THE SECOND LHAASO COLLABORATION CONFERENCE IN 2026

Self-supervised Denoising and Calibrated Significance for LHAASO-KM2A Counting Skymaps

Yijun Wu (吴奕君)
Supervisor: Songzhan Chen (陈松战)
Institute of High Energy Physics, CAS
University of Chinese Academy of Sciences
wuyj@ihep.ac.cn

22 October 2026
The Academic Symposium for the 10th Anniversary of the LHAASO Collaboration

Denoise the counting skymaps with a 3-D transformer; measure the significance of an excess with a flux-plus-shape likelihood calibrated on the data itself.

01 Counting skymaps of 14 energy bands
02 3-D Restormer denoising
03 All-sky stitching
04 Flux + shape likelihood
05 Null dist. calibration
06 Significance of sources

14 · 0.1°	**8 mo ≈ 60 mo**	**11 / 12**	**same false-positive rate**
energy bands · pixel size	denoised ≈ raw fluctuation level	known-source cases improved	the comparison that matters

OPENING

- Denoise the KM2A counting skymaps
- Measure the significance of the sources
- Goal: weak and extended sources, reliable statistics
- Two halves: denoise the maps · recompute “5σ”

WHAT TO EXPECT

- Noise, not exposure, hides weak emission
- No clean map, no simulation
- Gain at the same false-positive rate
- Faint sources: not established yet

TAKEAWAY 8 months of denoised data reach the fluctuation level of 60 raw months; with a calibrated flux-plus-shape likelihood that gives **higher significance for known bright and medium-bright sources at the same false-positive rate**.

The whole chain, and why we need it

Two rebuilds: the images, and the statistics.

- **14 bands · 0.1°/pixel · 61 months** of counting skymaps
- This figure: **the whole analysis chain**
- The limit is **background noise**, not exposure
- Two parts: **denoise the maps**, **recompute 5σ**

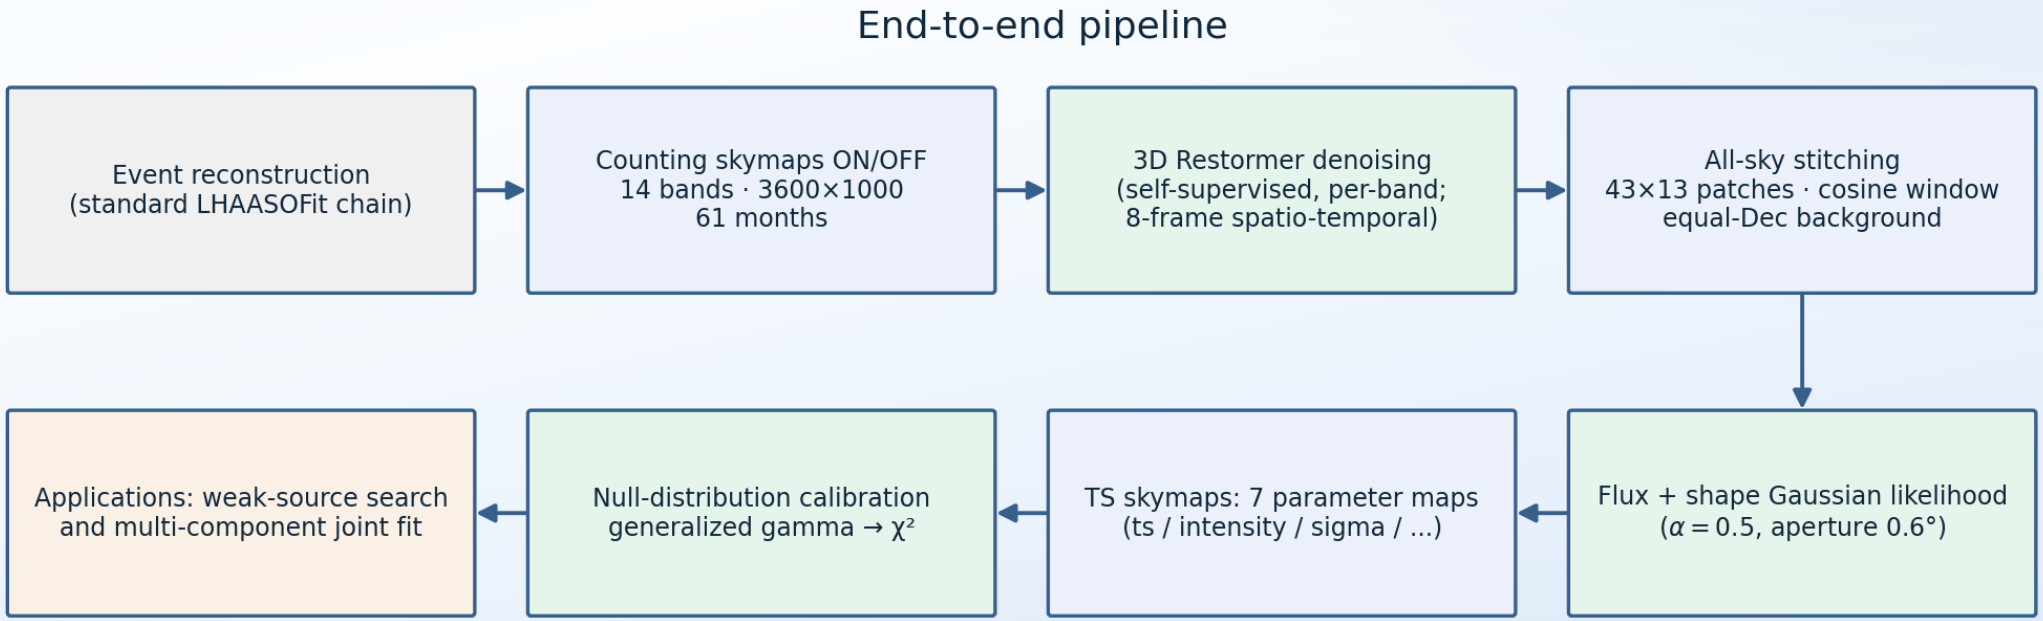
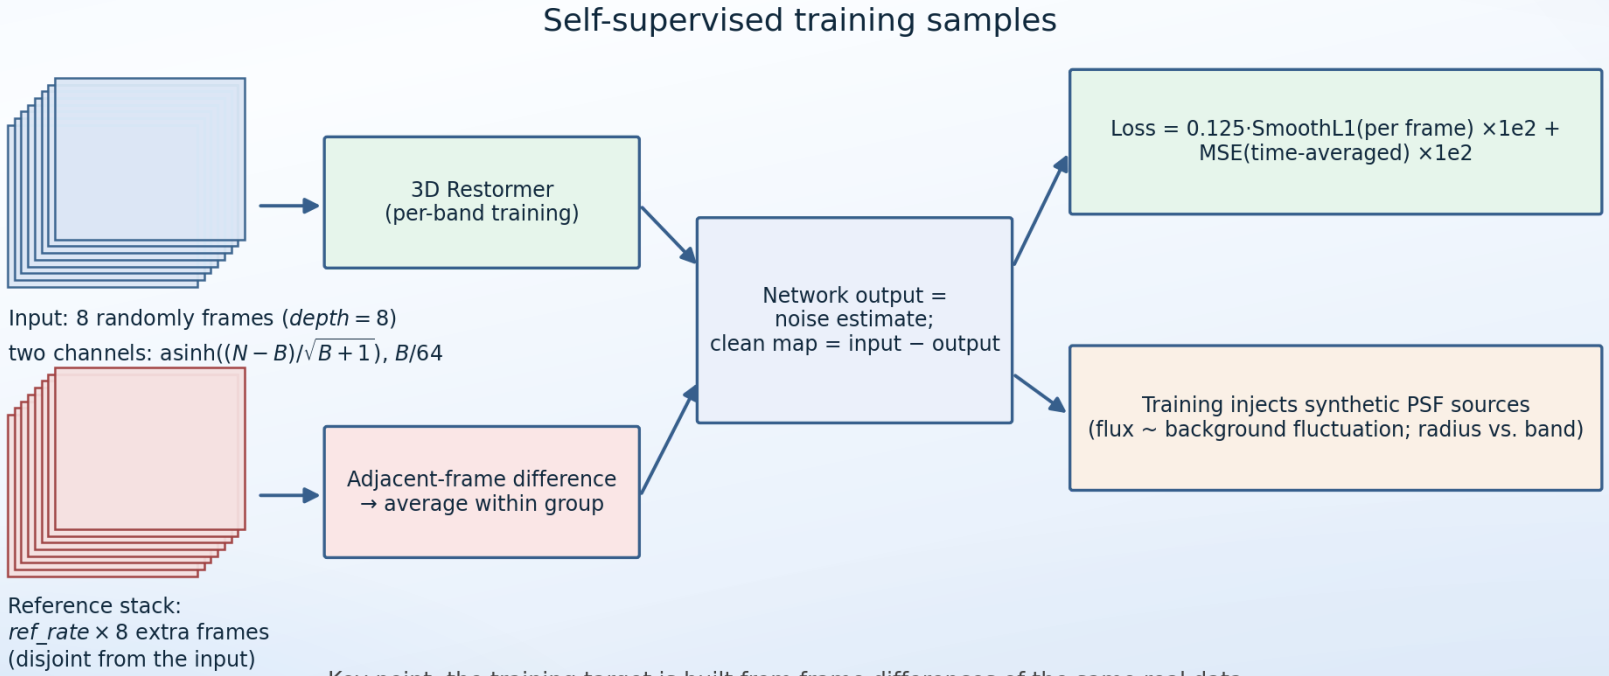


Fig. 2 The end-to-end pipeline: monthly maps → denoising → all-sky stitching (top); likelihood → TS maps → null-distribution calibration (bottom).

01 Event reconstruction | 02 Monthly maps and 61 months | 03 3-D Restormer denoising | 04 All-sky stitching 43×13 patches | 05 Flux + shape likelihood | 06 TS skymaps | 07 Null distribution calibration

Learns from the data itself — no clean map, no simulation

- Input: **8 monthly maps** (blue)
- Green: loss + artificial point sources
- Reference: **another 8 random months** (red)
- Restormer reads **several months at once**



Key point: the training target is built from frame differences of the same real data —

no clean ground-truth skymap and no synthetic background are needed: self-supervised learning.

Fig. 3a Training sample: 8-frame input, disjoint reference stack, loss and synthetic-source injection.

- **J0358+5412**, 16–25 TeV: top = traditional, bottom = ours
- **Same region, same colour scale**: peak √TS **14.5** → **49.8**

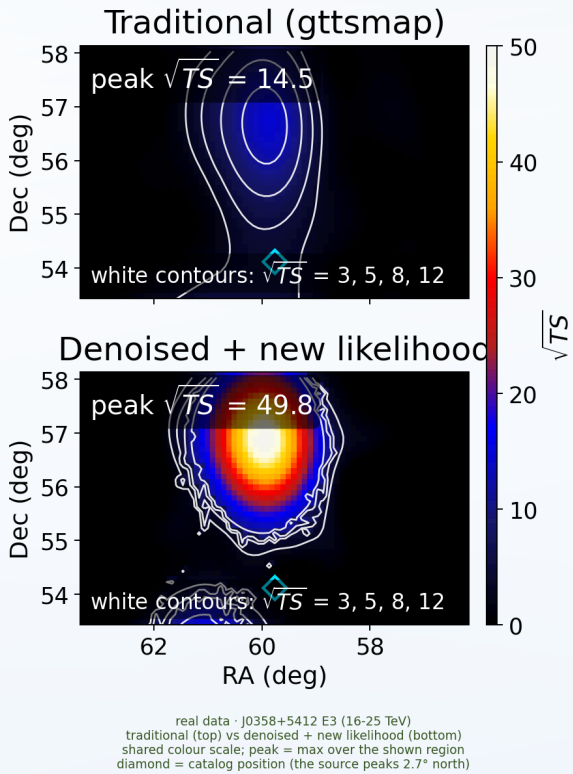


Fig. 3b √TS maps, real data, J0358+5412 (16–25 TeV): traditional (top) vs denoised + new likelihood (bottom), same colour scale. Diamond = catalog position; the source peaks 2.7° north of it.

Flux *and* shape, with thresholds measured on the data

- Likelihood = **flux** + **shape**
- Standard thresholds **no longer apply**
- Black: excess · green: flux · orange: shape
- Fit the TS tails → new 5σ per band

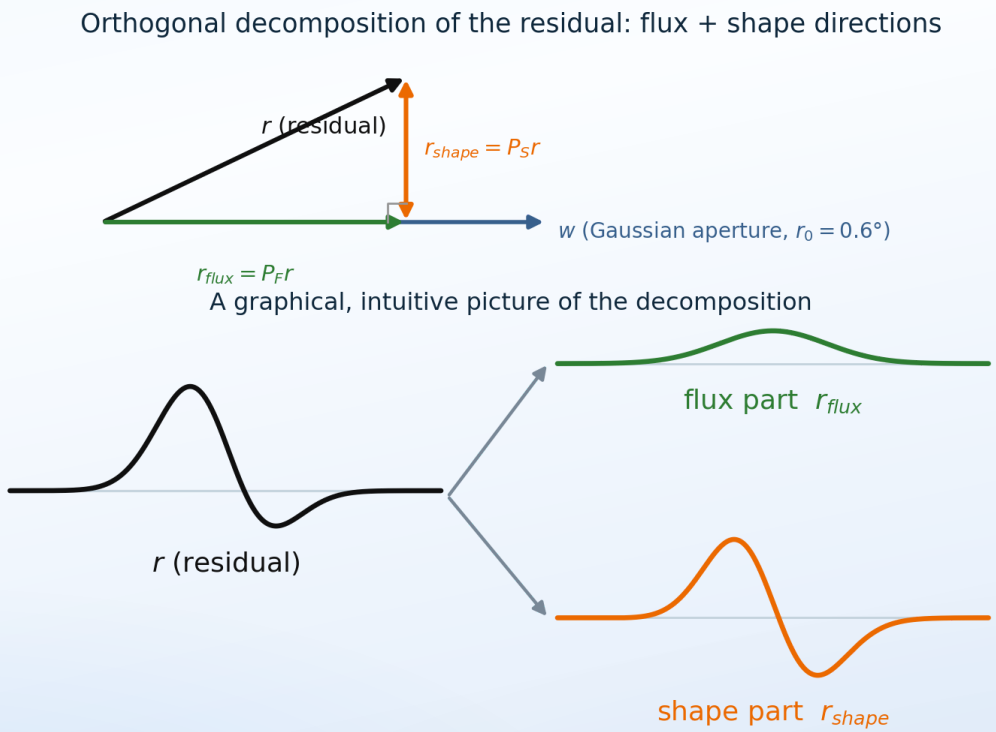


Fig. 4a Residual split into a flux part along the Gaussian aperture and an orthogonal shape part. α = 0.5, r_0 = 0.6°.

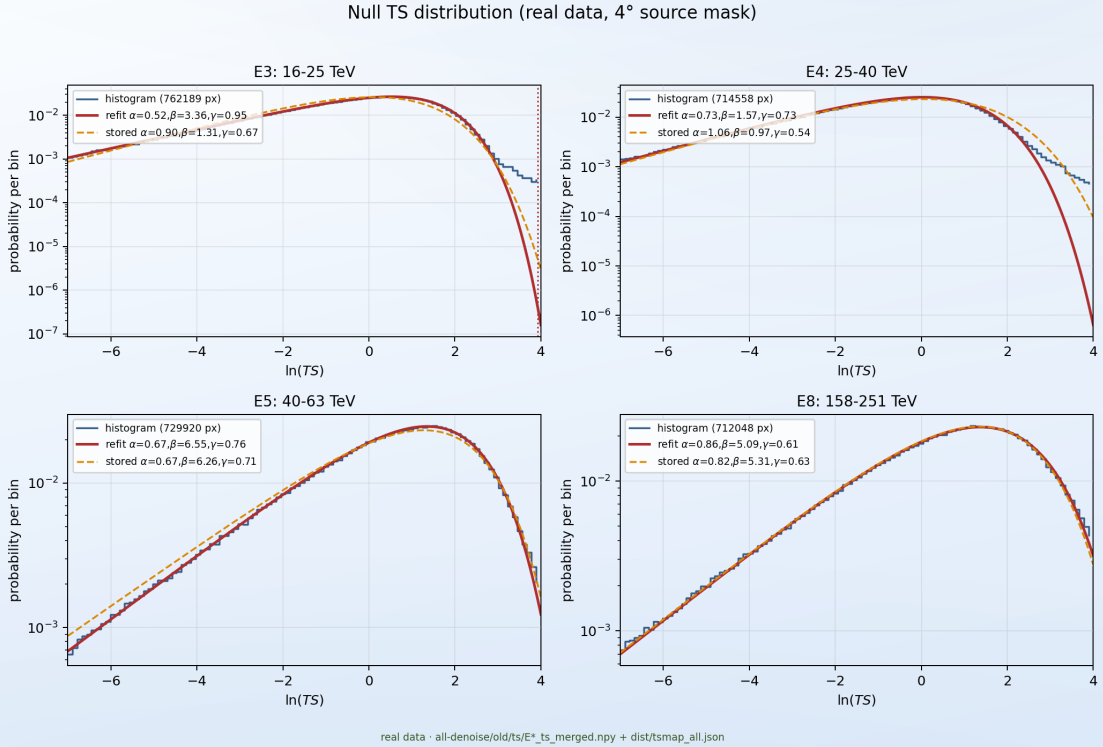


Fig. 4b Measured TS distributions in four bands; red = our fit, dashed = previous calibration — they separate in the tails.

70.8	151.1	256.5	29.4 → 66.1
$TS_{5\sigma}$ · 15–25 TeV	$TS_{5\sigma}$ · 25–40 TeV	$TS_{5\sigma}$ · 40–63 TeV	traditional → merged, 9 bands

Gain at the same false-positive rate

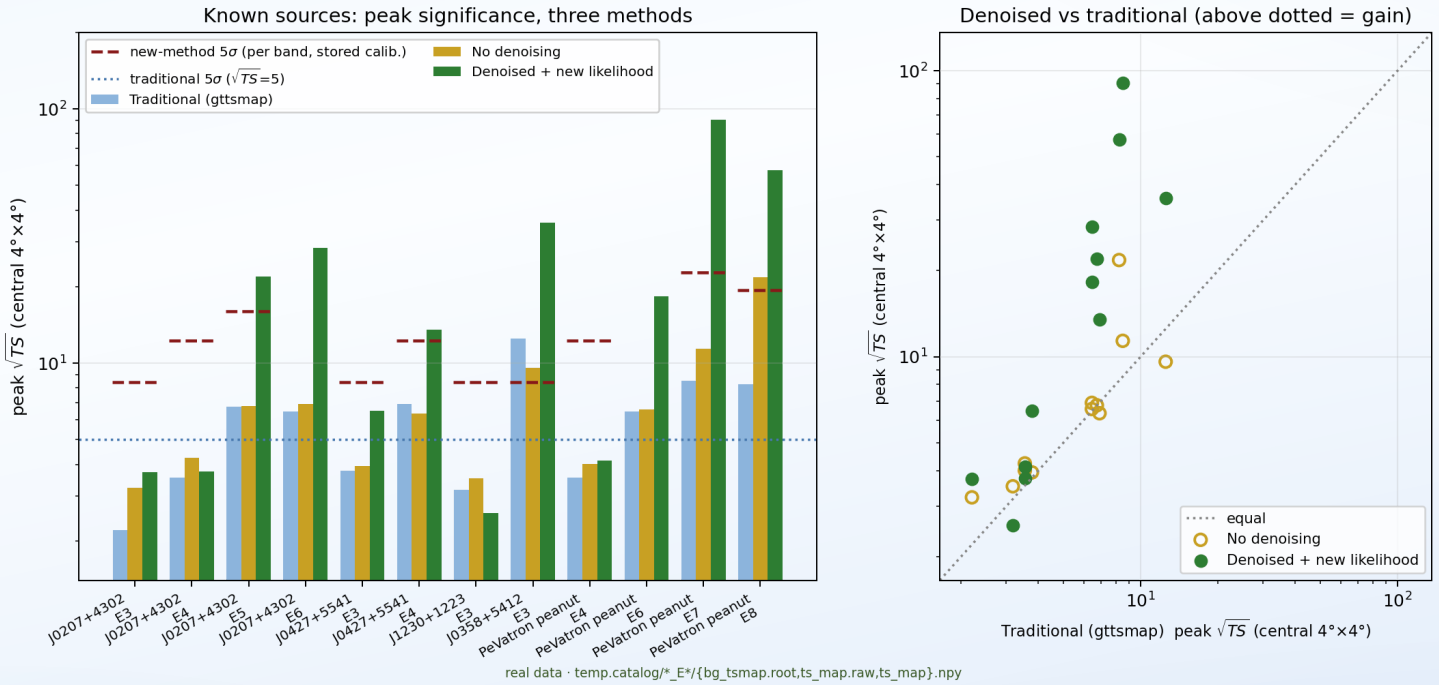


Fig. 5 Twelve (source, band) cases in the central 4°×4° window: bars = traditional / no denoising / ours, dashed red = our 5σ; right, denoised vs traditional — above the dotted diagonal is a gain.

WHAT THE FIGURE SAYS

- **11 of 12 cases** higher after denoising
- Almost every **green point above the diagonal**
- **J0358+5412**, 15–25 TeV: **≈15σ → ≈50σ** (ROI max; 12.5 → 35.8 in this window)
- A higher value alone is **not** better sensitivity — the thresholds move too
- Gain confirmed for **bright / medium-bright**; **faintest & lowest bands open**

11 / 12
cases improved (central 4°×4°)

15σ → 50σ
J0358+5412, 15–25 TeV (ROI max)

same false-positive rate
the comparison that matters

Across energy bands, and several sources at once

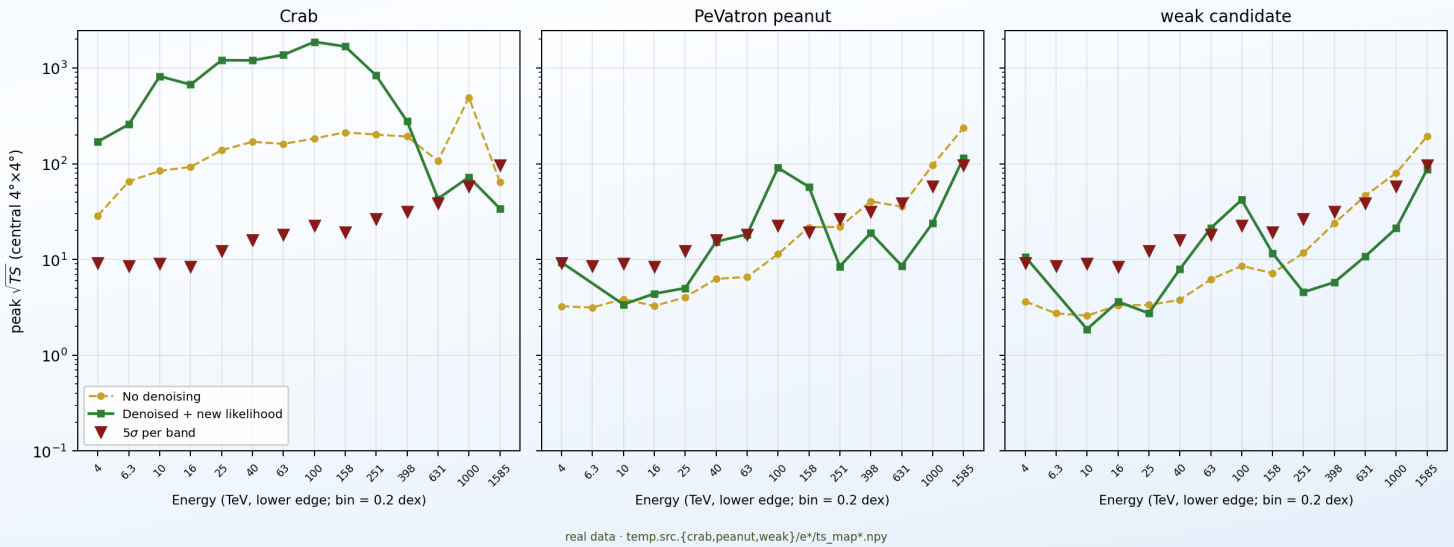


Fig. 6 Peak significance vs energy for Crab, the PeVatron Peanut and a weak candidate.

READING THE FIGURE

- Green = ours · gold = no denoising
- Red triangles = **5σ per band**
- Green above red = a gain

WHAT COMES NEXT

- ① **Joint spectral fit** — one shared spectrum across the bands; a faint source adds up, so the 5σ threshold becomes lower
- ② **Multi-source fit** — several sources in one likelihood, to separate overlapping emission in crowded regions

A complete method for KM2A counting skymaps

IN ONE LINE Denoising + a flux-and-shape likelihood + a **data-based significance calibration**: 8 denoised months reach the fluctuation level of 60 raw months, and the calibrated significance gives a robust gain for **bright and medium-bright sources at the same false-positive rate**.

① The chain

- Self-supervised denoising
- Flux + shape likelihood
- Null distribution calibrated on the data

② What we measured

- 8 denoised months ≈ 60 raw months
- Gain at the same false-positive rate
- Faint sources: not established yet

③ Still open

- Joint multi-band spectral fit
- Multi-source fitting in crowded regions
- Flux calibration, edge and baseline systematics

Thank you for your attention.

Yijun Wu · wuyj@ihep.ac.cn · Institute of High Energy Physics, CAS / UCAS