagrams, whose interference generates a genuine splitting function,

$$\begin{aligned} \langle P_{q \rightarrow \bar{q} q q} \rangle &= \frac{s_{123}}{s_{12}} \langle P_{q \rightarrow q}^{\mu\nu}(p_3, p_{12}) \rangle P_{g \rightarrow q, \mu\nu}(p_1, p_2) \\ &\quad + \frac{s_{123}}{s_{13}} \langle P_{q \rightarrow q}^{\mu\nu}(p_2, p_{13}) \rangle P_{g \rightarrow q, \mu\nu}(p_1, p_3) + \langle P_{q \rightarrow \bar{q} q q}^{(p)} \rangle \end{aligned}$$

- $\langle P^{(p)} \rangle$: **not** reconstructible from lower orders

Azimuthal dependence retained **if and only if** no kinematic limit is taken.

Scalar radiators: two-gluon emission



Abelian *and* non-abelian components: **complete** and validated

- **Leading double-soft**, after colour-charge conservation:

$$\mathcal{S}_{i,k;l,m}^{(ab)} = \frac{1}{4} \mathcal{S}_{i;l}(q_1) \mathcal{S}_{k;m}(q_2) = \frac{s_{il} s_{km}}{s_{i1} s_{l1} s_{k2} s_{m2}}$$

- the **squared one-gluon current**: the structure a shower iterates
- two further abelian terms, sub-leading and sub-sub-leading
- **Non-abelian**: double-soft limit reproduces Catani-Grazzini [Catani, Grazzini (1999)]
 - the difference is sub-leading, but **required for the ϵ -poles**
 - the full radiator does **not** factorise

All $1 \rightarrow 3$ splittings decomposed, quark *and* gluon initial states.
 They agree with the known results; the decomposition is new.

Scope of the decomposition

- **All $1 \rightarrow 3$ splittings** = **scalar radiators** + **pure remainders**
 - $q \rightarrow qgg, q \rightarrow qq'\bar{q}', g \rightarrow ggg, g \rightarrow gq\bar{q}$
- **Scalar part**: a **full amplitude**, physical propagators, no limit taken $\rightsquigarrow$ valid over the **entire** phase space
- **Remainders** sub-leading, with **two exceptions**
 - $P_{\tilde{q} \rightarrow gg\tilde{q}}^{(\text{ab},p)}, P_{\tilde{q} \rightarrow gg\tilde{q}}^{(\text{pnab},p)}$ keep double-soft poles; reclassified as radiators
- Same decomposition for the **one-loop $1 \rightarrow 2$ splittings**, needed for the real-virtual

The ingredients are complete.

What remains is to turn them into a subtraction scheme.

A.II

Constructing the Scheme

Overlap removal, kinematics, integration, and numerical validation

Campbell, Höche, MK, [arXiv:2606.26247](https://arxiv.org/abs/2606.26247)

Structure of the double-real subtraction

$$\mathcal{M}_{RR} = \mathcal{M}_{n+2} - \sum \mathcal{M}_n \otimes \mathcal{S}_2 - \sum \left[\mathcal{M}_{n+1} - \mathcal{M}_n \otimes \mathcal{S}_1 \right] \otimes \mathcal{S}_1$$

$\mathcal{M}_k$: **squared** tree-level matrix element, k partons. $\otimes$: colour and spin correlations.


- $\mathcal{S}_2$: **double-unresolved** divergences; $\mathcal{S}_1$: **single-unresolved**
- Nested bracket: the **overlap** of iterated NLO with pure NLO [Gehrmann-De Ridder, et al. (2004)]
 $\rightsquigarrow \mathcal{S}_2$ **need only supply what iterated NLO does not**
- **Intermediate** polarisation tensor: consistent between the two clusterings

Open question: **to which emitter does each counterterm belong?**

Answered here by partial fractioning, not by sectorisation.

Overlap removal at NLO

$\mathcal{S}_{i;k}(q_1) = -4p_i p_k / (p_{i1}^2 p_{k1}^2)$ is singular for $q_1 \parallel p_i$ **and** $q_1 \parallel p_k$: two collinear directions in one object, no unique emitter.




$$\mathcal{S}_{i;k}(q_1) = \mathcal{S}_{i;k}^{(i)}(q_1; n) + \mathcal{S}_{i;k}^{(k)}(q_1; n), \quad \mathcal{S}_{i;k}^{(i)}(q_1; n) = \frac{\eta_{k1}}{\eta_{k1} + \eta_{i1}} \mathcal{S}_{i;k}(q_1)$$

$$\eta_{ij} = \frac{2p_i p_j}{\zeta_i \zeta_j n^2}, \quad \zeta_i = \frac{2p_i n}{n^2}$$

- An **angular** weight $\rightsquigarrow$ a unique emitter, a well-defined branching
- A **smooth reweighting of the integrand, not a partition of phase space**

$\Rightarrow$ the construction the ALARIC shower uses for its kernels [Herren, et al. (2022)]

Overlap removal at NNLO



Two emissions off one dipole: **four** emitter assignments


- Leading double-soft radiator **fully factorised**: each factor partial fractioned **exactly as at NLO**

$$\mathcal{S}_{i,k;l,m}^{(\text{ab}), (i,k)} = \frac{1}{4} \underbrace{\frac{\eta_{l1}}{\eta_{l1} + \eta_{i1}} \mathcal{S}_{i;l}(q_1)}_{\mathcal{S}_{i;l}^{(i)}(q_1;n)} \underbrace{\frac{\eta_{m2}}{\eta_{m2} + \eta_{k2}} \mathcal{S}_{k;m}(q_2)}_{\mathcal{S}_{k;m}^{(k)}(q_2;n)}, \quad n = K + q_{12}$$

- the product of the two weights selects one of the four assignments
- $[\mathcal{M}_{n+1} - \mathcal{M}_n \otimes S_1] \otimes S_1$ needs a **hierarchy** of the two emissions [Dulat, et al. (2018)]
 - a **smooth weight** in transverse momentum, not a Θ -function
 - no region excluded from ordered evolution

$\Rightarrow$ a unique branching topology for every counterterm, without sectorisation:
 each counterterm $\leftrightarrow$ one parton-shower branching

Kinematics: iterated NLO mappings



The emitter keeps its direction $\hat{n}$ and is only **rescaled**; the spectator system absorbs the recoil.

- The identified-particle mapping [Catani, Seymour (1997)] applied **twice**, on shell at every step
- **For matching:** the **same iterated $1 \rightarrow 2$ kinematics** as the ALARIC shower $\rightsquigarrow$ the recoil matches an NLL-safe shower exactly [Herren, et al. (2022)]
- Other schemes: a $3 \rightarrow 2$ antenna or a direct $4 \rightarrow 2$ map
 - fine for a fixed-order result, but **neither splits into two successive $1 \rightarrow 2$ branchings**
- **A natural path towards NNLO+PS** [Campbell, et al. (2023)]

Integration of the counterterms

$\int d\Phi_{+2} S_2$ exposes the $1/\epsilon$ poles that cancel against the real-virtual and double-virtual contributions. For the **scalar radiators** in **back-to-back kinematics**, two independent routes:

1. Analytic

- Reverse unitarity
[Gehrmann-De Ridder, et al. (2003)]
- IBP reduction with KIRA
- Closed forms for all radiators, e.g.

$$\mathcal{I}_{i,i;k,k}^{(\text{ab})} = \frac{1}{4\epsilon^4} + \frac{1}{\epsilon^3} + \frac{1}{\epsilon^2} \left(\frac{9}{2} - \frac{3\pi^2}{8} \right) + \frac{1}{\epsilon} \left(19 - \frac{37}{6}\zeta_3 - \frac{3\pi^2}{2} \right) + \dots$$

2. Numerical



- Laurent expansion in the iterated two-body phase-space variables
- Sector decomposition [Czakon (2010, 2011)], with PYSECDEC
- Only the *sector variables* adopted: the Born phase space stays factorised
- **Agreement order by order in ϵ**

Partial-fractioned terms integrate to the **expected fractions**.

Outstanding: the pure remainders, and general kinematics.

Pole structure after partial fractioning

- Partial fractioning relocates singularities, it does not remove them
- Each term keeps the poles of a single branching topology
- Sectors labelled by **collinear topology**
 - an integration device; the scheme is not sectorised
- Abelian two-gluon radiator at ϵ^{-4} :
 (i, i) **exactly zero** in S^{II} , (i, k) in S^I



coeff. of ϵ^{-4}	S_1^I	S_2^I	S_3^I	S_1^{II}	S_2^{II}	ratio
$\mathcal{I}_{i,i;k,k}^{(ab)}$	$\frac{1}{16}$	$\frac{1}{48}$	$\frac{1}{24}$	$\frac{1}{24}$	$\frac{1}{12}$	1
$\mathcal{I}_{i,i;k,k}^{(ab),(i,i)}$	$\frac{1}{32}$	$\frac{1}{96}$	$\frac{1}{48}$	0	0	$\frac{1}{4}$
$\mathcal{I}_{i,i;k,k}^{(ab),(i,k)}$	0	0	0	$\frac{1}{48}$	$\frac{1}{24}$	$\frac{1}{4}$


S_4^I, S_5^I vanish at this order; the ratio is taken with respect to the *unpartitioned* integral.

The one-to-one correspondence with a shower branching holds
 also for the *integrated* counterterm.

Validation procedure



$$\Delta(\lambda) = \frac{|\mathcal{M}_{RR}|}{\mathcal{M}_{n+2}}, \quad \mathcal{M}_{RR} = \mathcal{M}_{n+2} - \sum \mathcal{M}_n \otimes S_2 - \sum [\mathcal{M}_{n+1} - \mathcal{M}_n \otimes S_1] \otimes S_1$$

- A random phase-space point scaled into a **singular limit** by λ
[Asteriadis (2020); Bothmann, et al. (2024)]
- **Local** subtraction $\rightsquigarrow \Delta \rightarrow 0$ proportional to λ
 - **the slope of $\log |\Delta|$ is the diagnostic**
 - -1 : leading power; -2 : sub-leading too
- **Arbitrary precision**, down to $\lambda \sim 10^{-20}$
- In SHERPA and PEPPER; numbers from PEPPER
[Bothmann, et al. (2024)]



Schematic. The steeper the fall, the larger the fraction of the singular behaviour removed.

Validation: convergence in the double-soft limit



- Abelian $e^+e^- \rightarrow q\bar{q}gg$, scaled into the **double-soft** region
- **Left:** $|\Delta|$ for 250 points, each vanishing for $\lambda \rightarrow 0$
- **Right:** local slope over 10^4 points, **all stabilising at -2** $\rightsquigarrow$ leading *and* sub-leading subtracted; other channels reach -1



Convergence in all limits and channels

Slope of $\log |\Delta|$ against $\log(1/\lambda)$ at $-\log_{10} \lambda = 20$, over 10^4 points:

Limit	$q\bar{q}Q\bar{Q}$	$q\bar{q}q\bar{q}$	$q\bar{q}gg$ (ab)	$q\bar{q}gg$ (nab)
Single Soft	–	–	-1 ± 10^{-17}	-1 ± 10^{-18}
Double Collinear	-1 ± 10^{-17}	-1 ± 10^{-18}	-1 ± 10^{-17}	-1 ± 10^{-17}
Double Soft	-1 ± 10^{-18}	-1 ± 10^{-19}	-2 ± 10^{-18}	-1 ± 10^{-17}
Triple Collinear	-1 ± 10^{-19}	-1 ± 10^{-19}	-1 ± 10^{-19}	-1 ± 10^{-19}
Soft $\times$ Collinear	–	–	-2 ± 10^{-18}	–
Double Collinear $\times 2$	–	–	-2 ± 10^{-19}	–

- **Slope -1 throughout:** leading-power subtraction correct in every limit
- Uncertainties are the **spread over points**, not statistical noise: the convergence is uniform
- **Performance:** $5.1 \cdot 10^8$ points $\rightarrow 1\sigma$ of **654 fb**, below $2 \cdot 10^{-5}$ of the Born

Differential results: Durham jet rates



- Durham 2 $\rightarrow$ 3, 3 $\rightarrow$ 4 rates at $\sqrt{s} = 91.2$ GeV, against NLO Catani-Seymour for $e^+e^- \rightarrow 3$ jets
- **Caveat: double-real contribution only**
 - real-virtual and double-virtual still missing $\rightsquigarrow$ **not yet physical**
- Shown as evidence of a stable subtraction and a converged integral

B


Robust Calibration

History Matching applied to the non-perturbative models

Iskauskas, MK, Krauss, Schumann, arXiv:2602.22324, JHEP 08 (2026) 066

Non-perturbative models and their calibration

- **20** parameters for hadronisation alone; **40 to 50** for a full tune
- Standard output: one **best-fit point** plus a local ellipsoid
- Redone whenever the perturbative input improves
- Very different parameter sets describe the same data



Schematic objective-function landscape: two acceptable regions, one reported.

Locating a global minimum is not by itself a calibration.

The Hessian at one minimum is not the parametric uncertainty; the object of interest is the whole compatible region.


The test case: hadronisation models in SHERPA

Two physically distinct models in the *same* generator, on an identical hard process and shower.

AHADIC: cluster fragmentation

[Chahal, Krauss (2022)]


- **Colour preconfinement:** gluons split, singlet clusters form and decay
- **19** parameters: PT_MAX, ALPHA_x, BARYON_FRACTION, ...



PYTHIA: Lund string fragmentation

[Andersson, et al. (1983)]

- The **linear confining potential**, through an interface
- **23** parameters: aLund, bLund, Sigma, ...



Hadronisation is only the **example**.

The method is generic: any set of model parameters, in any generator.

Limitations of the established tuning methodology

Established procedure

[PROFESSOR: Buckley, et al. (2009)]

- $\mathcal{O}(1000)$ runs, analysed with RIVET
- **Polynomial surrogate** per data bin
- χ^2 minimisation, ranges narrowed by hand
- Uncertainty from $\Delta\chi_{\text{red}}^2 \leq 1$

Where this becomes insufficient

- **Multiple minima and disjoint regions not resolved**
- Generic at $\mathcal{O}(10\text{--}100)$ parameters: the optimiser runs to the **boundary of the model manifold**

[Transtrum, Machta, Sethna (2010); Gutenkunst, et al. (2007)]

- in tuning: best fits at the edge of the sampled range [Buckley, et al. (2009)]

A χ^2 minimisation is insufficient where the landscape is poorly understood.

History Matching characterises the entire acceptable region.

Model discrepancy and the implausibility measure

Measurement z , simulator f at a parameter point x :

$$z = f(x) + \epsilon(x) + e ,$$

- $f(x)$: **one generator run**; e : **observational error**
- $\epsilon(x)$: the **model discrepancy** against reality; here 1 to 3% of the observation plus generator statistics

The **implausibility** measures the distance to the data in units of all of it:

$$I(x)^2 = \frac{(\mathbb{E}_D[g(x)] - z)^2}{\text{Var}_D[g(x)] + \text{Var}[e] + \text{Var}[\epsilon(x)]}$$

- $g(x)$: a **fast emulator** replacing $f(x)$, trained on runs D ; $\mathbb{E}_D$, Var_D its prediction and own uncertainty (**next slide**)
- Large $I(x)$: **implausible**, discarded. Small $I(x)$: only **not yet ruled out**. Threshold $I = 3$ [Pukelsheim (1994)]

This is **not** a χ^2 : it establishes **badness-of-fit** only.


History Matching excludes the unacceptable regions. [Craig, et al. (1997)]

Bayes Linear emulation

- $f(x)$ is needed everywhere $\rightsquigarrow$ a **Bayes Linear emulator** [Goldstein, Wooff (2007)], one per output bin:

$$g(x) = \sum_{i=1}^p \beta_i h_i(x_A) + u(x_A) + w(x)$$

- h_i : polynomials to quadratic order, as in PROFESSOR
- u : stationary process for the residual surface
- x_A : the **active** parameters of that output
- w : the residual from the inactive ones
- Expectations and variances only, no distribution
- **An evaluation costs $\sim 10^{-7}$ s**, and returns $\text{Var}_D[g(x)]$, which enters $I(x)$



Emulator prediction $\mathbb{E}_D[g(x)]$ (left) and its uncertainty $\text{Var}_D[g(x)]$ (right) for one DELPHI output at wave 1, over two AHADIC parameters. The uncertainty increases towards the edges of the domain.

Calibration setup and the wave procedure

Input

- LEP1 through RIVET: 29 observables, 432 bins
- 19 parameters (AHADIC), 23 (PYTHIA)

Iterative procedure (waves)

- ① space-filling design, 800 points
- ② generator evaluated at those points
- ③ emulators trained on the informative outputs
- ④ exclude $I > 3$, repeat


⇒ **stop once the emulator uncertainty is subdominant**

Outcome: 3 waves for AHADIC, 5 for PYTHIA

after wave	AHADIC	PYTHIA
initial box	1.0	1.0
1	$3.0 \cdot 10^{-3}$	$2.2 \cdot 10^{-2}$
2	$7.7 \cdot 10^{-4}$	$3.4 \cdot 10^{-3}$
3	$5.2 \cdot 10^{-4}$	$1.1 \cdot 10^{-5}$
4	—	$2.2 \cdot 10^{-6}$
5	—	$1.5 \cdot 10^{-13}$

- Volume relative to the initial domain
- A further test wave: only 23% / 13% more excluded
- >390 of 432 observations matched


The non-implausible parameter space



Two-dimensional projections of the final AHADIC wave. Densities from 10^4 sampled points; each contour is a 10% increment of probability.

- Flavour parameters of gluon splitting and cluster decay
 - The surviving region: far smaller than the initial domain, strongly correlated
 - **Bimodal: two maxima at opposite ends of a curved region**
 - both describe the data equally well
 - a best-fit tune settles in *one*; its ellipsoid misses the other
 - Similar structures for many parameter pairs, in both models
- ⇒ inaccessible to single-point calibration

Goodness of fit of the non-implausible parameter sets



- Implausibility gives badness of fit only $\rightsquigarrow$ a χ^2 a posteriori, for comparison with established methods
- Both distributions narrow, peaking near $\chi^2_{red} \approx 1.5$
 - no member above 2.5 (AHADIC) or 3.0 (PYTHIA)
 - **the two models describe the data equally well**
- All 800 members are tunes of comparable quality

A set of equally acceptable tunes, not one tune with an ellipsoid.

Reduced χ^2 over all 432 calibration bins, for each of the 800 members of the final wave.

Summary and outlook

A. NNLO subtraction

- All $1 \rightarrow 3$ splittings = **scalar radiators** + spin-dependent remainders
[Campbell, Höche, MK, Preuss, Reichelt (2025)]
- **Partial fractioning replaces sectorisation**: counterterms = ALARIC shower kernels [Campbell, Höche, MK (2026)]
- Correct scaling in all limits; integrated back-to-back

Outlook:

- remainders, general kinematics, real-virtual
- e^+e^- at NNLO, then hadron colliders

B. Robust calibration

- History Matching with Bayes Linear emulators, on hadronisation
[Iskauskas, MK, Krauss, Schumann (2026)]
- **A region, not a best-fit point**: **multi-modal**, disjoint, beyond a single fit
- Every non-implausible point fits within the stated model discrepancy

Outlook:

- the region propagated as an uncertainty, like PDF variations
- a complete, uncertainty-quantified SHERPA-3.1 tune

Both parts ask: **what is the accuracy of the generator, and how is it quantified?**

Thank you!