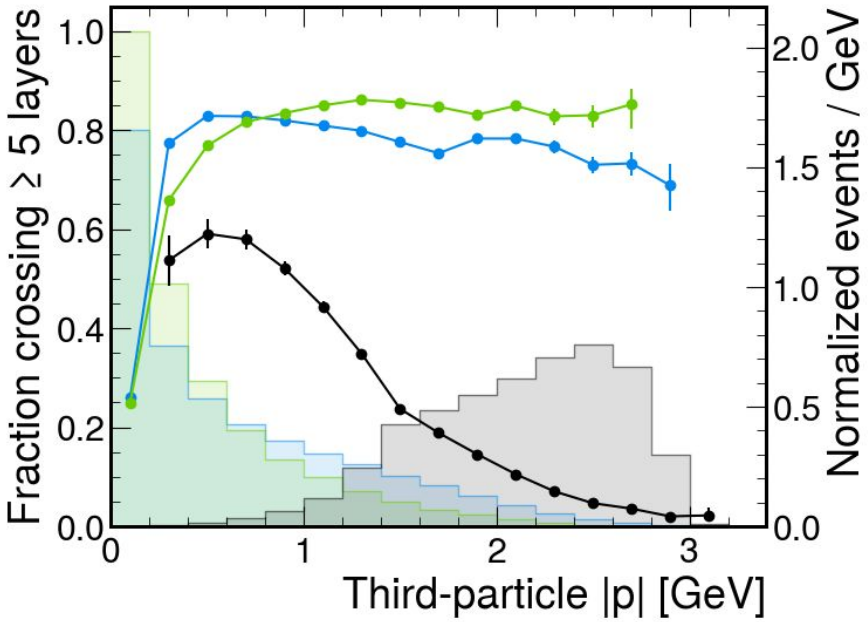
Third Track Reconstruction

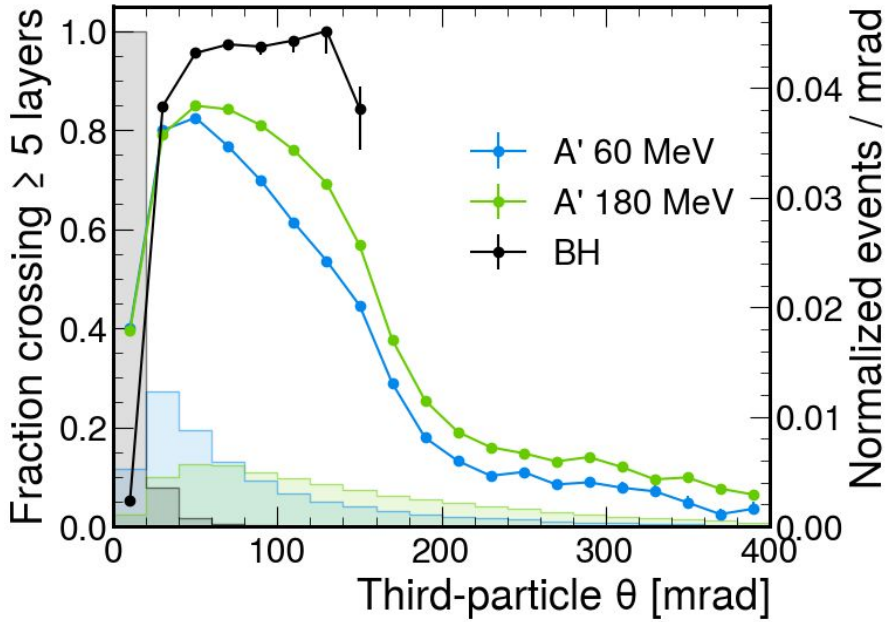
Elizabeth Berzin
Collaboration Meeting
October 20, 2026

U.S. DEPARTMENT OF ENERGY
Office of Science
Stanford University
SLAC NATIONAL ACCELERATOR LABORATORY

Motivation

- Acceptance to the recoil track is very different between signal and BH tridents.
- Existing recoil track variables provide a handle, but resolution is low—goal was to understand if dedicated third track hit finding, for small numbers of hits, can help discriminate signal vs. background.
- Previous iterations of this work did not factor in preselection, mass window effects—this talk attempts to propagate hit finding to what things might look like with a final selection.



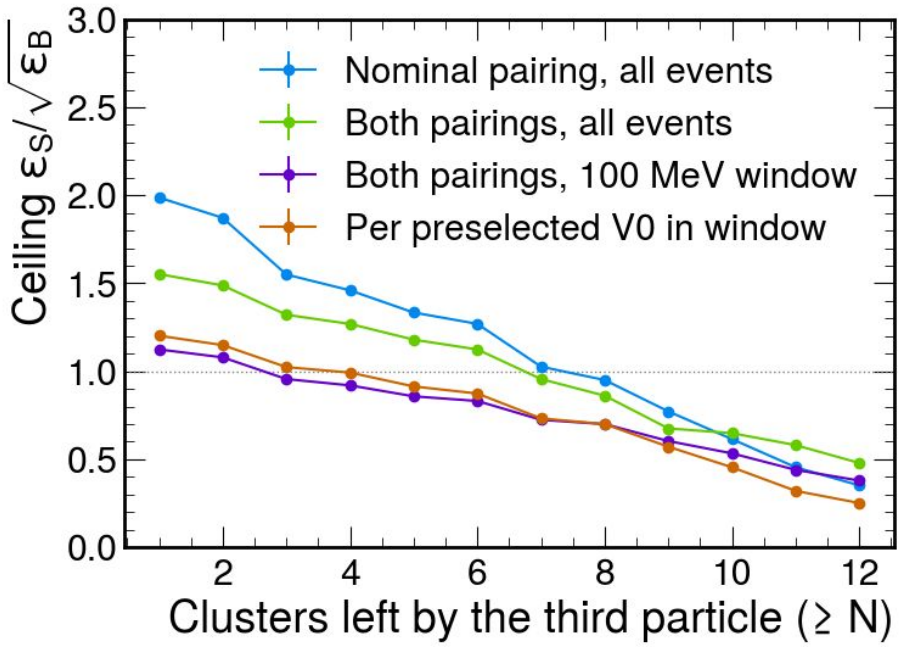


Outline

- Seed calculation:
  - Including a digression about pulser contamination on tracks, effects of preselection.
  - Seed pulls.
- Hit finding:
  - Method and configuration.
  - “Road” diagnostics.
  - Beamspot constraint.
- Refitting.
- Performance metrics.

Expected Performance

- Most aggressive way to apply a cut is to require a third track with >= N hits in the signal selection.
- Evaluating best possible performance by looking at all reconstructed V0s, truth level information on the third track.
- After preselection, among events in the 100 MeV mass window, we get a possible improvement only for very short track segments.
- Note: this is also before considering correlations with other selections: psum, pT, all of which will select for events that lower this ratio.
- This suggests that third track identification will provide limited improvement in a cut-based selection, but could potentially be used as an additional variable in a discriminator.
    - Potentially also as a way to identify when we have a mispaired vertex?



Calculating Seeds

- From the vertex fit, we get the vertex position and momentum for both tracks, contributing 9 fit parameters and a 9 x 9 covariance matrix.

$$\begin{array}{c} \\ \mathbf{x}_{vtx} \\ \mathbf{p}_1 \\ \mathbf{p}_2 \end{array} \begin{array}{c} \begin{array}{ccc} \mathbf{x}_{vtx} & \mathbf{p}_1 & \mathbf{p}_2 \end{array} \\ \begin{bmatrix} C_{vtx} & \mathrm{cov}(\mathbf{x}, \mathbf{p}_1) & \mathrm{cov}(\mathbf{x}, \mathbf{p}_2) \\ \cdot & C_1 & C_{12} \\ \cdot & C_{12}^{\top} & C_2 \end{bmatrix} \end{array}$$

- The vertex pair momentum is given by:

$$p_{pair} = p_1 + p_2$$
$$C_{pair} = C_1 + C_2 + 2C_{12},$$

- Recoil track momentum is given by:

$$p_{recoil} = p_{beam} - p_{pair}$$

Nuclear recoil $\boxed{C_{beam} = \mathrm{diag}(\sigma^2_{beam})}$

$$C_p = C_{pair} + C_{beam}$$

- The position covariance C_x can be given either by propagating uncertainties through a projection of the V0 to the target, or externally through the beamspot constraint.
  - The majority of work was done before the beamspot bug in MC was identified, so all verification of seed calculations here uses the back-projection calculation.

Calculating Seeds

- If back-projecting from the vertex to the target, the position covariance is given by:

$$G = \partial \mathbf{x}_{\text{recoil}} / \partial (\mathbf{x}_{\text{vtx}}, \mathbf{p}_{\text{pair}})$$

$$C_x = G \ \text{blockdiag}(C_{\text{vtx}}, C_{\text{pair}}) \ G^{\top}$$

$$C_x[2][2] \mathrel{+}= \sigma^2_{\text{bs},z}$$

  - A cross term comes from the fact that both the recoil position and momentum depend on the pair momentum:

$$\text{cov}(\mathbf{x}_{\text{recoil}}, \mathbf{p}_{\text{recoil}}) = G_p \ C_{\text{pair}} \left(\frac{\partial \mathbf{p}_{\text{recoil}}}{\partial \mathbf{p}_{\text{pair}}} \right)^{\top} = - \ G_p \ C_{\text{pair}}$$

- This is missing the $\text{cov}(\mathbf{x}, \mathbf{p}_1)$ $\text{cov}(\mathbf{x}, \mathbf{p}_2)$ terms, which don't get propagated in the vtx-fit output.
- To correct for this, the position covariance is inflated by some factor f:

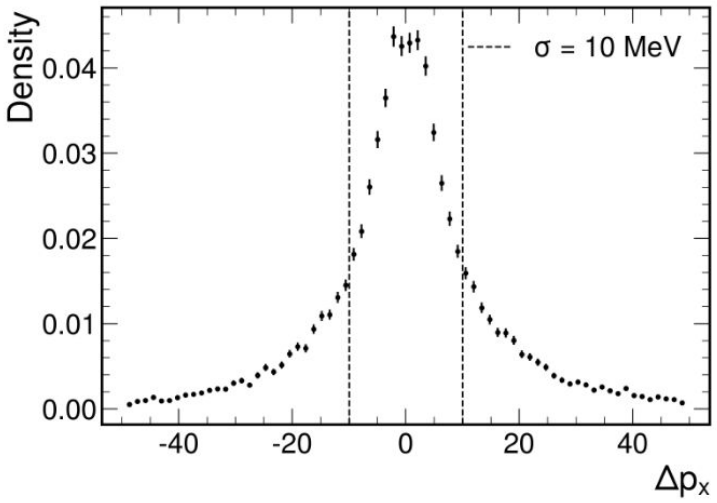
$$C_x \rightarrow f^2 \ C_x$$

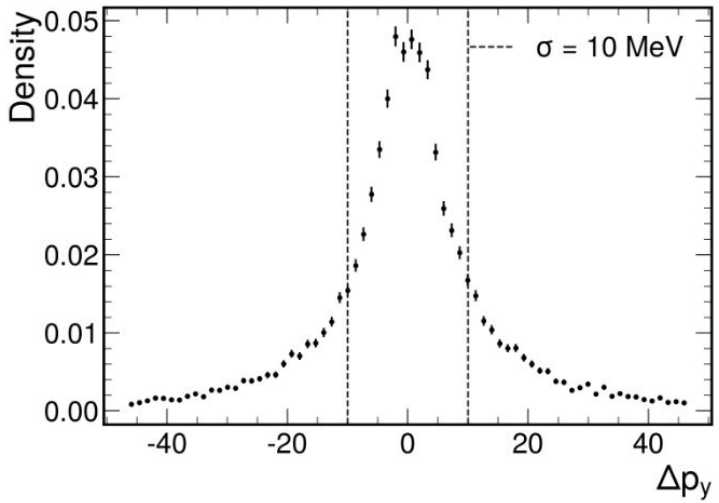
$$C_{xp} \rightarrow f \ C_{xp}$$

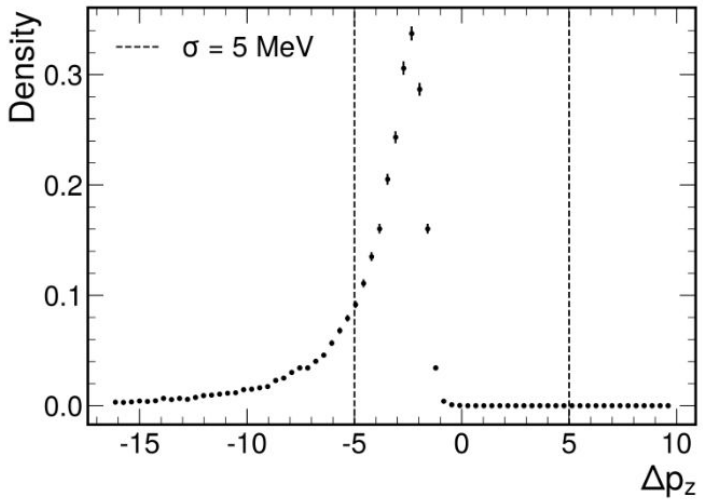
- If using the beam spot, C_x is the beam spot width from the conditions database.

Nuclear Recoil

- The unmeasured nuclear recoil smears the transverse direction.
- Impact is estimated using generator-level truth information from the e+e- pair and recoil electron in a 100 MeV A' sample.
- Estimate ~10 MeV in x, y.

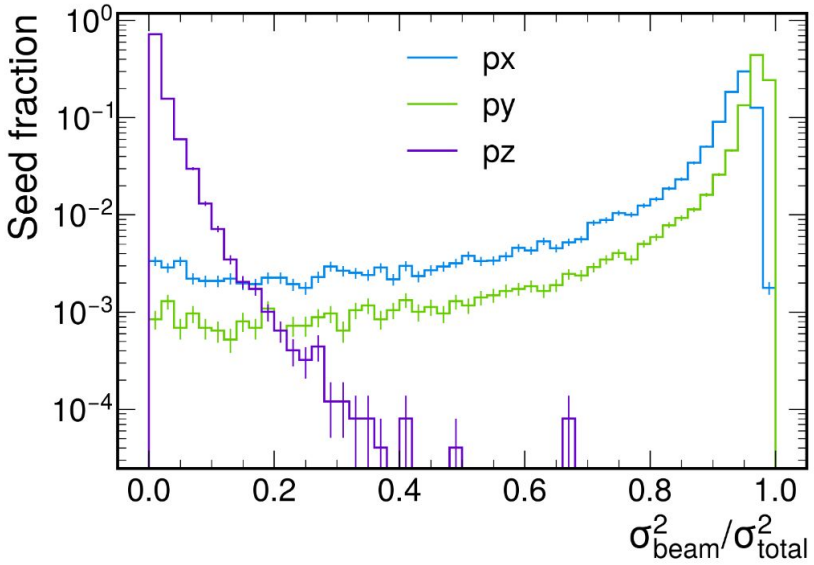


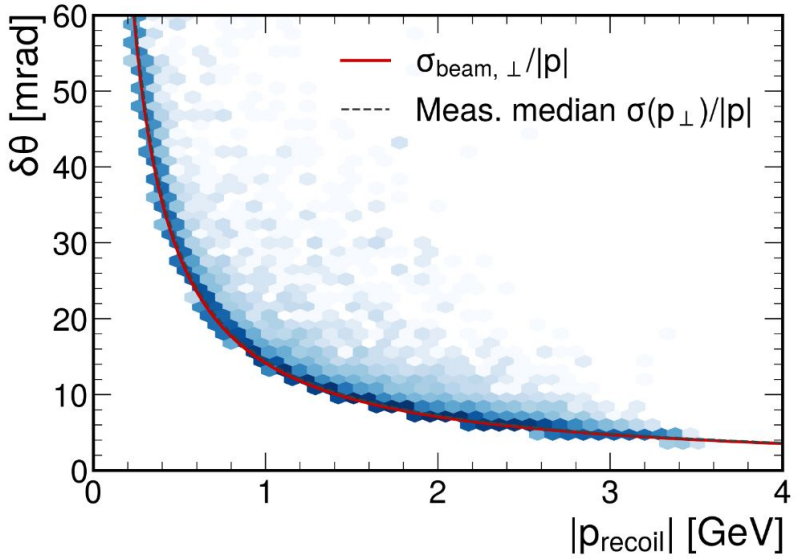




Nuclear Recoil

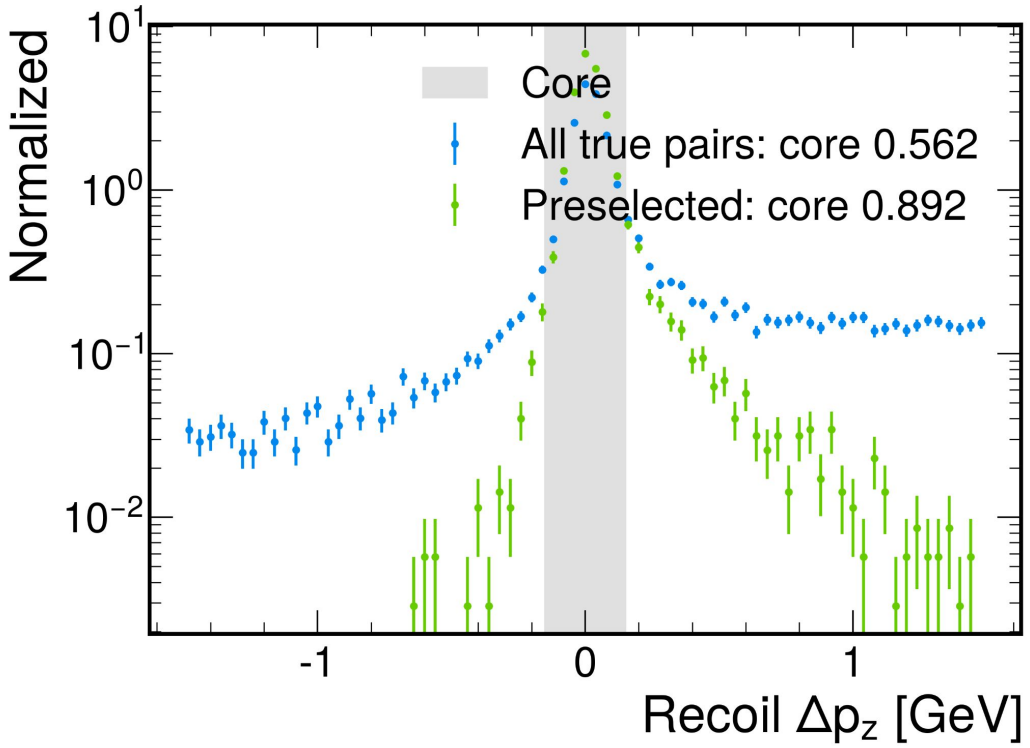
- The unmeasured nuclear recoil smears the transverse direction.
- Impact is estimated using generator-level truth information from the e+e- pair and recoil electron in a 100 MeV A' sample.
- Estimate ~10 MeV in x, y.
- This term dominates the uncertainty in the transverse direction, particularly for low-momentum seeds.





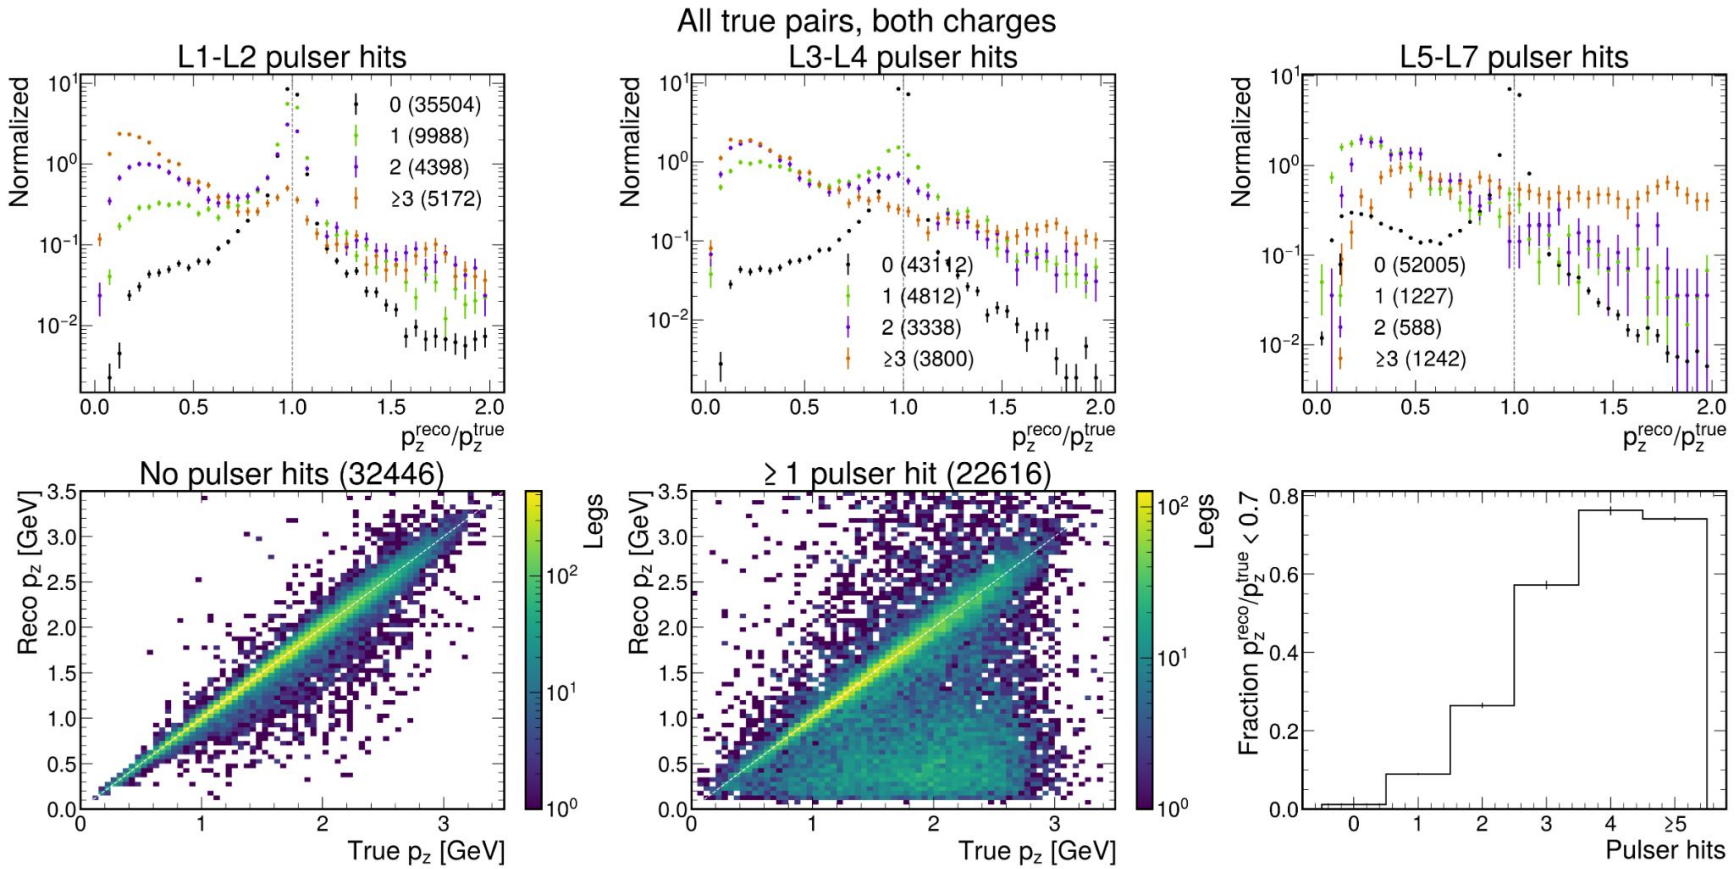
Momentum Mismeasurement

- Previous updates on this work noted a significant fraction of V0s mismeasured the momentum, causing the recoil seed momentum to be overestimated.
- This was found to arise primarily from pulser hit contamination in the outer layers.
- Preselection cleans this up significantly, significantly reducing number of pulser hits in outer layers.



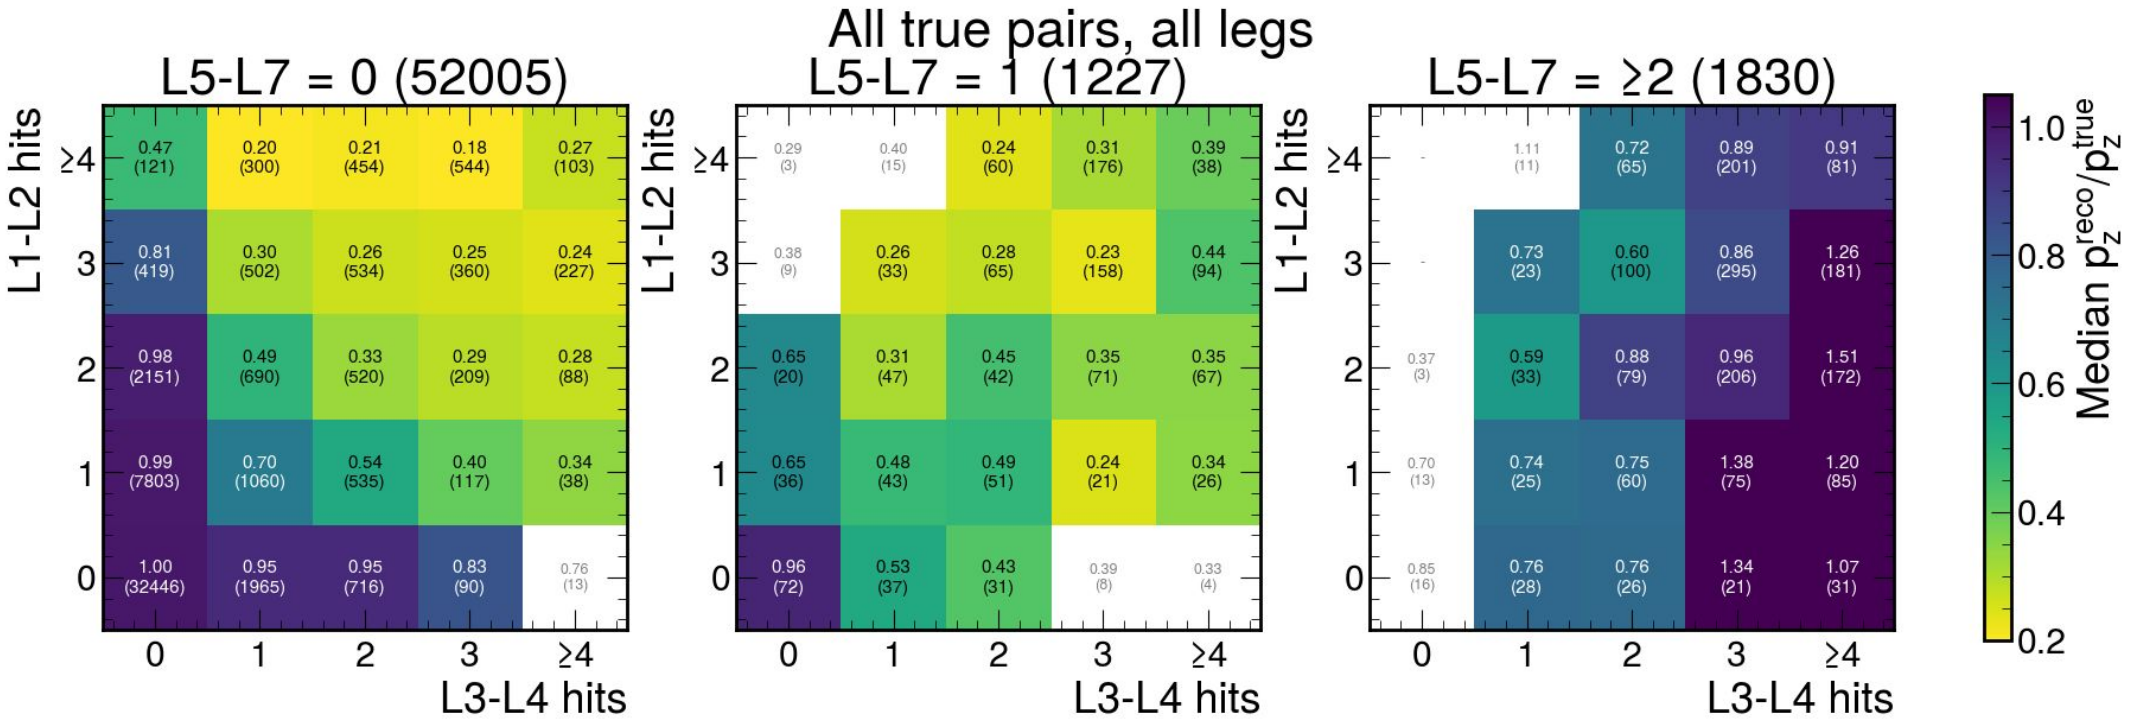
Pulser Hit Contamination

- When looking at the momentum of the V0 legs binned in the number and location of pulser hits, we see that incorrect hits in the outer layers, and combinations of inner/outer incorrect hits, significantly degrade the momentum measurement.



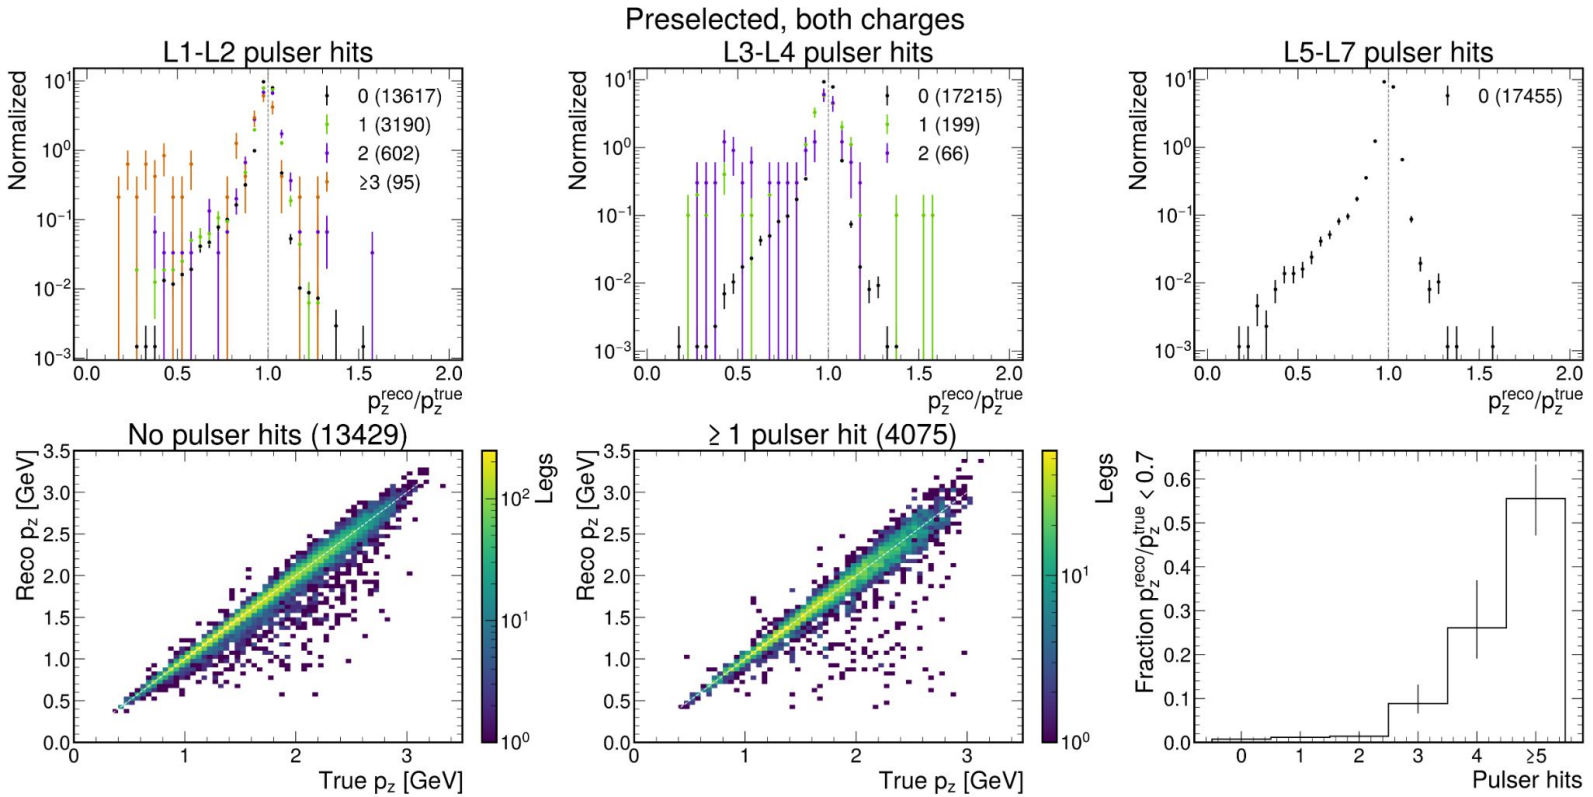
Pulser Hit Contamination

- When looking at the momentum of the V0 legs binned in the number and location of pulser hits, we see that incorrect hits in the outer layers, and combinations of inner/outer incorrect hits, significantly degrade the momentum measurement.



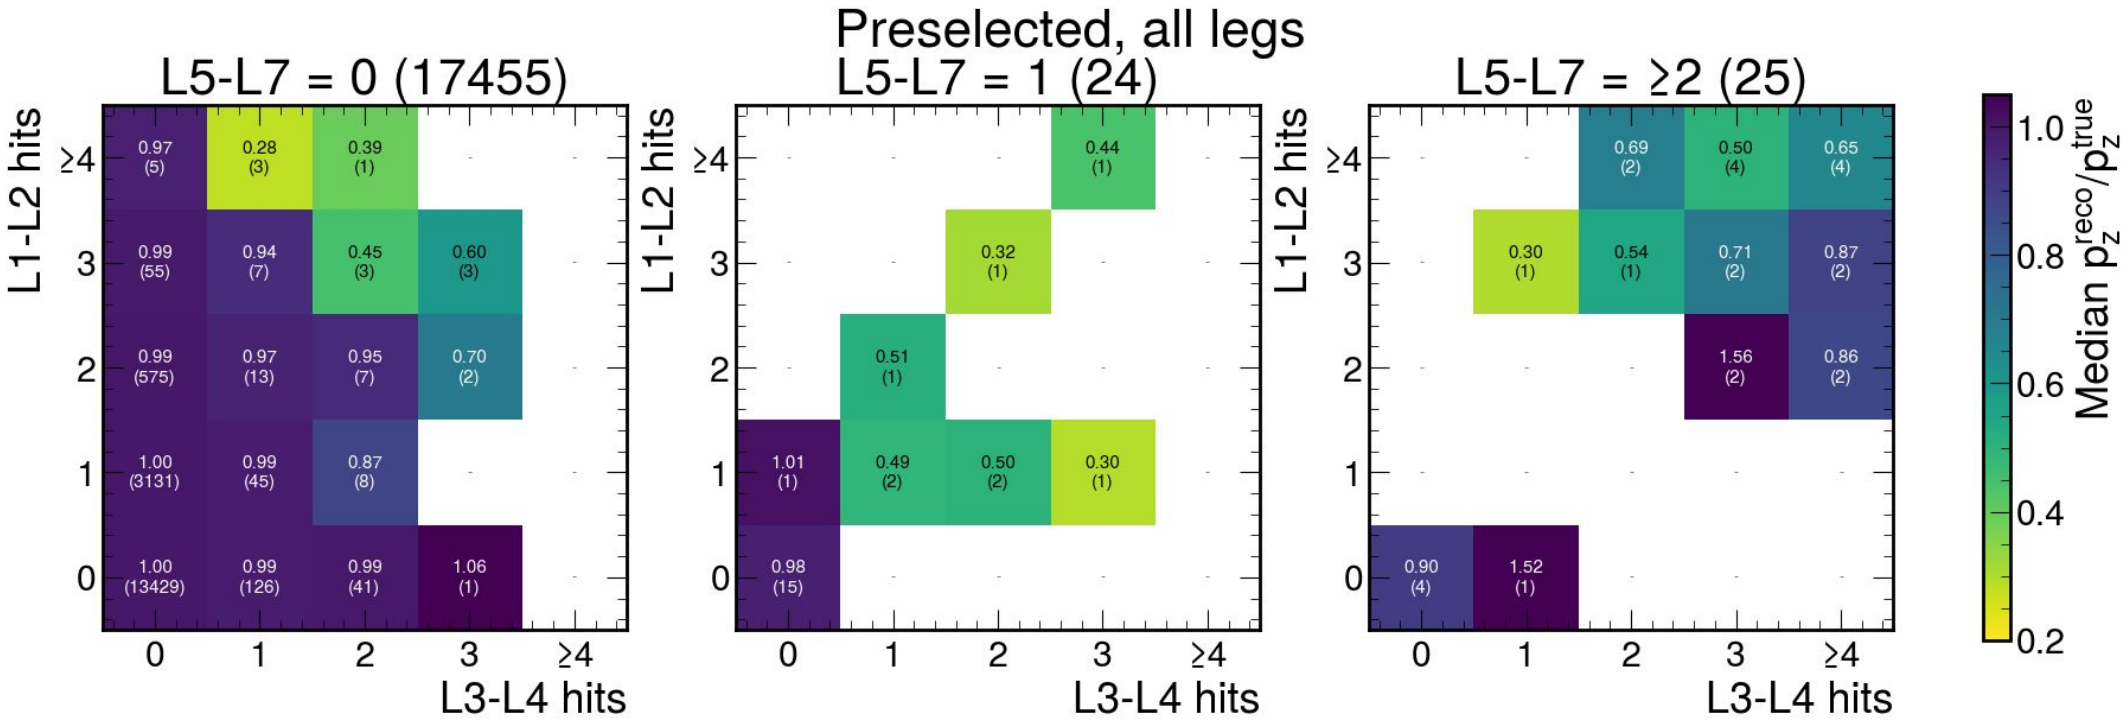
Pulser Hit Contamination

- When looking at the momentum of the V0 legs binned in the number and location of pulser hits, we see that incorrect hits in the outer layers, and combinations of inner/outer incorrect hits, significantly degrade the momentum measurement.
- Preselection improves the momentum resolution, even for similar numbers of incorrect hits.



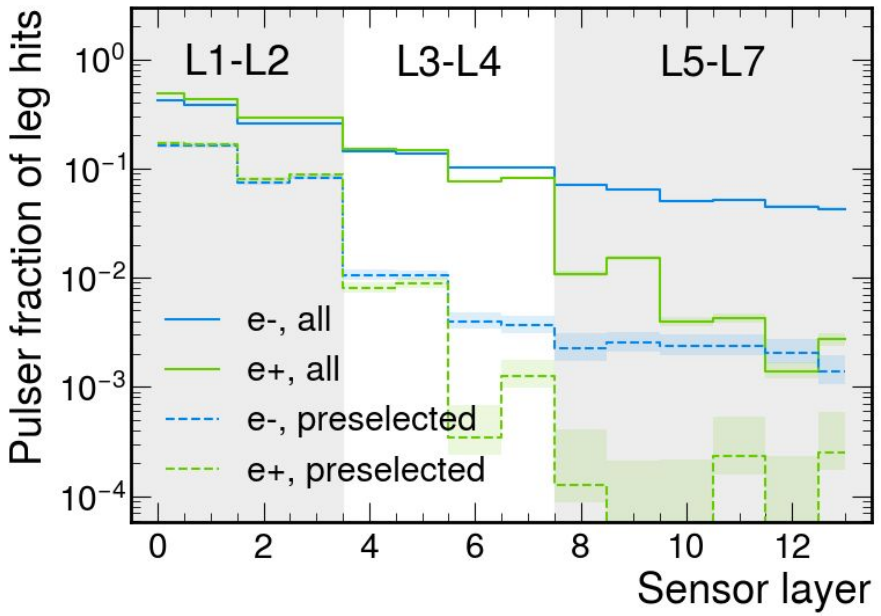
Pulser Hit Contamination

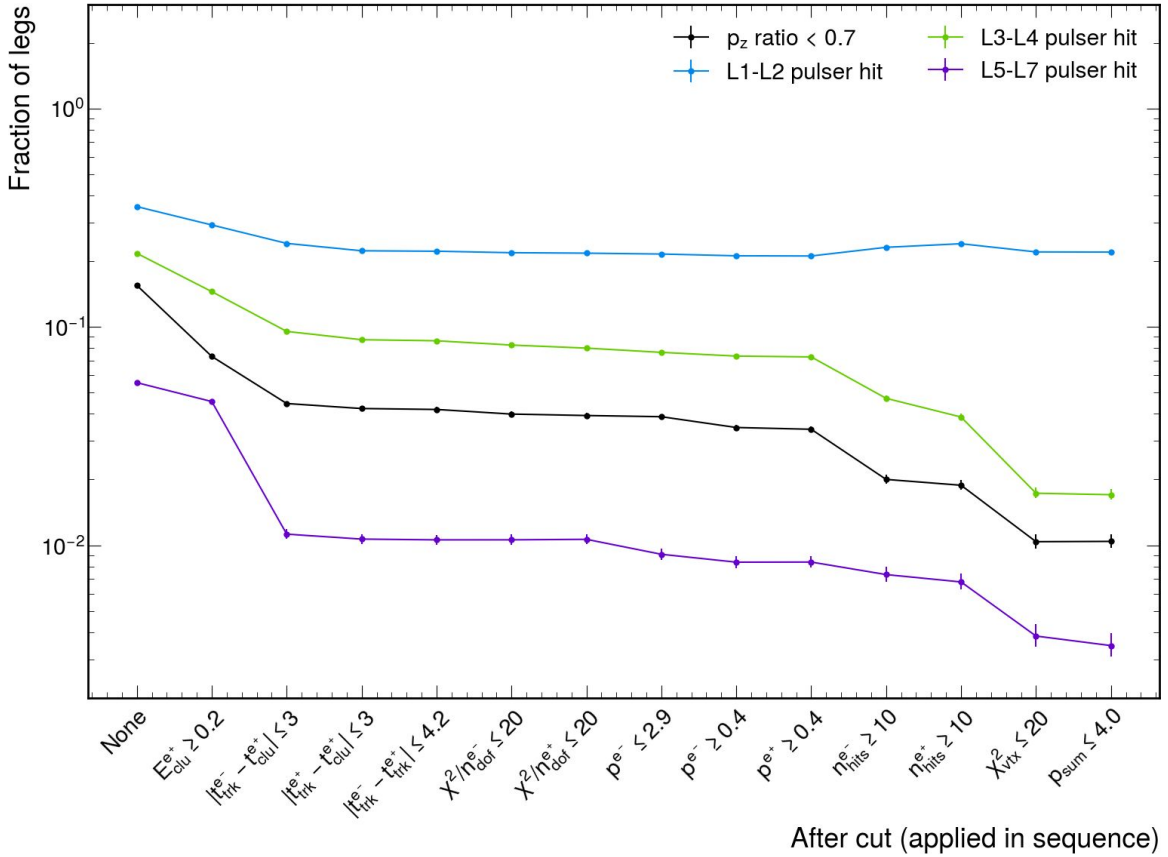
- When looking at the momentum of the V0 legs binned in the number and location of pulser hits, we see that incorrect hits in the outer layers, and combinations of inner/outer incorrect hits, significantly degrade the momentum measurement.
- Preselection improves the momentum resolution, even for similar numbers of incorrect hits.



Pulser Hit Contamination

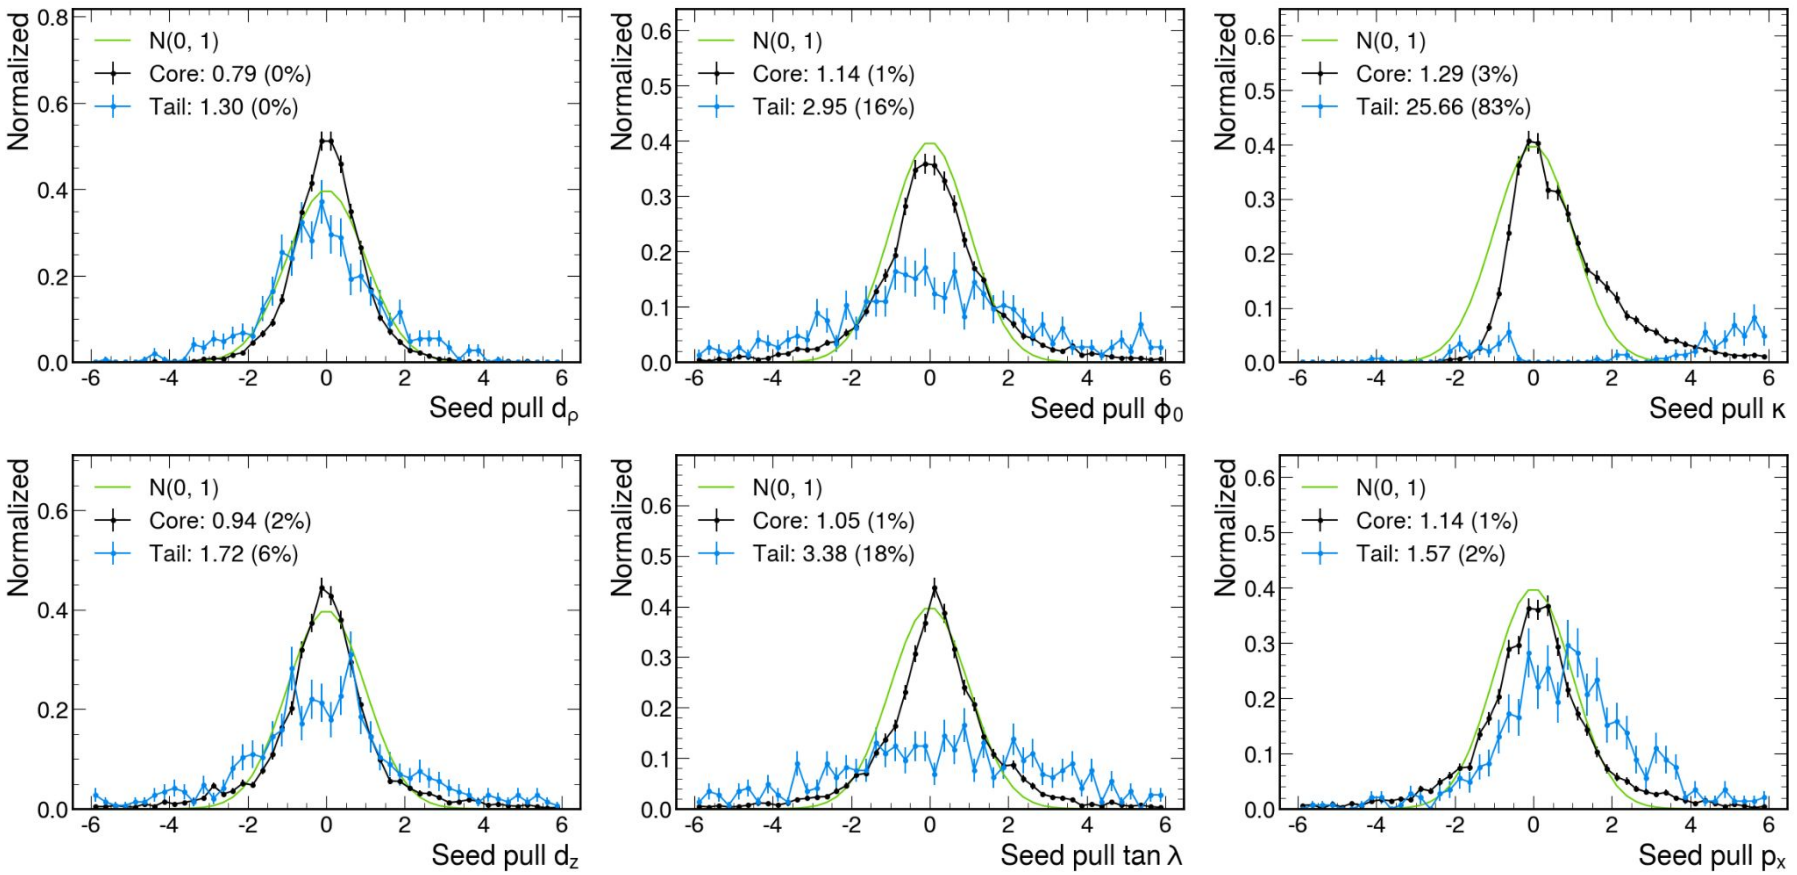
- The cuts most responsible for reducing the pulser hit fraction include timing cuts, χ^2 cuts, cuts on the number of hits.





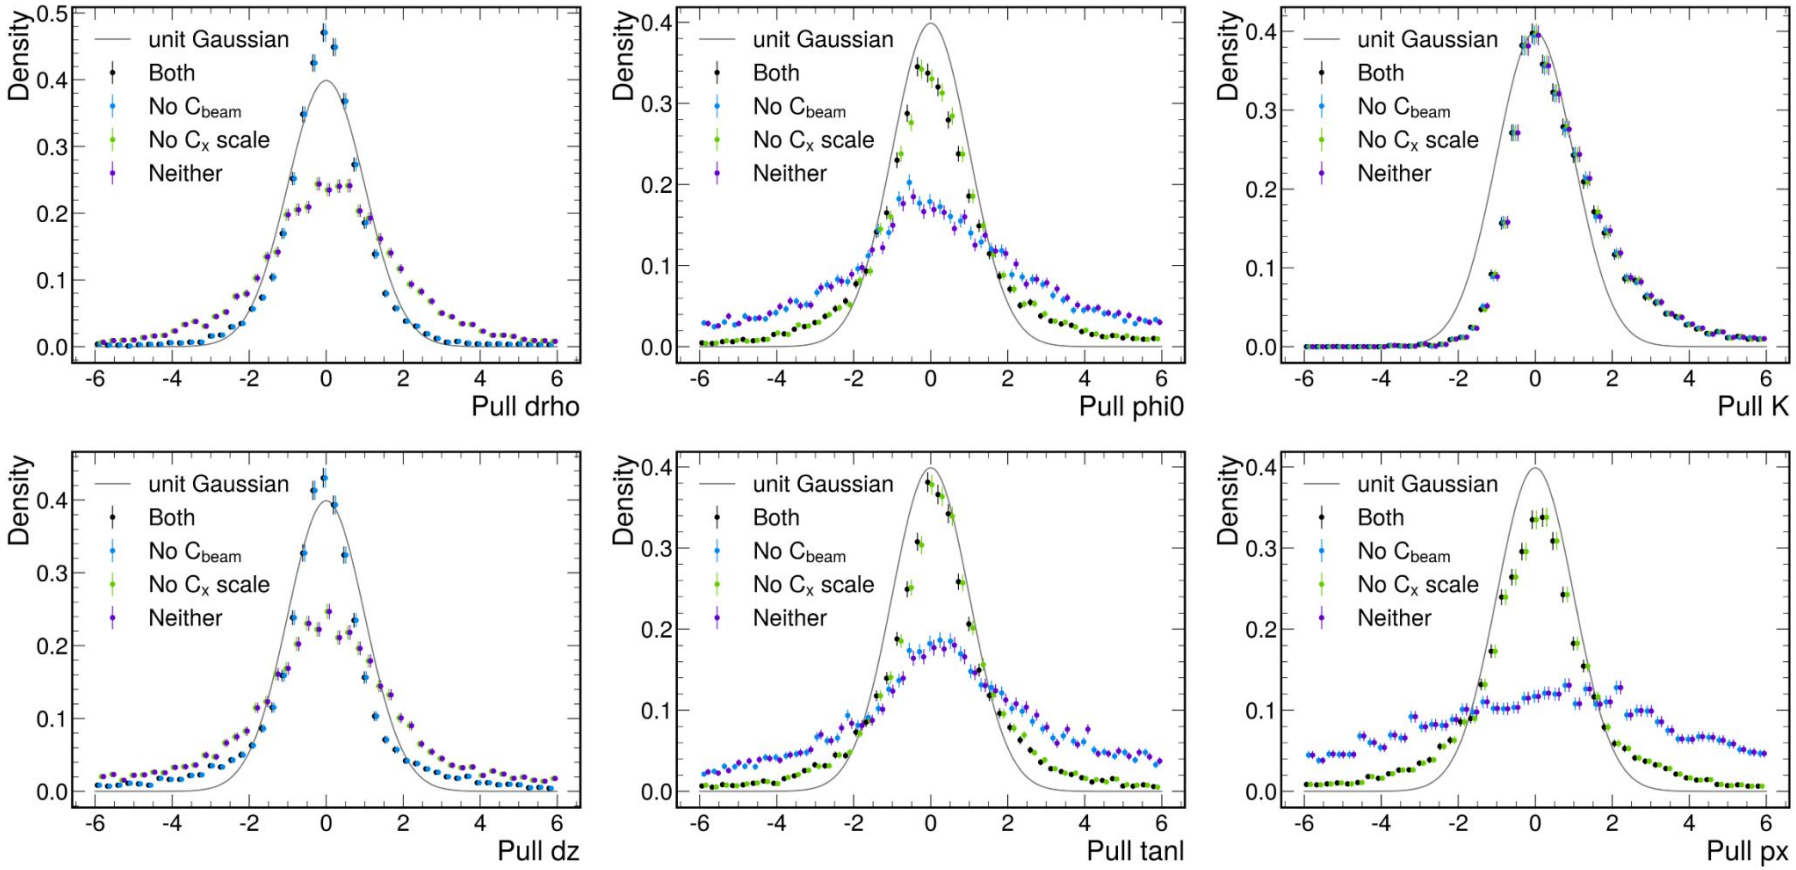
Recoil Seeding

- Seeds in core of pz residual distribution show good agreement with standard Gaussian in helix parameter pulls (on 100 MeV A').



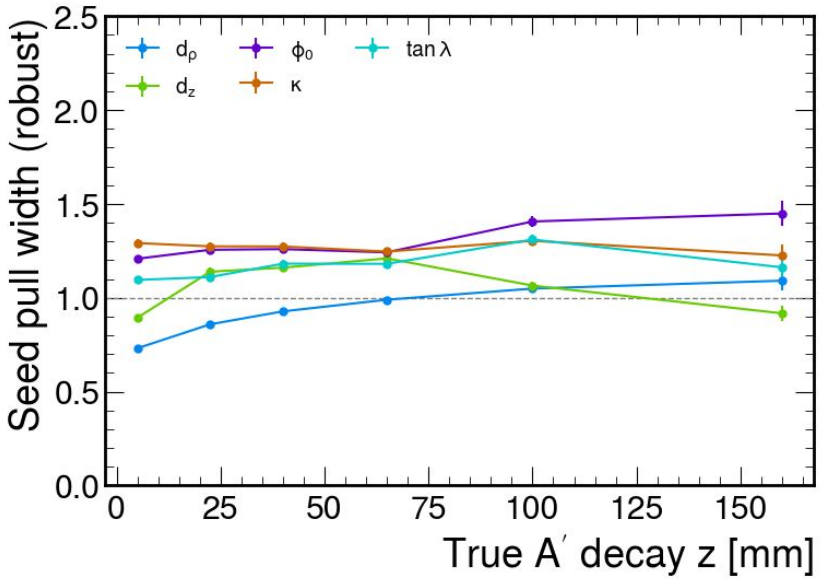
Recoil Seeding

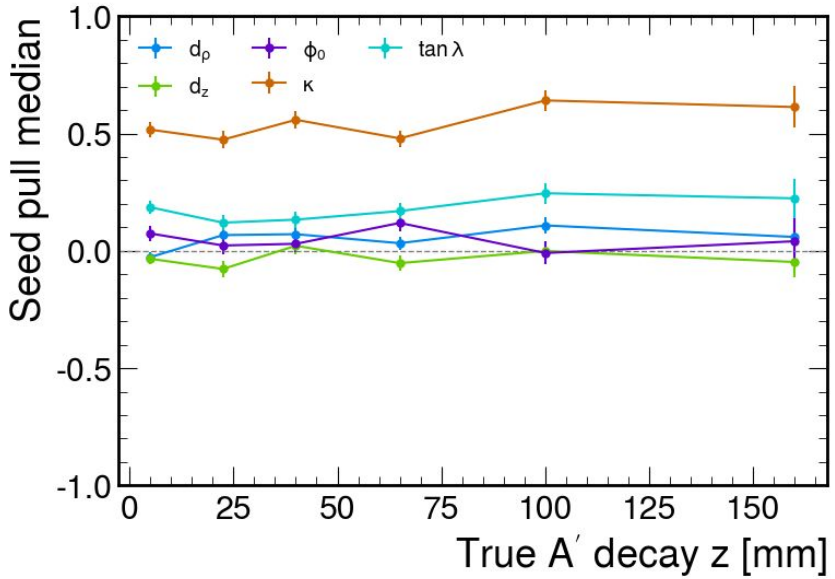
- Seeds in core of pz residual distribution show good agreement with standard Gaussian in helix parameter pulls, after tuning of position covariance (on 100 MeV A').



Recoil Seeding

- Pulls are stable with displacement.





Hit Finding

- Seeds are Runge-Kutta propagated to the first layer of the tracker to determine starting state for hit finding. Hits that are already on reconstructed tracks are vetoed.
- At each layer, the hits are subject to the same set of cuts as the standard Kalman filter: timing, charge, split-strip side. Hits surviving cuts are sorted by their χ^2 increment:

$$C_{\text{in}} = \begin{bmatrix} C_x & C_{xp} \\ C_{xp}^{\top} & C_p \end{bmatrix}$$

$$R_{\text{extrap}} = H^T C H \qquad \text{track covariance's contribution}$$

$$R = \sigma^2_{\text{hit}} + R_{\text{extrap}} \qquad \text{total predicted-residual variance}$$

$$\chi^2_{\text{inc}} = r^2 / R \qquad r = v_{\text{hit}} - \text{aP.mPred}$$

- A hit is added to a candidate track if the χ^2 increment is < maxHitChi2Inc (16).
- Additional candidate cuts:

$$n_{\text{ChargedHoles}} > \texttt{maxHoles} \qquad \rightarrow \text{dead}$$

$$n_{\text{Hits}} + \texttt{layersRemaining}(\text{volume}, \text{lyr} + 1) < \texttt{minHits} \qquad \rightarrow \text{dead, cannot reach minHits}$$

$$n_{\text{Hits}} > 0 \text{ and } \chi^2 / n_{\text{Hits}} > \texttt{maxChi2PerHit} \qquad \rightarrow \text{dead}$$

- Surviving candidates are sorted according to:

$$\text{score} = \chi^2 - \texttt{hitRewardChi2} \cdot n_{\text{Hits}} \qquad \text{lower is better}$$

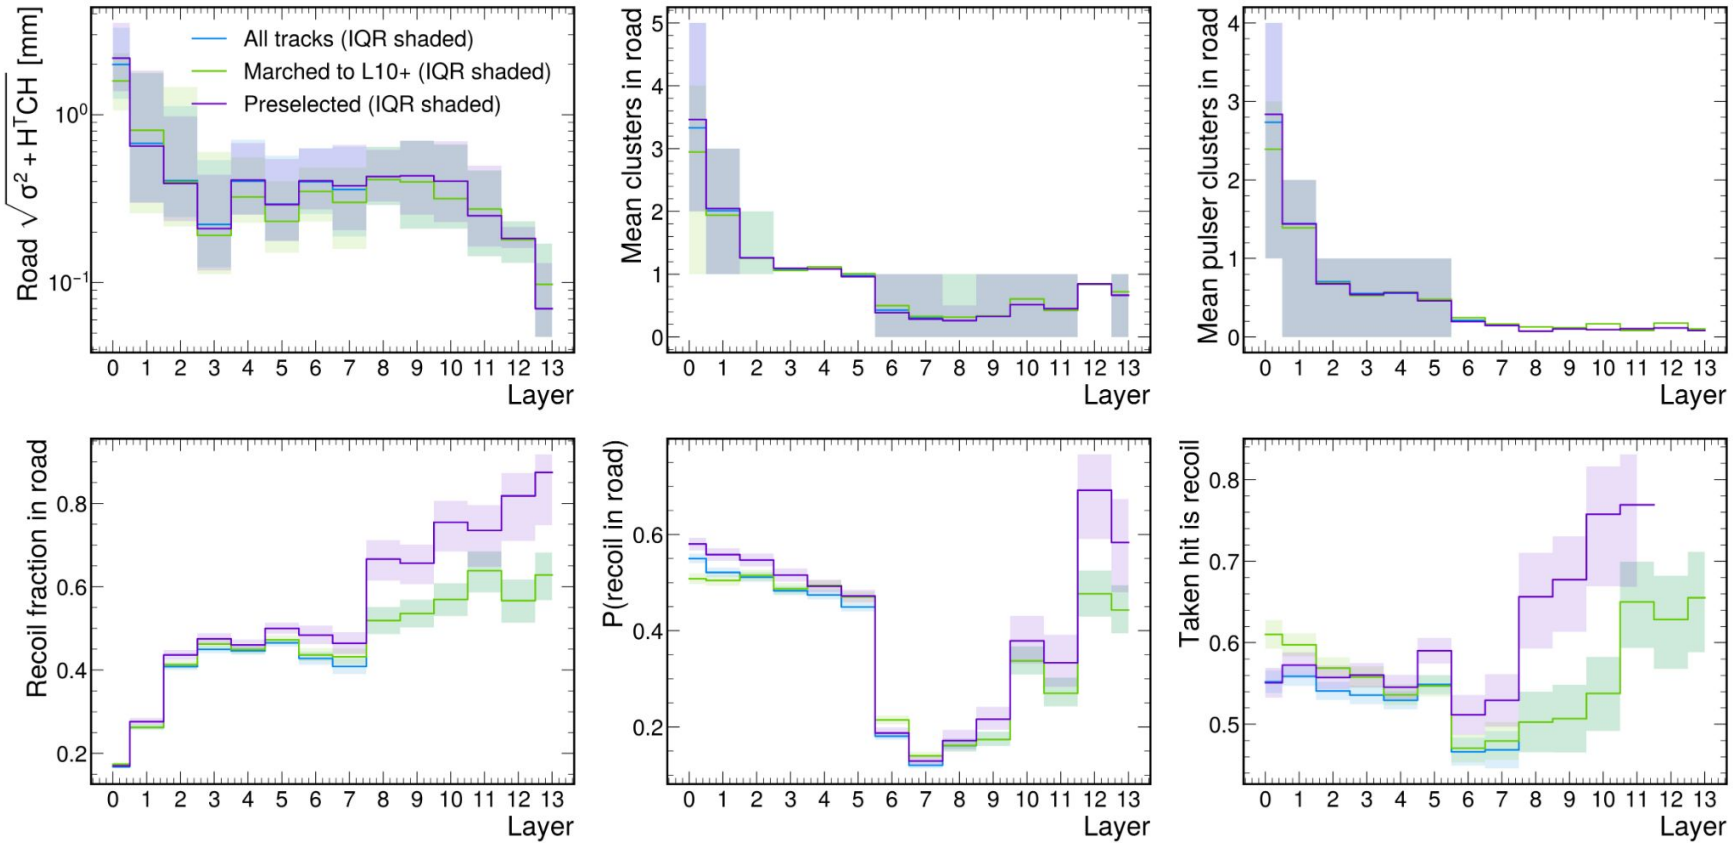
Hit Finding

Parameter	Default	Description
`minHits`	4	Minimum hits required to keep a road.
`minAxialHits`	2	Minimum axial hits required.
`minStereoHits`	2	Minimum stereo hits required.
`maxHoles`	3	Budget on **interior** holes only
`maxTrailingHoles`	4	Termination criterion on number of holes at the end of a candidate
`maxTrailingHolesKept`	3	Maximum number of holes at the end of a candidate
`maxFirstLayer`	13	Latest layer the first hit may sit on.
`maxLayer`	13	Last layer marched.
`nBranch`	3	Candidates kept per layer, best χ^2 first
`maxCandidates`	20	Total allowed candidates after per-layer sorting
`maxHitChi2Inc`	16.0	A hit is a candidate if $\chi^2_{\text{inc}} < 16$, i.e. a 4σ window
`maxChi2PerHit`	8.0	Running quality cut, χ^2/n_{Hits}
`minSeedEnergy`	0.6	Cluster charge in keV. Production uses 1.3
`hitRewardChi2`	4.0	For candidate ranking only: score $= \chi^2 -$ `hitRewardChi2` $\cdot\, n_{\text{Hits}}$
`hitTimeWindow[14]`	12.0 ns	Per-layer $\lvert t_{\text{hit}} - t_{\text{pair}} \rvert$ window
`searchBothVolumes`	true	Search top *and* bottom rather than selecting a single volume
`concurrentHoleBranch`	true	At each layer, propagate a candidate without a hit.
`beamSpotConstraint`	false	Apply the beam spot as a 3D Kalman measurement.

- **Note:** these are default values, not the values implemented in the optimal selection.

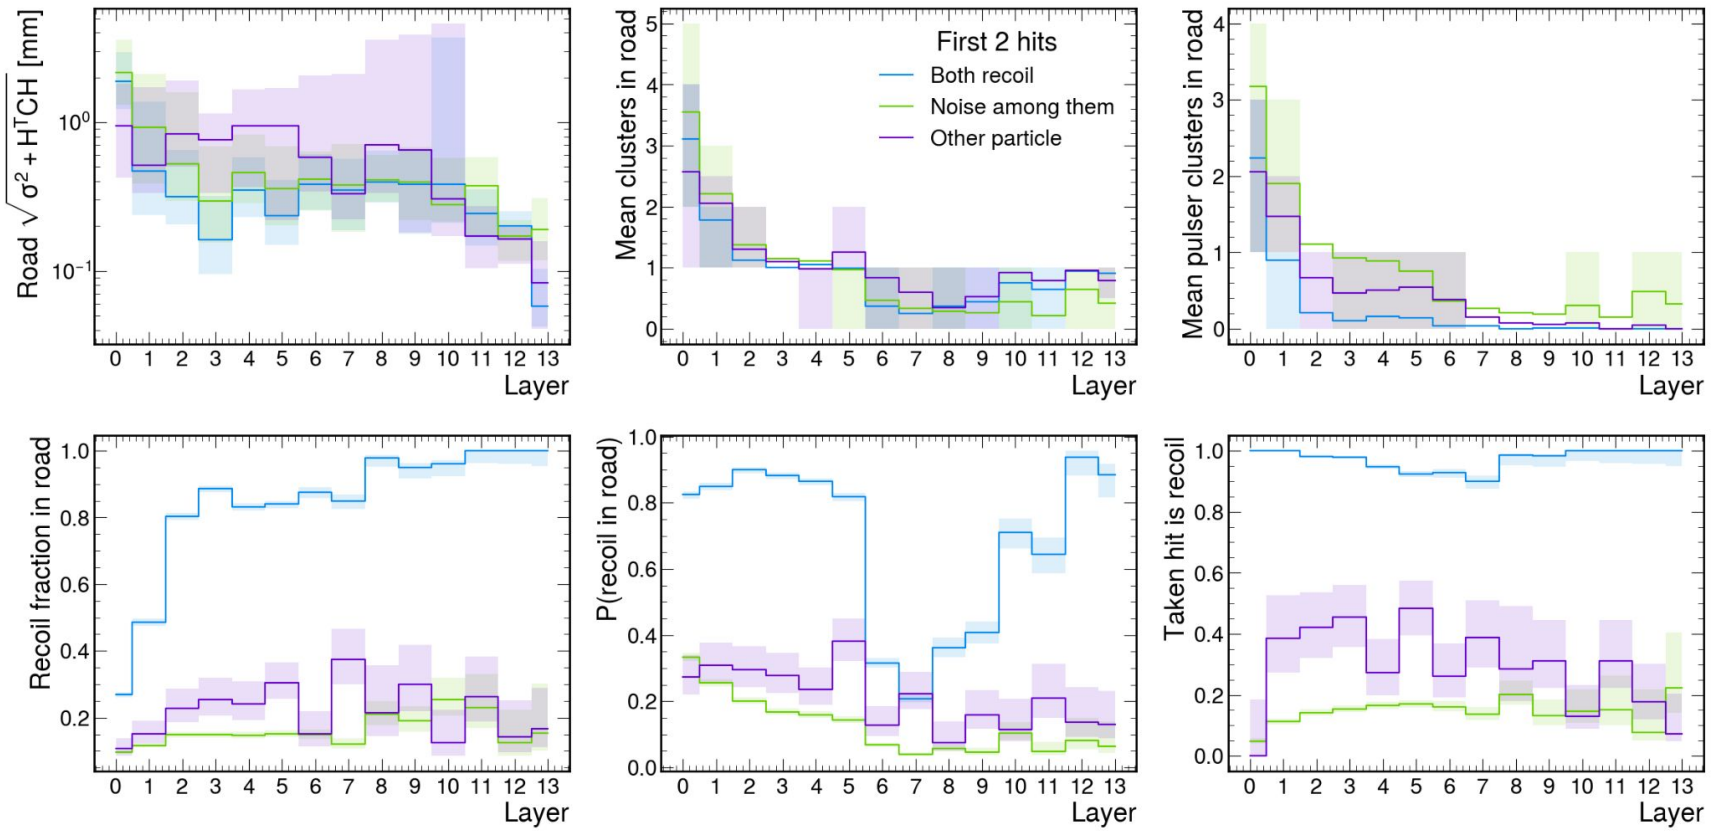
Road Diagnostics

- Looking at the road configuration of the best candidate track.
- Road 1σ width is about 2 mm at L0, 200 um at L4.
- In early layers, roads consist of multiple clusters, dominated by pulser hits—fake tracks consisting of pulser hits are a significant fraction of candidates.
- Preselected tracks are more likely to find the recoil.



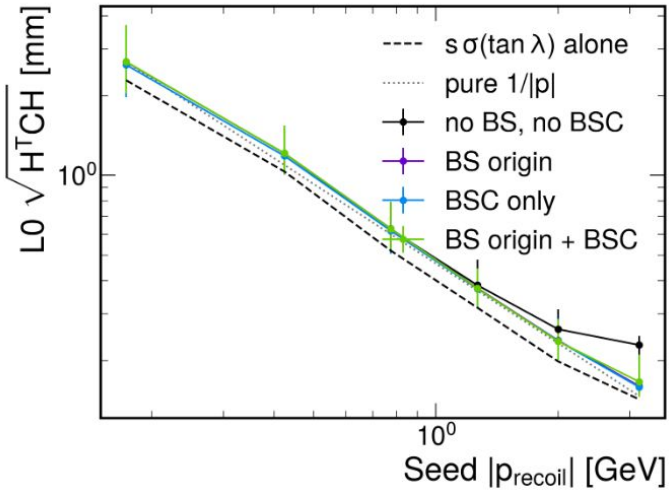
Road Diagnostics

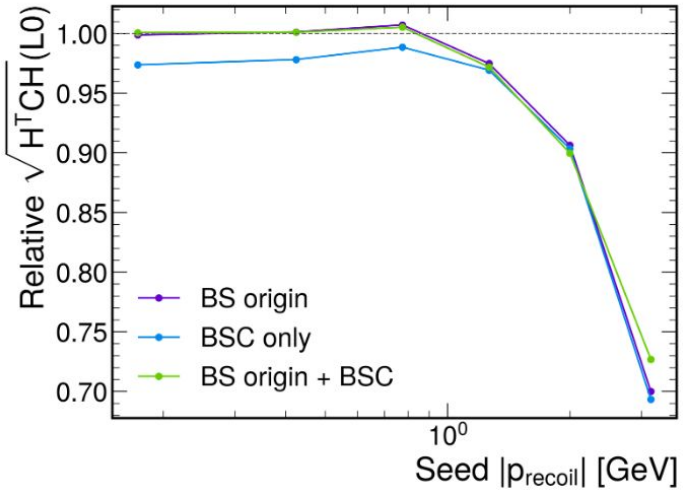
- A road that contains a noise hit in one if it's first two layers does not recover.
- Tracks whose first two hits are both recoil keep a high purity throughout.

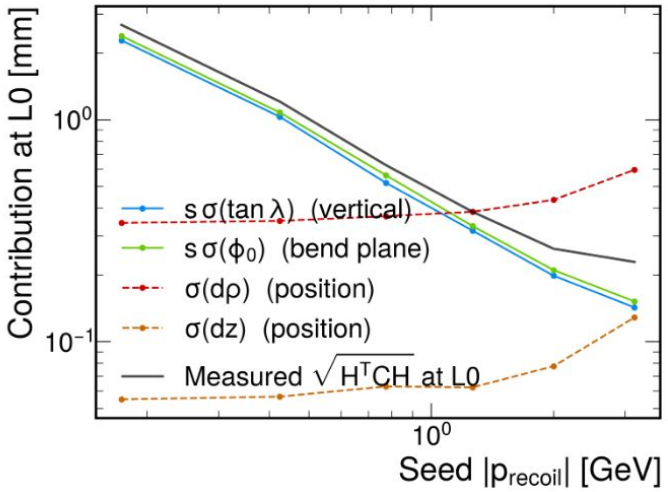


Road Diagnostics (with BSC)

- Beam spot constraint help only among higher momentum seeds, where directional uncertainty from the nuclear recoil is subdominant to position-uncertainty terms.

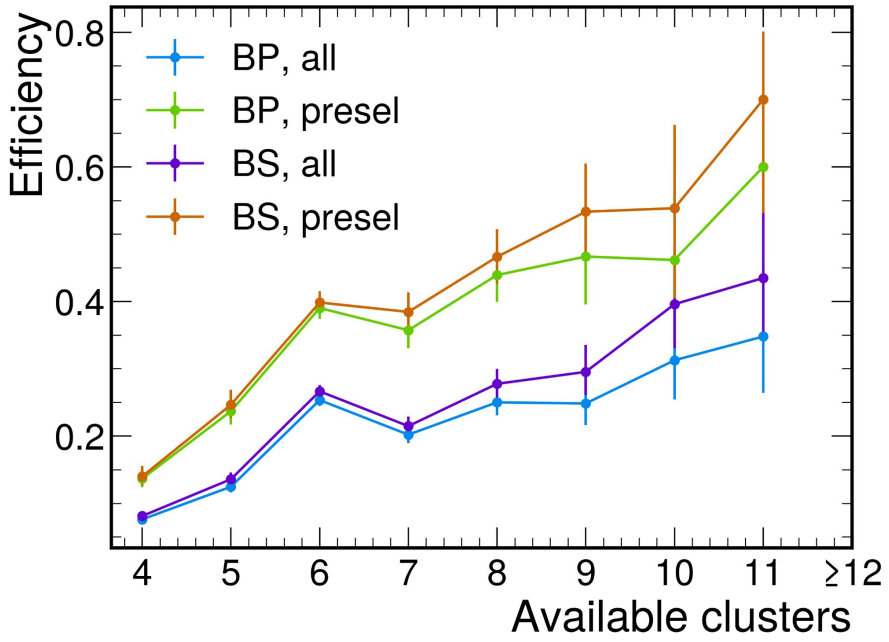


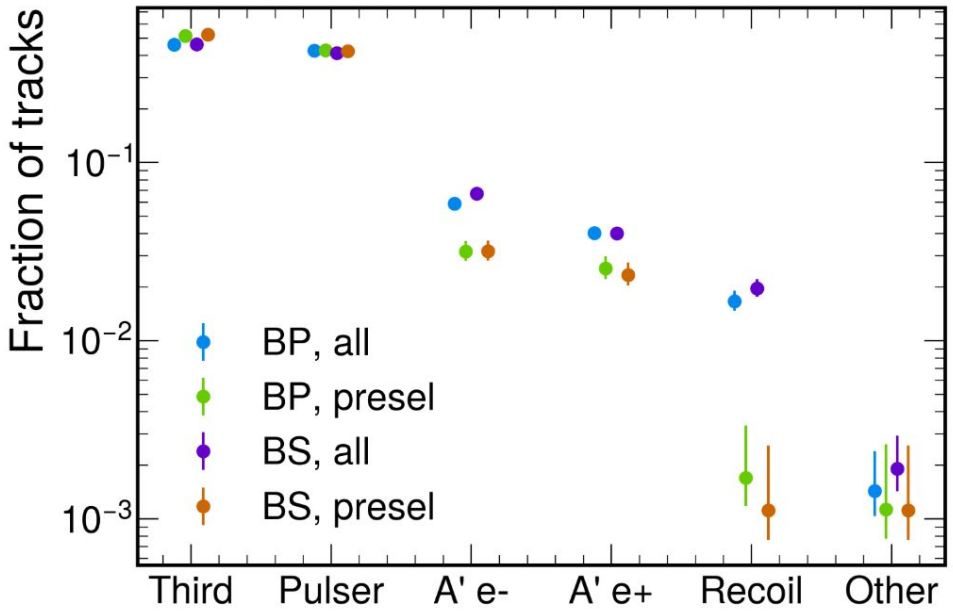




Finder Performance (minHits = 4)

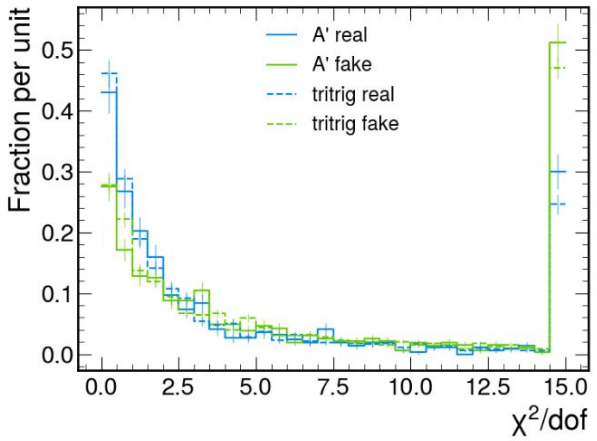
- A real track has a purity (fraction of a track's hits that carry the true recoil) ≥ 0.5. A fake track has purity < 0.5.
- The efficiency is the fraction of events where the recoil left ≥ minHits usable clusters not already on a reconstructed track, where a real track is also reconstructed.
- Efficiencies are evaluated on a single A' sample.
- Efficiencies are low, and fake pulser tracks are found almost as often as true third tracks.

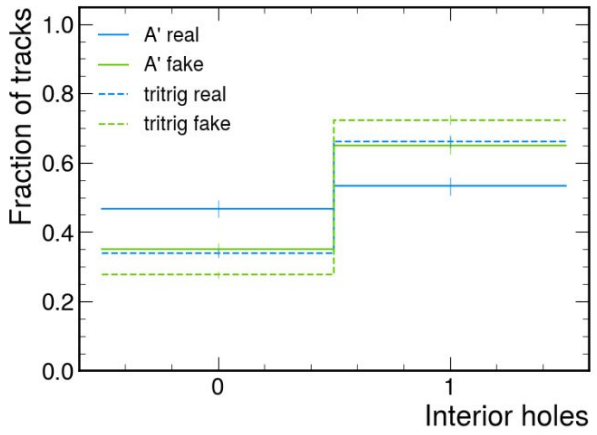


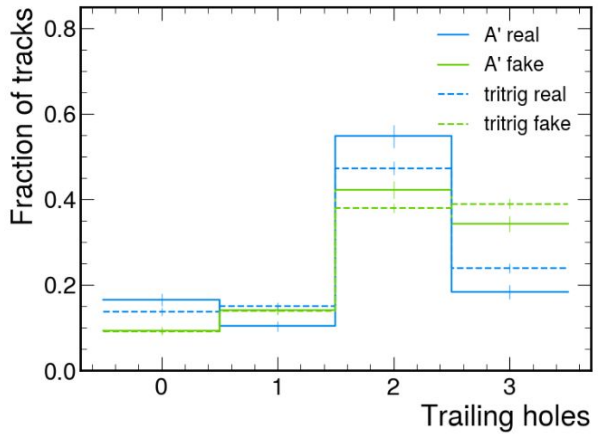


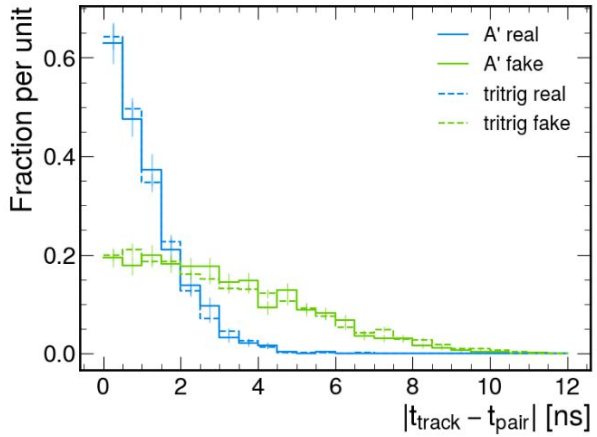
Real vs. Fake Tracks

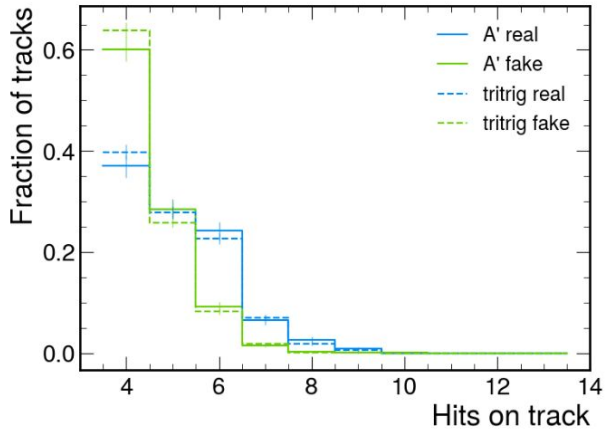
- Fake tracks show worse fit quality, timing, more holes, etc. than real tracks, but are not easily distinguishable in a cut-based way.





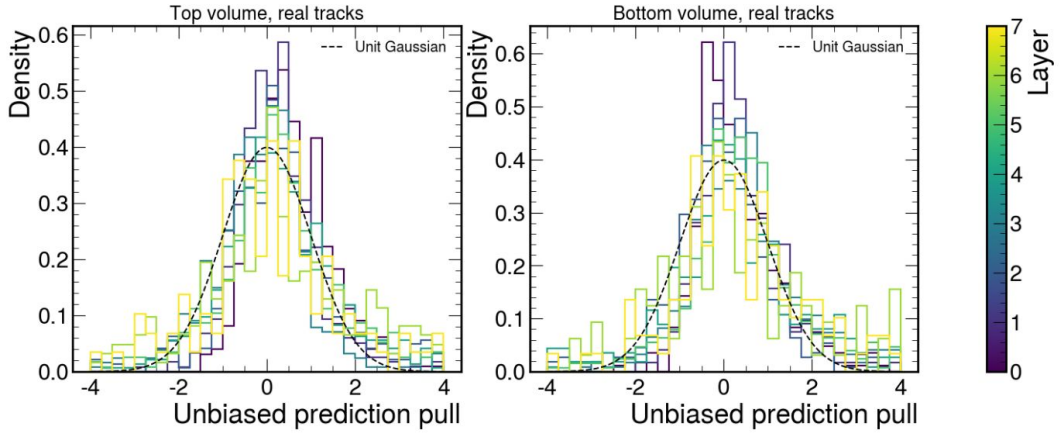


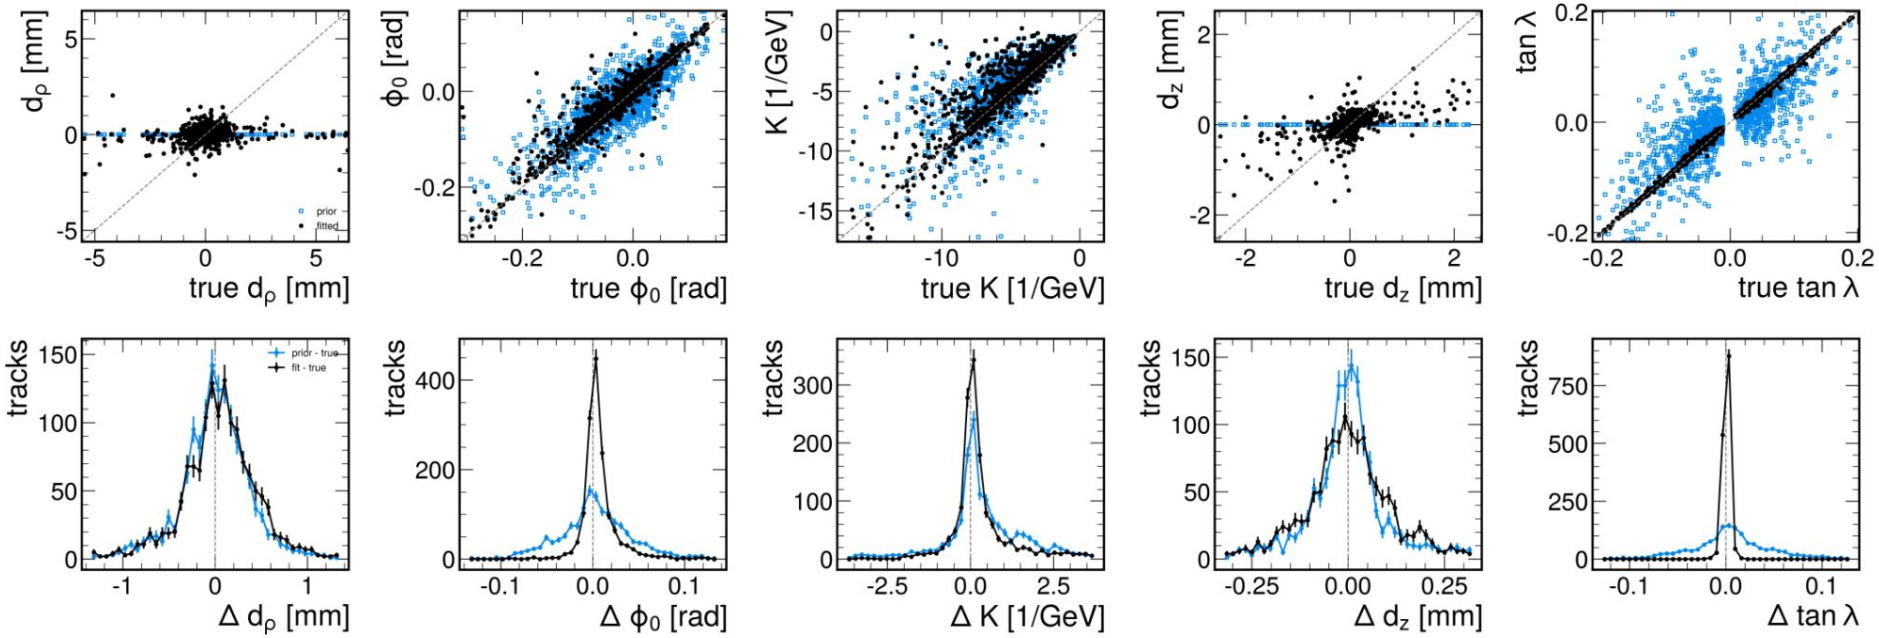




Fitting Validation

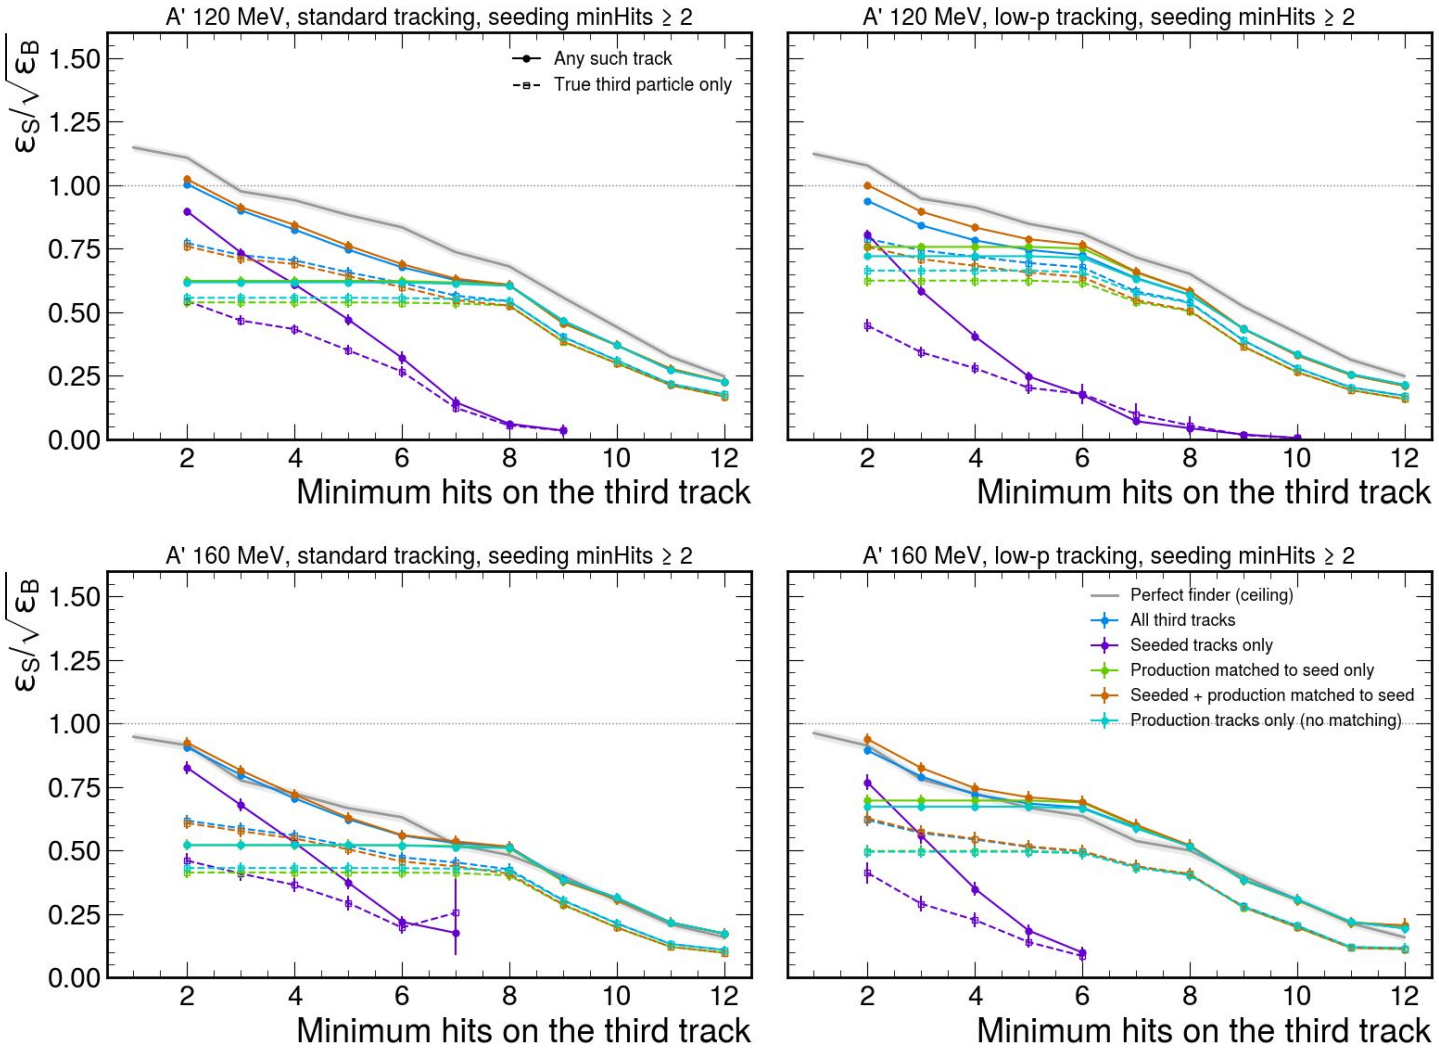
- Among real tracks, unbiased pulls and refit parameters behave as expected.
- Additional hits on short tracks constrain recoil direction, momentum shows little change.





Signal/Background Performance

- Ran recoil tracking on signal and tritrig MC, with two-hit minimum configuration.
- With real tracks alone, significance drops.
- Fake tracks are a significant contribution, barely break even.
- With low-p tracking, existing production contribution increases relative to seed contribution.
- Potential improvement drops at higher masses.



Summary

- Finding real tracks is challenging for short track segments in the early layers.
  - Unlikely to be helpful for a cut-based requirement on the presence of a track.
  - Limitations: transverse direction smearing from nuclear recoil, pulser contamination in early layers.
- Might have limited applications as an additional discriminatory variable, but is correlated with other kinematic variables.
- Potentially has some usefulness is selected the correct electron/positron pair among surviving preselected vertices, but χ^2 and target-consistency cuts already help with this.