

**Search for QCD instantons
in proton–proton collisions at ATLAS
using Inner and ALFA detectors**
action plan for diffractive variant of ANA-STDM-2021-01

Radek Vavříčka, supervisor T. Sýkora

Institute of Particle and Nuclear Physics
Faculty of Mathematics and Physics, Charles University

11.09.2026

Building on ANA-STDM-2021-01-INT1 v0.9, *Search for QCD Instantons*
Barr, Carli, Harris, Schott, Sykora, Vavricka

- 1 Where we stand
- 2 QCD instantons: theory and signature
- 3 The ALFA sub-detector
- 4 LHC runs, datasets and run conditions
- 5 Analysis strategy
- 6 Uncertainties and statistical analysis
- 7 Plan and risks

Strategy of ANA-STDM-2021-01

- 1.616 nb^{-1} of very low-pileup pp data at $\sqrt{s} = 13 \text{ TeV}$, June 2015 Run-2 commissioning; minimum-bias trigger HLT_noalg_mb_L1MBTS_1
- Target: gluon-initiated instantons, $20 < m_I < 150 \text{ GeV}$; signal SHERPA 3.0.1, CT18NNLO, seven mass slices
- Background: EPOS LHC, PYTHIA 8 JZ0–JZ2, SHERPA 3 MB, reweighted to data in N_{trk} and S_T
- Discriminants from tracks only: $\langle \eta \rangle$, $\langle |\Delta\eta(i, j)| \rangle$ dominate
- ABCDisCo: two decorrelated DNN classifiers $\rightarrow$ binned ABCD background estimate; profile likelihood fit in TREXFITTER v1.7.0

The original analysis: result and known soft spots

Result

- $\hat{\mu} = 2.33 \pm 0.64$; observed 3.5σ , expected 1.54σ
- EPOS LHC-reweighted variant:
 $\hat{\mu} = 1.88 \pm 0.83$, 3.0σ
- Excess low in N_{trk} , broad in $\langle\eta\rangle$, isotropic in thrust and sphericity; not described by any background MC considered
- Observed limit $\mu_{95}^{\text{obs}} = 42.8$ vs expected $15.4^{+9.6}_{-7.4}$

Known soft spots

- MPI in minimum-bias events *mimics* rest-frame isotropy
- Background models tuned on data that may already contain signal
- Event-shape mismodelling not fully absorbed by the signal

Cross-checking the 3.5σ against the literature


Analysis	Significance	Comment
ANA-STDM-2021-01 (this)	3.5σ obs. / 1.54σ exp.	Local. Observed/expected ≈ 2.3 : the classic signature of residual soft-QCD background mismodelling.
H1, DIS, EPJC 25 (2002) 495	no claim	Instanton-like excess seen, <i>explicitly not claimed significant</i> (simulation uncertainty); UL 60–1000 pb.
ZEUS, EPJC 34 (2004) 255	none	No instanton-induced signal.
H1 high- Q^2 , EPJC 76 (2016) 381	none	351 pb^{-1} ; UL 1.5–6 pb at 95% CL, an order of magnitude better, “ <i>for the first time challenging predictions</i> ”.
ATLAS SUEP search (13 TeV)	1.7σ local	Nearest kinematic neighbour (isotropic, high multiplicity); anchor for where such excesses land.

What the 3.5σ tells us, and what it does not

What the number tells us

- 3.5σ is a **local** significance; no look-elsewhere correction applied. It is a **shape** result: normalisation carries no information, by construction. No ATLAS instanton result is public, so no external corroboration
- Mangano, arXiv:2101.02719: predicted cross-section uncertain by up to **two orders of magnitude**, so $\hat{\mu} = 2.33$ is neither in tension with theory nor constrained by it
- H1 constrains γ^*g with a hard scale Q'^2 regulating the instanton-size integral; the LHC gg process has *no* such regulator $\Rightarrow$ **HERA gives context, not a direct constraint on the LHC rate**

The honest reading: an intriguing excess whose credibility is limited by background modelling, not by statistics. The value of a second, *orthogonal* measurement is therefore high: this is the case for ALFA.



Effective inst. vertex
(not a particle)

S. Amoroso, D. Kar, M. Schott;
EPJC 81 (2021) 624

The process

$$g + g \rightarrow N_g \times g + \sum_{f=1}^{N_f} (q_{Rf} + \bar{q}_{Lf})$$

- Semi-classical Yang–Mills solutions: tunnelling between topologically distinct vacua
- Axial anomaly: vertex **violates chirality** by $\Delta Q_5 = 2N_f$; one $q\bar{q}$ pair per flavour $\Rightarrow$ flavour democracy

Amplitude

$$\mathcal{A} = \int d^4x \int_0^\infty d\rho D(\rho) e^{-S_I} \prod_i^{N_g+2} A_\rho(p_i) \prod_j^{2N_f} \psi_\rho^{(0)}(p_j)$$

External legs **factorise** $\Rightarrow$ rest-frame isotropy.

The experimental problem, and why forward protons help

Why the inclusive search is hard

- Non-resonant: a *spectrum* of masses, no peak to fit
- At low m_I the final state resembles a minimum-bias event
- A MB event is an assembly of independent low-mass $2 \rightarrow 2$ MPI scatters, which **also** look isotropic
- Backgrounds are phenomenological, tuned to data that may contain the signal

What a leading-proton tag buys (KKMR, PRD 104 (2021) 054013)

- Large Rapidity Gap via a leading proton in ALFA. Theory target $m_I \approx 20\text{--}60$ GeV with $\sum_i E_{T,i} > 15$ GeV, $N_{\text{ch}} > 20$ in $0 < \eta < 2$: $S/B \gtrsim 8$
- A gap survives only if **no** additional soft interaction fills it: gap survival $S^2 \simeq 0.05\text{--}0.1$, a property of the *event*, not a per-scatter weight
- $\Rightarrow$ MPI suppressed **kinematically at the source**, not statistically by a classifier

Two distinct forward-tagging regimes

	ALFA (this proposal)	AFP / CT-PPS
Reference	Khoze $\times$ 2, Milne, Ryskin, PRD 104 (2021) 054013	Tasevsky, Khoze, Milne, Ryskin, EPJC 83 (2023) 35
ξ acceptance	$-4.0 < \log_{10} \xi < -1.6$, i.e. $10^{-4} \lesssim \xi \lesssim 0.025$	$0.02 < \xi < 0.15$
$ t $ acceptance	$0.016 < t < 0.43 \text{ GeV}^2$	—
Theory target m_l	20–60 GeV (PRD 104 (2021) 054013)	$> 60 \text{ GeV}$
Optics	high- β^* special runs	standard β^* , high luminosity
Pileup	$\langle \mu \rangle \sim 0.002\text{--}0.2$	$\langle \mu \rangle \sim 1\text{--}50$
Limiting factor	integrated luminosity	combinatorial pileup background

Why ALFA, not AFP

The dominant background of the original analysis is attacked by a completely different mechanism: an orthogonal cross-check of the 3.5σ . **The tag does not select m_l :** 20–60 GeV is simply where $\hat{\sigma}(m_l)$ is largest.

Why ALFA

- ALFA's low- ξ reach covers the excess mass window: **the tag makes the sample cleaner, it does not select m_l**
- EPJC 83 (2023) 35: AFP degrades badly with pileup, S/B from 2.3 (generator level) to 0.3 at $\langle\mu\rangle = 5$
- High- β^* runs give $\langle\mu\rangle \ll 1$: gaps survive, ID tracks clean

Inherited lesson

Tasevsky et al.: combinatorial background from *fake* pileup protons dominates once $\langle\mu\rangle \gtrsim 1$. At $\langle\mu\rangle = 0.1\text{--}0.2$ (2018) the uncorrelated-proton probability per crossing is small but **not** negligible: our leading systematic.

Single tag or double tag? An explicit choice

The kinematics must not be conflated

Double tag (central diffraction, both protons):

$$M_X = \sqrt{\xi_1 \xi_2 s} \approx 40 \text{ GeV for } \xi_{1,2} \simeq 3 \times 10^{-3}$$

The central mass *is* fixed by the two momentum losses.

Single tag (one proton, other side dissociates):

$$M_X = \sqrt{\xi s} \approx 700 \text{ GeV for } \xi = 3 \times 10^{-3}$$

The instanton is a sub-process *inside* a dissociating system of that mass. m_I is set by the steeply falling $\hat{\sigma}(m_I)$, not by the tag.

Our choice: single tag + forward veto (primary)

- Exactly one ALFA proton; on the *untagged* side, veto activity with MBTS, FCal and LUCID
- **Why:** rate is the binding constraint (yield slide). Double tagging costs a further Pomeron flux $\times$ acceptance factor we cannot afford

Single tag: what it costs; double tag as secondary channel

Cost and consequence of the single tag

- **Cost:** the untagged proton dissociates. Pomeron and proton remnants fragment, and part of that activity lands in $|\eta| < 2.5$. Gap-survival suppression of MPI applies **only on the tagged side**
- **Consequence:** remnant contamination and veto efficiency must be modelled, not assumed. Both enter the systematics

Double tag (secondary channel)

Genuine DPE: clean central system, mass fixed by $\xi_1 \xi_2$, theory in hand (**PRD 105 (2022) 036008**, k_T -factorisation with H1 Fit B DPDFs). But the elastic-overlay background moves **from CR1 into the signal region**, and the rate falls further. Carried as a cross-check, not the headline.

Order-of-magnitude yield: the cost of the tag

Factor	Value / range	Source, status
Luminosity ratio $664/1.616 \text{ nb}^{-1}$	$\times 411$	PoS(LHCP2023)250; before prescales (A1)
Trigger prescale factor	$\times (\leq 1)$	unknown until the 2018 menu is recovered
Diffraction fraction $f_{\mathbb{P}}$	$\times 10^{-2}-10^{-3}$	PRD 104 (2021) 054013 / PRD 105 (2022) 036008
Gap survival S^2	$\times 0.05-0.1$	KKMR
ALFA (ξ, t) acceptance $\mathcal{A}$	$\times 0.1-0.5$	SD fiducial, JHEP 02 (2020) 042
Trigger $\times$ selection ε	$\times 0.3-0.6$	SD paper; forward veto to be added
Net $N_{\text{exp}}^{2018 \text{ tag}} / N_{\text{exp}}^{2015}$	$\sim 10^{-2}-10^0$	factor ~ 10 band, both ways

Order-of-magnitude yield: what this means

$$N_{\text{exp}} = \sigma_{\text{inst}}(m_l > m_l^{\text{min}}) \times f_{\mathbb{P}} \times S^2 \times \mathcal{A}(\xi, t) \times \varepsilon_{\text{trig}} \times \varepsilon_{\text{sel}} \times \int \mathcal{L} dt$$

Rather than an absolute cross-section, build upon the *signal expectation of the original analysis* and track only the ratio.

What this means


The $\times 411$ in luminosity is **cancelled, and probably over-cancelled**, by the diffractive fraction, gap survival and acceptance. **We expect comparable or fewer signal events than the original analysis, not more.** The tag buys *purity*, not rate. Any slide that leads with “400 $\times$ ” is misleading and has been removed.

Why this belongs at Phase 0

Every factor above is estimable *today* from published material. None requires ALFA reconstruction or an MC production request. The go/no-go criterion (expected tagged signal below $\mathcal{O}(10)$ events $\Rightarrow$ stop) has therefore been **moved from Phase 4 to Phase 0**, before any resource is committed.

The ALFA sub-detector

Section 3 — The ALFA sub-detector



S. Abdel Khalek et al.; JINST 11 (2016)
P11013

Layout

- Absolute Luminosity For ATLAS: 8 Roman Pots in 4 stations, both sides of IP1
- Inner stations A7L1/A7R1 at $z = 237$ m; outer B7L1/B7R1 at $z = 245$ m
- Outer stations moved in LS1: lever arm $4 \rightarrow 8$ m, angular resolution improved by $\sim 2\times$
- Two spectrometer arms (upper-A + lower-C, and the mirror)

Why high- β^* is mandatory; ALFA acceptance in (ξ, t)

Why high- β^* is mandatory

- Parallel-to-point focusing in y : $y_{RP} \propto \theta_y^*$, independent of vertex position
- $|t_{\min}| = p^2 n^2 \epsilon_N / \beta^*$: large β^* and small n open up small angles
- Pots retract during ramp, approach the beam only in stable, scraped, low-intensity fills; beam intensity capped at 3×10^{11} p/beam in the 2.5 km fills

Correcting a common shorthand

$\xi = 3 \times 10^{-3}$ is the *worked example* of Khoze, Khoze, Milne & Ryskin ($x_L \simeq 0.997$). It is not a detector property. The measured ALFA fiducial region (ATLAS SD analysis, JHEP 02 (2020) 042) is a **broad band**:


$$-4.0 < \log_{10} \xi < -1.6 \quad \text{and} \quad 0.016 < |t| < 0.43 \text{ GeV}^2$$

ALFA acceptance in (ξ, t) : what follows

What follows

- Single tag: the dissociating-system mass $M_X = \sqrt{\xi s}$ spans 130–2000 GeV across the band ($\log_{10} \xi = -4.0 \rightarrow -1.6$ at 13 TeV). It is not m_I
- Double tag: $\xi_{1,2} = 3 \times 10^{-3} \Rightarrow M_X \approx 40$ GeV; $\xi_{1,2} \leq 0.025$ gives $M_X \lesssim 325$ GeV, so the accessible central mass range is **broader than a single worked point suggests**
- Acceptance is not flat across the band (ξ -acceptance $\approx 30\%$ where at least half of maximum; $\geq 10\%$ across the $|t|$ range); the map must be taken from the SD analysis per (ξ, t) bin

The schematic (ξ, t) acceptance map of the long version was LLM-generated and no equivalent public figure exists (the SD paper states the fiducial band in text only), so it is omitted here and its content is given in text.



Khalek et al., JINST 11
(2016) P11013

Main detectors

10 modules $\times$ 2 $\times$ 64 square scintillating fibres, $0.5 \times 0.5 \text{ mm}^2$, in u/v at $\pm 45^\circ$ on a Ti plate, layers staggered by 1/10 pitch. Edgeless; MAPMT readout.

Performance

Resolution $\approx 30\text{--}35 \mu\text{m}$ per station; alignment $20\text{--}30 \mu\text{m}$. A track needs ≥ 6 overlapping fibres in each of u and v ; an elastic proton fires 18–19. Single-layer efficiency $\approx 90\%$.

ODs and trigger

Overlap detectors (3 layers of horizontal fibres) measure the upper–lower pot separation with halo protons: key to vertical alignment. 3 mm scintillator tiles trigger, $\varepsilon_{\text{trig}} > 99.9\%$.

ALFA datasets relevant to this proposal

Section 4 — LHC runs, datasets and run conditions

$\sqrt{s}$	Period	β^* [m]	L_{int}	Precision	$\langle\mu\rangle$
7 TeV	Oct 2011	90	$78.7 \mu\text{b}^{-1}$	2.3%	0.035
8 TeV	Jul 2012	90	$496.3 \mu\text{b}^{-1}$	1.5%	< 0.08
8 TeV	Jul 2012	90	1.67nb^{-1}	1.5%	< 0.08
8 TeV	2012	1000	—	—	low
13 TeV	Oct 2015	90	—	—	—
			TBC		
13 TeV	Sep 2016	2500	$339.9 \mu\text{b}^{-1}$	2.15%	0.002–0.006
13 TeV	Jul 2018	90	664nb^{-1}	1.0%	0.1–0.2

ALFA datasets relevant to this proposal: purpose and streams

$\sqrt{s}$	Period	β^* [m]	Purpose / stream
7 TeV	Oct 2011	90	elastic; 1 bunch/beam, 4 h run
8 TeV	Jul 2012	90	elastic (σ_{tot} , PLB 761 (2016) 158)
8 TeV	Jul 2012	90	MBTS $\times$ ALFA SD stream (JHEP 02 (2020) 042)
8 TeV	2012	1000	low- $ t $ elastic
13 TeV	Oct 2015	90	RP's inserted at 10σ (JINST 11 (2016) P11013); no ATLAS physics publication, status to be confirmed
13 TeV	Sep 2016	2500	CNI, ρ -parameter (EPJC 83 (2023) 441)
13 TeV	Jul 2018	90	elastic dip + low-mass diffraction

The primary dataset

The primary dataset

The July 2018 $\beta^* = 90$ m campaign delivered $664 \pm 7 \text{ nb}^{-1}$ at $\sqrt{s} = 13 \text{ TeV}$ with ALFA inserted, the largest ALFA sample at our energy. **This is not a $400\times$ gain in signal.** The diffractive fraction, gap survival and (ξ, t) acceptance cancel it (yield slide). The dataset is chosen for *purity*, not rate.

Caveats before we get excited

- $\langle\mu\rangle = 0.1\text{--}0.2$, higher than 2015's 0.028: pileup protons are now a real background
- The MB trigger in these fills was almost certainly **prescaled**; the effective MB $\times$ ALFA luminosity must be established from the trigger menu, not assumed
- ALFA acceptance in ξ and t costs a large factor

Run conditions in detail (I)

2018, $\beta^* = 90$ m — *primary dataset*

- ~ 80 – 700 colliding bunches; 100 ns spacing for the first two days at ALFA's request, then 50 ns
- $\langle\mu\rangle = 0.1$ – 0.2 ; $\mathcal{L}_{\text{inst}} \approx (2.4\text{--}10) \times 10^{30} \text{ cm}^{-2}\text{s}^{-1}$; luminosity from LUCID Bi2EventOR_A, 1.0% precision
- Reaches the diffractive dip in t : good ALFA acceptance across our ξ range

2016, $\beta^* = 2.5$ km — *ultra-clean control*

- Fills 5313, 5314, 5317, 5321 $\rightarrow$ runs 308979, 308982, 309010, 309039, 309074, 309165, 309166
- 4–5 colliding bunches, $\langle\mu\rangle = 0.002$ – 0.006 : **effectively zero pileup**; RPs at 3σ , ≈ 0.5 mm from the orbit; beams scraped to 2σ , re-scraped every 30–90 min
- Only $339.9 \mu\text{b}^{-1}$: too little for a search, ideal for validation

Run conditions in detail (II)

2012, $\beta^* = 90$ m — *methodological template*

- $\sqrt{s} = 8$ TeV, $\langle\mu\rangle$ never above 0.08; RPs at $9.5\sigma = 7.5$ mm from the beam centre
- Trigger: ≥ 2 MBTS counters one side $\times$ ALFA near+far coincidence on the *opposite* side; $1.67 \pm 0.03 \text{ nb}^{-1}$ after prescales
- The ATLAS SD measurement: **the exact ALFA+ID combination we need**, already validated

Sourcing

2018 luminosity: PoS(LHCP2023)250; 2016 fills, $\langle\mu\rangle$: EPJC 83 (2023) 441; 2012 trigger, prescales: JHEP 02 (2020) 042; detector: JINST 11 (2016) P11013. **Cite the source, not the slide.**

Action item A1

Recover the 2018 GRL, trigger menu and prescale history for the $\beta^* = 90$ m fills: the most important unknown is the effective MB $\times$ ALFA luminosity. *Owner: analysis contact.*
Milestone: 4 weeks.

Proposed dataset

- 1 **Signal region dataset:** 2018 $\beta^* = 90$ m, $\sqrt{s} = 13$ TeV, 664 nb^{-1} delivered
- 2 **Zero-pileup control:** 2016 $\beta^* = 2.5$ km, $339.9 \mu\text{b}^{-1}$: validates the gap and proton selection free of pileup
- 3 **Cross-energy check:** 2012 $\beta^* = 90$ m, 8 TeV, 1.67 nb^{-1} : instanton rate scales differently from diffraction with $\sqrt{s}$
- 4 **Legacy overlap:** the 2015 $\mu \approx 0.028$ sample remains available for the untagged arm of a combined fit

Data quality

GRL must require ALFA *and* ID simultaneously good. Exclude scraping periods and any LB with dead-time fraction $> 5\%$ (following EPJC 83 (2023) 441). Apply the same event-cleaning as the original note.

Inherited from ANA-STDM-2021-01

- Tight Primary tracks, $p_T > 500 \text{ MeV}$, $|\eta| < 2.5$, $|d_0^{\text{BL}}| < 1.5 \text{ mm}$, $|z_0^{\text{PV}} \sin \theta| < 1.5 \text{ mm}$
- ≥ 1 primary vertex; a second PV allowed only if it has ≤ 4 tracks
- Retain the $10 < S_T < 20 \text{ GeV}$ window *as a cross-check region*, but see below

New for ALFA

- Exactly one reconstructed ALFA proton, in one armlet
- Track segment $\geq 6 \text{ } u$ and $\geq 6 \text{ } v$ fibre layers; fiducial y between ~ 8 and 20 mm from the beamline
- 3σ cut in the (x, θ_x) correlation ellipse to reject halo
- **Reopen the S_T window**: the tag changes the kinematics, so the ML optimisation must be redone rather than inherited

Section 5 — Analysis strategy

Inherited from the ID

- Mass-sensitive: N_{trk} , S_T , leading track p_T , $\langle p_T \rangle$, m_T , M
- Isotropy: $\langle \eta \rangle$, $\sigma(\eta)$, $\sigma(\phi)$, $\langle |\Delta\eta(i,j)| \rangle$
- Event shapes: thrust, sphericity, transverse sphericity
 $S_T = 2\lambda_2/(\lambda_1 + \lambda_2)$

New from ALFA

- t from the proton scattering angle, resolution $\approx 15\%$
- $\xi^{(\text{ALFA})} = 1 - E'_p/E_p$, independent of the ID
- $\xi^{(\text{central})} \simeq \sum_i (E_i \pm p_{z,i})/\sqrt{s}$ built from calorimeter clusters + MBTS, not ID tracks (next slide)
- $\Delta\eta$: the rapidity gap, measured with calorimeter clusters and MBTS out to $|\eta| \approx 4.9$, following the SD analysis
- Forward veto on the untagged side (MBTS, FCal, LUCID)

Flavour democracy: promoted to Phase 0

One $q\bar{q}$ pair per light flavour predicts an **enhanced strange-to-non-strange ratio** in the excess region, relative to any MPI-driven background: the closest thing ATLAS has to a smoking gun. V^0 reconstruction is well established at low p_T and **does not depend on ALFA at all**. A K_S and Λ yield versus N_{trk} (or versus classifier score) can be run on the 2015 sample now, and would sharpen or undercut the 3.5σ within weeks, before any ALFA work begins.

Why not ID tracks for $\xi^{(\text{central})}$?

With charged tracks only, $|\eta| < 2.5$, $p_T > 500$ MeV, the estimator misses all neutrals and everything forward, and would correlate poorly with $\xi^{(\text{ALFA})}$. The track-only restriction was inherited from the original analysis, where it was motivated by the isotropy discriminants; **nothing in the instanton hypothesis forbids calorimeter information for ξ , the gap, or the veto**.

The most powerful new variable, and a new approach to classifiers

$\xi^{(\text{ALFA})}$ vs $\xi^{(\text{central})}$ correlation

For a *genuine* single-tagged event the two must agree within resolution: the same dissociating system produced the central activity and the proton's energy loss. For **combinatorial pileup** (an ALFA proton from one collision overlaid on a minimum-bias system from another) they are uncorrelated. A data-driven, MC-independent discriminant against the background that ruins the AFP approach. It did not exist in the original analysis.

New approach to classifiers

Apart from the observable-fed DNN, we introduce Convolutional Neural Networks (CNN) using raw track information (distribution on the η/ϕ cylindrical surface), effectively turning the tracks of each event into an image fed to a network class designed for image recognition. Further, and perhaps more importantly, forced independence is obtained not with ABCDisCo but with the **Numerical Rosenblatt Method** (arXiv:2509.25521).

What carries over unchanged

- SHERPA 3.0.1 instanton generator, run card built from the tabulated $\hat{\sigma}(\sqrt{s'})$ of JHEP 2020, 201, smoothed by polynomial interpolation to 100 points
- Massive RAMBO phase space $\Rightarrow$ rest-frame isotropic decay; N_g Poisson-distributed about $\langle N_{\text{gluons}} \rangle(s')$
- CT18NNLO nominal (positive-definite, survives the downward μ_f variation); NNPDF3.0 NNLO alternative
- Flavour thresholds INSTANTON_B/C_PRODUCTION_THRESHOLD = 100/20 GeV
- AHADIC cluster hadronisation nominal; PYTHIA 8 Lund string alternative

What must change

- 1 Diffractive production mode.** Instantons from a *proton-Pomeron* collision, not gg from proton PDFs. Convolve with diffractive PDFs (H1 2006 Fit B) and apply S^2 . **This is exactly the calculation of PRD 105 (2022) 036008** (k_T -factorisation, Pomeron-Pomeron single-gluon coupling dominant): use it as the reference for the MC request
- 2 Mass slices retuned.** The steeply falling $\hat{\sigma}(m_l)$, *not* the tag, makes 20–60 GeV dominant: finer slicing there, higher slices kept for the shape tail

What must change

- 3 **Proton transport** through the LHC lattice to $z = 237$ and 245 m via the thin-lens transfer matrix, then ALFA digitisation
- 4 **Simulate with pileup.** The original samples had **zero** additional interactions: unacceptable at $\langle\mu\rangle = 0.1\text{--}0.2$

Action item A2: raise an ATLMCPROD request for diffractive instanton samples with proton transport and $\langle\mu\rangle$ overlay: the long pole in the schedule, start it first.

Physics backgrounds: new hierarchy

- 1 **Single diffractive dissociation** $pp \rightarrow Xp$: now the **dominant** background, not minimum bias. PYTHIA 8 A3 tune (Donnachie–Landshoff, $\alpha(0) = 1.07$); A2 tune (Schuler–Sjöstrand) as alternative
- 2 **Central diffraction** $pp \rightarrow pXp$: naturally produces forward protons; $\sim 10\%$ of the SD rate in the ATLAS SD analysis
- 3 **Double diffractive and non-diffractive**: negligible as correlated sources, but feed the overlay
- 4 **Hadronisation systematic**: HERWIG 7 7.1.3 cluster model against PYTHIA 8 string

Note: ATLAS measured the fiducial SD cross-section as 1.59 ± 0.13 mb, well below PYTHIA 8's 2.52 mb (A3) and HERWIG 7's 4.96 mb. **Normalise the SD background to the ATLAS measurement, not to the generator default.**

Combinatorial (overlay) background: data-driven

An ALFA proton uncorrelated with the central activity: from elastic pileup, CD, showering, or beam halo. Method (following JHEP 02 (2020) 042): select a control sample with hits in all 32 MBTS segments and tracks near both edges of the ID acceptance, a topology that heavily suppresses genuinely correlated forward protons. The rate of ALFA protons in that sample fixes the normalisation: in the 2012 data 0.77% per crossing, with a 9% correction for residual signal. **Must be re-derived for 2018:** at $\langle\mu\rangle = 0.1\text{--}0.2$ rather than < 0.08 , expect a larger value.

Pseudo-data

Build the Asimov dataset from SD+CD+overlay at the measured normalisations. Run signal-injection tests at $\mu = 0, 1, 2.33$: the last to check we would recover the original analysis's fitted strength if it is real.

Background estimation and the ML approach

Keep ABCDisCo, with changes

- Two classifiers trained or modified so they are independent for background:
 $N_A^{\text{pred}} = N_B N_C / N_D$; binned extrapolation into the upper-right of the classifier plane, as before
- **New:** $\xi^{(\text{ALFA})}$, t , $\Delta\eta$ and $\xi^{(\text{ID})}$ enter as inputs; the ξ -consistency variable is held *out* of the classifiers and used as an orthogonal validation axis

Known failure mode to fix

In the original analysis ABCD closure degraded to $> 10\%$ towards the top right of the plane, because the distance correlation is dominated by the background peak at low classifier values. With a much smaller sample after the ALFA tag this will **get worse, not better**. Mitigations (larger batch, reweighted DisCo loss, mass-decorrelated architecture) are reasonable but untested. **The simpler fallback, a template fit in $\Delta\eta$ and S_T , is developed in parallel from Phase 5:** for a cross-check, a fully closed background estimate beats a stronger classifier carrying a 10% closure systematic.

Control and validation regions

- **CR:** $S/B \lesssim 0.2\%$, low-classifier corner: constrains the background model
- **VR:** intermediate, excluded from the fit, tests extrapolation
- **New CR1:** *two* ALFA armlets with protons: dominated by elastic overlay, directly tests the combinatorial model
- **New CR2:** two armlets + 2–10 MBTS sectors: enriches central diffraction
- **PVR:** orthogonal region of the classifier plane, unbiased extrapolation check

Both new CRs are lifted directly from the ATLAS SD analysis, where they were shown to work.

Blinding

Blind the SR until the CR/VR description is signed off, exactly as before. Given that we are checking a 3.5σ excess, the blinding discipline matters more here than in a null search.

Source	Size (ref.)	Treatment
Category: Signal theory		
Renorm. scale μ_r , custom variation of $(\rho\mu_r)^{b_0}(\bar{\rho}\mu_r)^{b_0}e^{-4\pi/\alpha_S}$	orders of mag. on norm.	On-the-fly weight; shape-only in the fit
Factorisation scale $\mu_f = v/\rho$, $v = 2^{\pm 1}$	—	On-the-fly PDF-ratio weight
CT18NNLO 59 Hessian eigenvectors + $\alpha_S = 0.117, 0.119$	$\lesssim 2\%$ on shapes	Envelope; NNPDF3.0 as alternative set
Hadronisation: AHADIC cluster vs Lund string	few %	Two-point

Systematic uncertainties I: diffractive modelling

Source	Size (ref.)	Treatment
Category: Diffractive modelling (<i>new</i>)		
Pomeron flux: Donnachie–Landshoff vs Schuler–Sjöstrand	$\alpha(0) = 1.07$ vs 1.00	Two-point, correlated with the SD background
Gap survival factor S^2	0.05–0.1	Normalisation only; absorbed by μ <i>within</i> the tagged channel, which is why μ cannot be shared with the untagged 2015 result

As in the original analysis, the enormous normalisation uncertainty on the predicted cross-section does *not* degrade a discovery shape-fit. Only shape variations matter: S^2 , for all the attention it attracts, is not a problem here.

Systematic uncertainties II: tracking, luminosity, ABCD

Source	Size (ref.)	Treatment
Category: Tracking (inherited)		
TRK_EFF_TIGHT_{GLOBAL, IBL, PP0, PHYSMODEL}, TRK_FAKE_RATE_TIGHT	ID material ~ 5%	Quadrature sum, applied before Tight Primary
TRK_RES_D0/Z0_MEAS; TRK_BIAS_D0/Z0_WM	small	IP smearing; InDetTrackBiassingTool
Luminosity	1.0% (2018)	Flat; was 2.1% in the original analysis
Category: ABCD		
Non-closure, per generator	~ 10%, worse at high classifier	γ NPs per bin, as in the original fit

Systematic uncertainties II: ALFA

Source	Size (ref.)	Treatment
Category: ALFA (<i>new</i>)		
Horizontal / vertical alignment, rotation	$\sim 1\%$ on t	Offset method, correlated across t
Beam optics: effective k -values of Q1–Q6	up to 1% on matrix elements	Alternative optics model (Q5, Q6 only)
ALFA reconstruction efficiency	91–94% per armlet	Data-driven tag-and-probe on elastics
Trigger efficiency (MBTS $\times$ ALFA)	to be measured (2012 value not carried over)	Reference-trigger method, per bin, once the 2018 menu is known (A1)

Systematic uncertainties II: background (new, dominant)

Source	Size (ref.)	Treatment
Proton overlay subtraction	$< 5\%$ typical, $\rightarrow 20\%$ at t extremes	Bin-by-bin data/prediction in CR1
CD cross-section 1.12–1.66 mb; CD ξ reweighting	$\sim 5\%$; $\leq 2\%$	Normalisation and shape NPs
Untagged-side remnant contamination in $ \eta < 2.5$	to be evaluated	Two-point on the diffractive fragmentation model
Forward veto efficiency (MBTS, FCal, LUCID)	to be evaluated	Data-driven, from the SD control samples

Overlay subtraction: the leading systematic, as in the ATLAS SD measurement (data-driven, no MC fallback). CR1 statistics should drive the dataset selection.

Likelihood — unchanged in form

$$L(\mathbf{n}|\mu, \boldsymbol{\theta}) = \prod_{i=1}^N \text{Pois}(n_i|\nu_i) \prod_{j=1}^M \text{Pois}(n_j|\nu_j) \prod_{\chi \in \mathcal{X}} c_{\chi}(\mathbf{a}_{\chi}|\chi), \quad \nu_i(\mu, \boldsymbol{\theta}) = \left(\prod_{\kappa} \kappa \right) \left(\nu_i^0 + \sum_{\Delta} \Delta \right)$$

- POI: signal strength μ , unconstrained; Gaussian constraints on all α (correlated) and γ (per-bin) NPs
- TREXFITTER from STATANALYSIS, as before

Deliverables

Expected sensitivity; expected and observed upper limits on μ and on the visible cross-section; NP pulls, ranking and correlation matrix; post-fit distributions in CR, VR, LVR, PVR and SR.

Statistical analysis: changes specific to this analysis (I)

Changes specific to this analysis

- 1 Report the global significance.** The original quotes a local number. With a mass-sliced signal and a two-classifier plane, a look-elsewhere correction is required before any claim
- 2 Pre-register a *shape* criterion.** Since normalisation carries no information, the criterion must be a statement about shapes. Pre-registered as: *a significant excess in the same isotropy variables ($\langle\eta\rangle$, $\langle|\Delta\eta(i,j)|\rangle$, transverse sphericity) in a sample where MPI is kinematically suppressed.* Test statistic: the profile-likelihood ratio q_0 on the shape-only fit, with the LEE correction applied

Statistical analysis: changes specific to this analysis (II)

Changes specific to this analysis

- 3 No shared μ in the combination.** The diffractive-to-inclusive ratio carries S^2 , the Pomeron flux normalisation and the diffractive PDFs, and is uncertain by a factor of several: **not** predicted by the inclusive calculation, so a shared μ has no physical meaning. **The 2015 and 2018 results are reported separately and compared in shape only**; a free transfer factor κ_{diff} is retained as a documented alternative, not the default
- 4 Do not tune on the excess.** Selection and classifier choices are optimised on expected sensitivity in Asimov data, never on observed SR yields

Ph.	Deliverable	Depends on	Duration
0a	Recover 2018 $\beta^* = 90$ m GRL, trigger menu, prescales; effective MB $\times$ ALFA luminosity	—	1 month
0b	Order-of-magnitude yield estimate and go/no-go from published inputs	0a	1 month
0c	K_S/Λ flavour-democracy test on the 2015 sample , no ALFA needed	—	1–2 months, in parallel
1	ALFA reconstruction, alignment and optics for the 2018 dataset; validate on elastics	Phase 0, ALFA experts	3 months
2	ATLMCPROD request: diffractive instanton signal, proton transport, μ overlay	Theory contact	2 months to submit, 4–6 to produce

Work plan (II)

Ph.	Deliverable	Depends on	Duration
3	SD/CD/DD/ND background samples, normalised to ATLAS measured σ_{SD}	Phase 2 queue	3 months
4	Overlay background model from CR1; re-derive the per-crossing fake-proton rate at $\langle\mu\rangle = 0.1\text{--}0.2$	Phase 1	2 months
5	Retrain classifiers with ALFA inputs; fix the high-classifier decorrelation failure	Phases 2–4	3 months
6	Systematics, Asimov fit, expected sensitivity; internal note and unblinding request	Phase 5	3 months
7	Unblind; combination with the 2015 untagged result	SM group approval	2 months

Critical path

Phase 2. Signal MC with diffractive production and proton transport does not exist anywhere; it must be built, and everything downstream waits on it. **Start it in week one.**

Go/no-go — now at Phase 0b

Every factor in N_{exp} is estimable today from published material. If the expected tagged signal is below $\mathcal{O}(10)$ events given achievable ABCD closure, **stop at month two** and write up the feasibility study plus the K_S/Λ result, rather than committing a year of effort first.

Risks, honestly stated (I)

Risk	Severity	Mitigation
Statistics after the ALFA tag are simply too small	High	The ξ and t acceptance costs a large factor on top of the diffractive cross-section and $S^2 \sim 0.1$. Phase-0 luminosity accounting and the go/no-go exist precisely to find this out early and cheaply.
Trigger was prescaled beyond usefulness in 2018	High	Unknown until the menu is recovered. If the MB stream is unusable, fall back to any ALFA-seeded stream with ID readout; failing that, the 2012 8 TeV SD dataset becomes the primary sample.
Combinatorial pileup background at $\langle\mu\rangle = 0.2$	Medium	The $\xi^{(\text{ALFA})}-\xi^{(\text{ID})}$ correlation is a strong, MC-independent handle unavailable to AFP-based searches. CR1 measures the rate directly.

Risks, honestly stated (II)

Risk	Severity	Mitigation
ABCD non-closure worsens with smaller samples	Medium	Mitigations as listed, <i>and</i> the $\Delta\eta/S_T$ template fit built in parallel from the start, so the choice between them is evidence-based rather than forced late.
Instanton theory in the diffractive mode is less developed than the inclusive case	Medium	Engage the authors of PRD 104 (2021) 054013 early. Treat the diffractive PDF and S^2 as explicit, documented model assumptions, not buried in a generator default.
We confirm nothing either way	Low	A null result at this luminosity is still the first forward-tagged instanton search at a hadron collider, and it constrains the interpretation of the 3.5σ . That is publishable.

The case

- ANA-STDM-2021-01: 3.5σ local excess on a 1.54σ expectation, intriguing and limited by **background modelling rather than statistics**. HERA (H1 2016: UL 1.5–6 pb) found nothing and already challenges the predictions
- An independent measurement with a different dominant background is worth more than more of the same data. **A cross-check proposal, not a discovery proposal**: a null result is the first forward-tagged instanton search at a hadron collider, constrains the 3.5σ , and is publishable

Summary: the method and the data

The method and the data

Tag a leading proton in ALFA ($10^{-4} \lesssim \xi \lesssim 0.025$) and veto the untagged side: MPI suppressed kinematically via gap survival $S^2 \simeq 0.05\text{--}0.1$; the tag does **not** select m_l . Data: 2018 $\beta^* = 90$ m, $\sqrt{s} = 13$ TeV, $664 \pm 7 \text{ nb}^{-1}$, $\langle \mu \rangle = 0.1\text{--}0.2$, luminosity to 1.0%; after $f_{\mathbb{P}}$, S^2 and (ξ, t) acceptance expect **comparable or fewer** signal events than 2015. Validated against 2016 $\beta^* = 2.5$ km (zero pileup) and 2012 8 TeV (existing ALFA+ID SD measurement).

Summary: the four things to do first

The four things to do first

- 1 Recover the 2018 trigger menu and prescales: decides whether the analysis is possible at all
- 2 **Order-of-magnitude yield estimate and go/no-go**, from published inputs, before committing resources
- 3 **K_S/Λ flavour test on the 2015 data**: no ALFA needed, and it can sharpen or undercut the 3.5σ within weeks
- 4 Submit the diffractive instanton MC request: the critical path thereafter

Backup: key references (theory)

- A. Belavin, A. Polyakov, A. Schwartz, Y. Tyupkin, *Pseudoparticle solutions of the Yang-Mills equations*, PLB 59 (1975) 85; G. 't Hooft, PRD 14 (1976) 3432; A. Mueller, NPB 348 (1991) 310
- V. V. Khoze, F. Krauss, M. Schott, *Large effects from small QCD instantons*, JHEP 2020, 201, arXiv:1911.09726
- V. V. Khoze, D. L. Milne, M. Spannowsky, PRD 103 (2021) 014017, arXiv:2010.02287
- S. Amoroso, D. Kar, M. Schott, *How to discover QCD Instantons at the LHC*, EPJC 81 (2021) 624, arXiv:2012.09120
- V. A. Khoze, V. V. Khoze, D. L. Milne, M. G. Ryskin, *Hunting for QCD Instantons at the LHC in Events with Large Rapidity Gaps*, PRD 104 (2021) 054013 — the basis of this proposal
- idem, *Central instanton production*, PRD 105 (2022) 036008 — k_T -factorisation with H1 Fit B DPDFs; the reference for the diffractive signal MC and for the double-tag channel

Backup: key references (theory, cont.; numerical techniques)

- M. Tasevsky, V. Khoze, D. Milne, M. Ryskin, *Searches for QCD instantons with forward proton tagging*, EPJC 83 (2023) 35, arXiv:2208.14089 — note: AFP/CT-PPS regime, $\xi = 0.02\text{--}0.05$, $m_I > 60$ GeV
- M. G. Ryskin, V. A. Khoze, *Instanton production in diffractive processes with a large rapidity gap*, arXiv:2606.10588 (2026)
- M. Mangano, *Remarks on the numerical impact of potential theoretical systematics. . .*, arXiv:2101.02719
- *A Numerical Rosenblatt Method for Forced Variable Independence*, arXiv:2509.25521

Backup: key references (experiment)

- H1 Collaboration, EPJC 25 (2002) 495; ZEUS Collaboration, EPJC 34 (2004) 255
- H1 Collaboration, *Search for QCD Instanton-Induced Processes at HERA in the High- Q^2 Domain*, EPJC 76 (2016) 381, arXiv:1603.05567
- S. Abdel Khalek et al., *The ALFA Roman Pot Detectors of ATLAS*, JINST 11 (2016) P11013, arXiv:1609.00249
- ATLAS Collaboration, *Measurement of differential cross sections for single diffractive dissociation in $\sqrt{s} = 8$ TeV pp collisions using the ATLAS ALFA spectrometer*, JHEP 02 (2020) 042, arXiv:1911.00453 — **the ALFA+ID template**
- ATLAS Collaboration, *Measurement of the total cross section and ρ -parameter. . . at $\sqrt{s} = 13$ TeV*, EPJC 83 (2023) 441, arXiv:2207.12246; *. . . from elastic scattering at $\sqrt{s} = 8$ TeV*, PLB 761 (2016) 158
- E. Sanzani (ATLAS), *Luminosity determination for σ_{tot} and ρ at 13 TeV*, PoS(LHCP2023)250 — **source of the 2018 664 nb $^{-1}$ figure**
- ATLAS Collaboration, *Charged-particle distributions in $\sqrt{s} = 13$ TeV pp interactions*, PLB 758 (2016) 67

Backup: the original dataset, for reference

Run	Delivered [μb^{-1}]	Prescaled L [μb^{-1}]	$\langle P \rangle$	$\langle \mu \rangle$
267358	81.1	72.3	1.12	2.4×10^{-3}
267359	240.7	108.9	2.21	6.0×10^{-3}
267360	320.9	76.7	4.18	1.6×10^{-1}
267367	1538.2	124.9	12.32	3.4×10^{-2}
267385	5932.2	707.2	8.39	3.0×10^{-2}
267599	5609.6	602.6	9.31	3.2×10^{-2}
Total [†]	13401.8	1615.9	8.30	2.8×10^{-2}

Backup: the original dataset, notes and lesson

Notes

† Run 267360 excluded: an inner-detector tracking defect disabled modules at $\eta < -2$, $\phi < -1.8$, distorting track-based observables. Removing it cost 4.5% of the luminosity and moved the model-dependent expected limit from $\mu = 15.99$ to 16.11. Trigger efficiency of HLT_noalg_mb_L1MBTS_1 rises from 98.8% at one track to 100% above five; the associated systematic (0.12%) was neglected as negligible.

The lesson for 2018

A single run with a subtle ID defect changed the distributions of exactly the variables the analysis depends on, and was only caught in a pileup-dependence study late in the review.