CONTINUOUS NORMALIZING FLOWS · NEURAL ODEs · XENON TPCs

Physics-informed continuous normalizing flows to learn the electric field within a TPC

Peter Gaemers, Ivy Li, Juehang Qin, Naija Bruckner, Maris Arthurs, Maria Elena Monzani, Christopher Tunnell

SLAC / KIPAC, Stanford, Rice University

arXiv:2511.01897 [physics.ins-det]

https://github.com/RiceAstroparticleLab/fieldflow

NATIONAL ACCELERATOR LABORATORY

MOTIVATION

How a liquid-xenon TPC sees an event

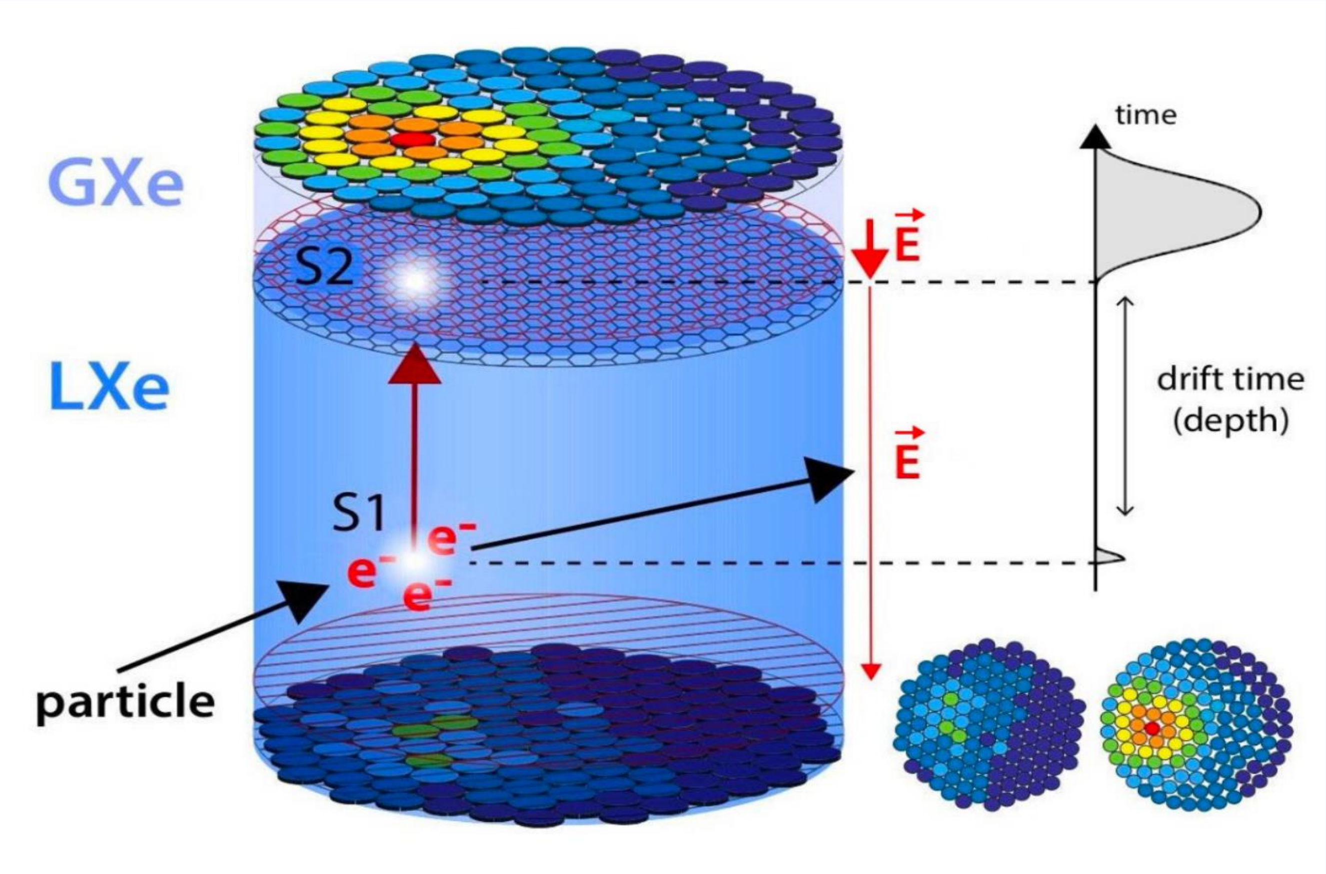

Particle → prompt light (S1) + drifting electrons → amplified S2. Hit pattern gives (x,y); drift time gives depth z.

Position is everything

WIMPs, 0νββ, CEνNS — every rare-event search fiducializes and rejects wall backgrounds on the reconstructed position.

…BUT THERE'S A CATCH

- Charge builds up on the PTFE walls once the detector is sealed.
- It bends the drift field inward → electrons land at the wrong radius, or are lost to the wall.
- You can't measure it in-situ.

MOTIVATION

Real detectors have warped fields

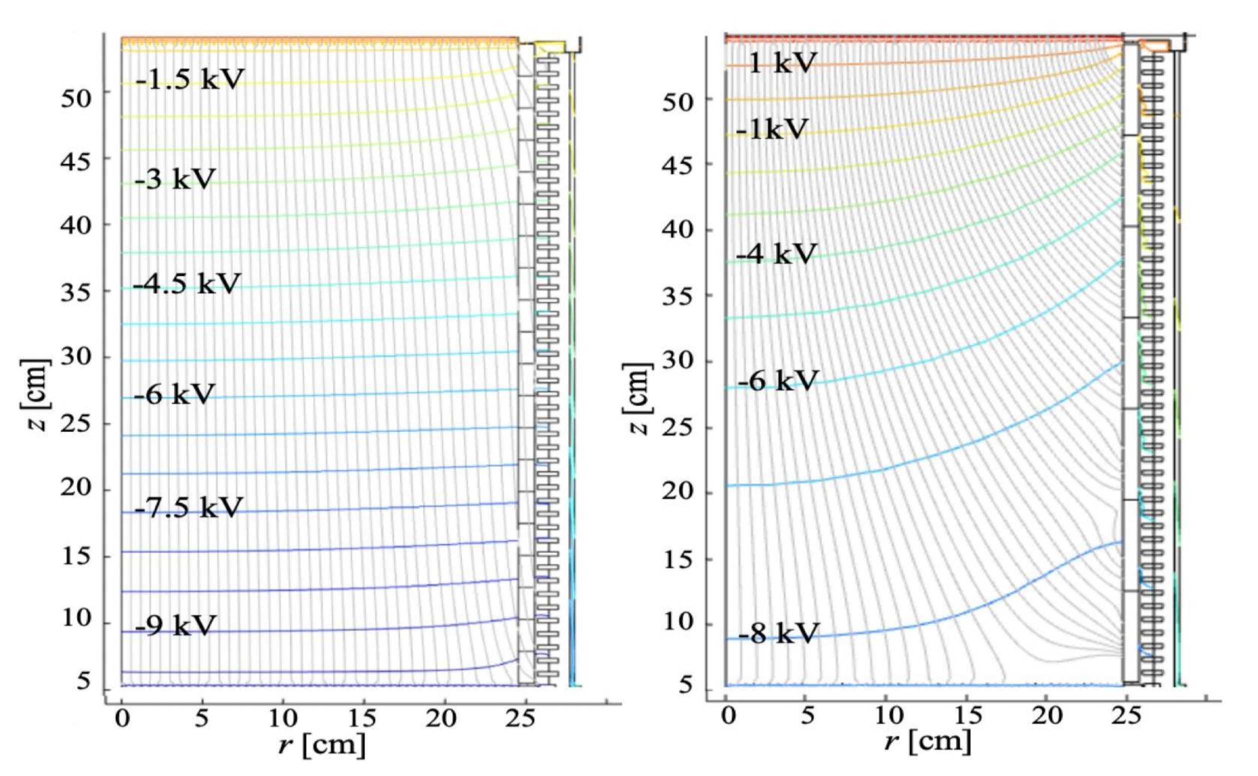

LUX: equipotentials, early run (left) vs. late run (right). The field visibly buckles inward as wall charge accumulates.

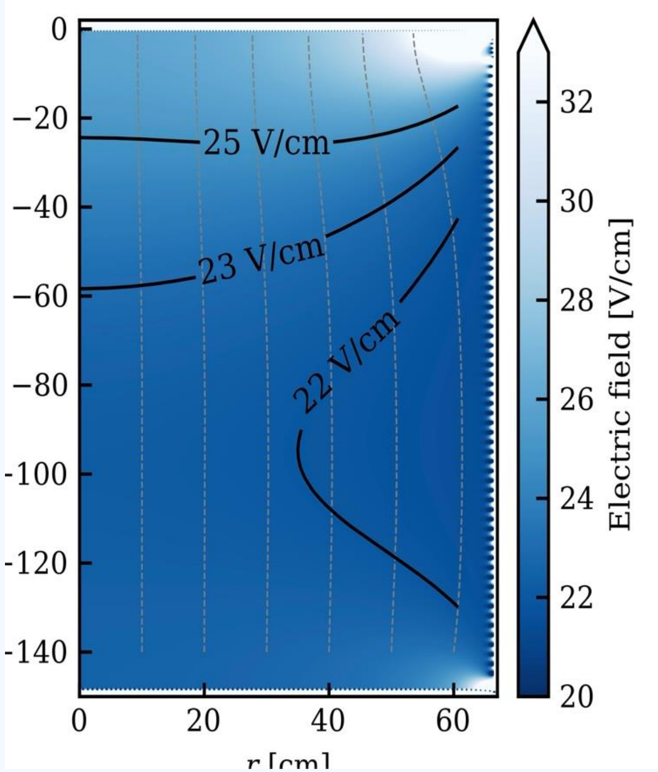

XENONnT: drift-field magnitude. Far from the uniform 23 V/cm you'd want — it sags toward the wall.

Today's fix: histogram-based field-distortion correction (FDC)

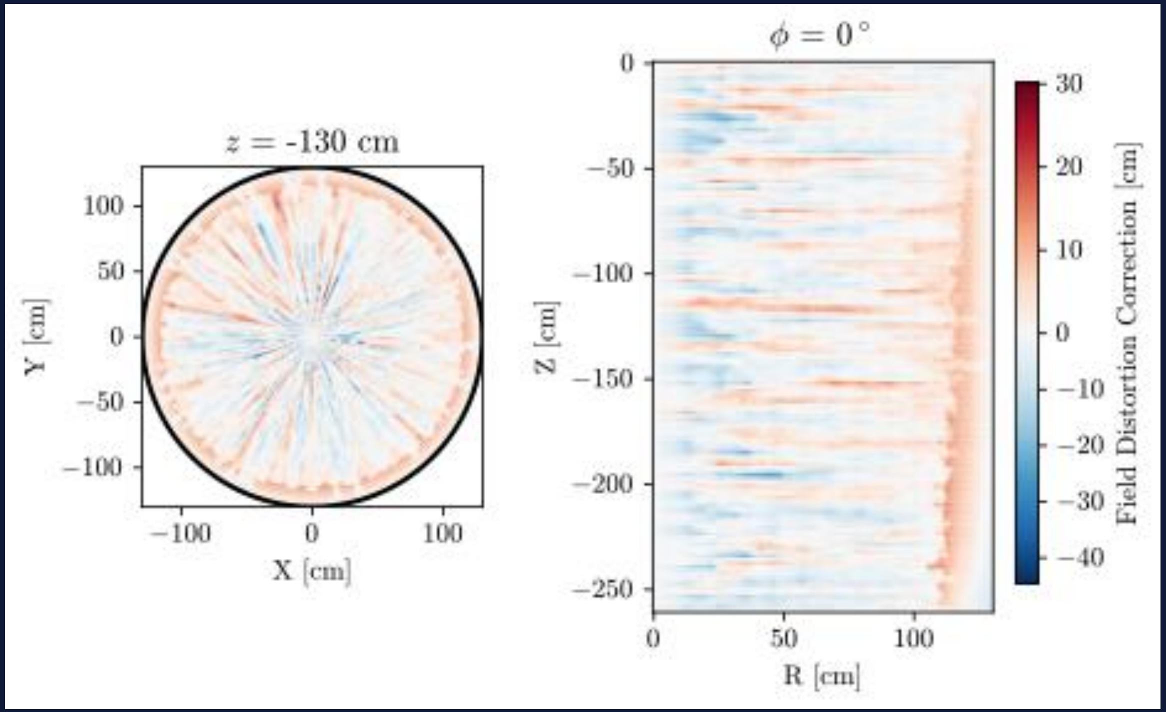

XENONnT-style FDC map — those radial spokes & stripes aren't physics, they're binning artifacts. This is terrible! 😱

Calibrate with uniform ^{83m}Kr → compare reconstructed vs. expected radial percentiles → tabulate Δr(r, ϕ, z), then interpolate.

Three structural problems

1 **Not smooth, not differentiable**

It's a lookup table. You can't propagate position uncertainties through it.

2 **Ignores Maxwell**

The transformation has no obligation to correspond to any real electrostatic field.

3 **Data-hungry**

Needs ~10^7 calibration events to fill every (r, ϕ, z) bin with enough statistics.

We want a field transformation that's **differentiable · curl-free · cheap · realistic**

THE KEY IDEA

Learn the drift, bake the physics into the architecture.

Physics in the loss (PINN-style)

- Add a curl-penalty term to the loss.
- Constraint only approximately satisfied.
- Needs a hand-tuned weight balancing data vs. physics.

We tested this — it works no better, and adds a knob.

Physics in the architecture (ours)

- Network outputs a scalar potential g_ϕ.
- Drift field = $-\nabla g_\phi$ → conservative by construction.
- Faraday's law ($\nabla \times E = 0$) holds exactly, no penalty, no knob.

The constraint can't be violated — it's structural.

METHOD · 1

An electron's drift written as a neural ODE

1. **Drift follows the field**

   $\partial r/\partial t = -\mu(|E|) \cdot E(r)$ — electrons ride the field lines.

2. **Reduce to the transverse plane**

   Drift is ~vertical & $|E|$ ~ constant (0.1%); solve a 2-D problem in s for each depth $z = v_d \cdot t$.

3. **Enforce conservativity**

   Set $f'_\phi = -\nabla_s g_\phi$. The network learns the potential g_ϕ, not the field.

4. **Reconstruct position**

   Integrate the ODE backward over the measured drift time → recover the interaction vertex.

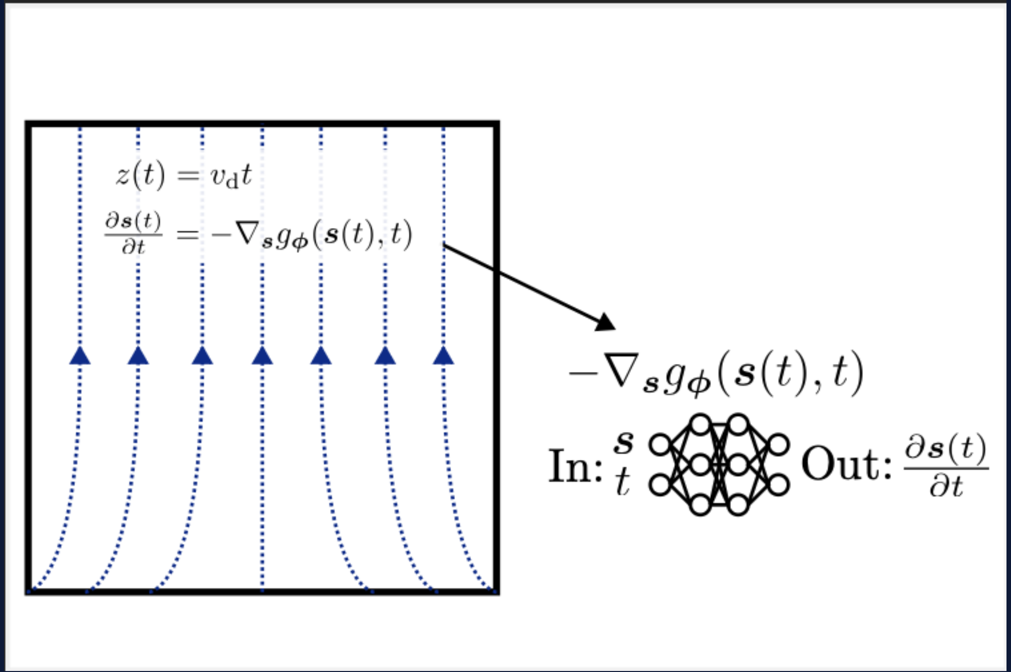

Input: transverse position s, time t (≙ depth).
Output: how that point moves.

Training with no ground truth: density estimation

The constraint

On real data we never know an event's true position. So we can't train on reconstruction error — there's no target to compare against.

Use what we DO know: ^{83m}Kr is uniform

- Train the flow so that corrected positions match the expected uniform distribution (modulo known charge loss).
- This is simulation-based inference — neural posterior estimation in disguise.
- Loss = negative log-likelihood

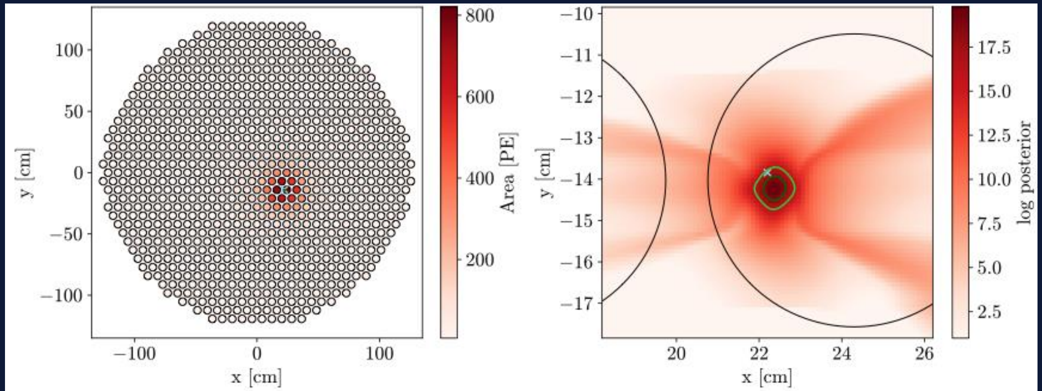

We even feed in a full position posterior per event (right), not a single point — uncertainties propagate through the flow.

survival prob. rolls off to 10^{-3}, never 0 → finite gradients in the charge-insensitive volume

VALIDATION SETUP

A realistic XLZD-scale simulation to test

To check the method honestly, we built an end-to-end simulation where we DO know the truth — then we hide it from the model.

XLZD-scale geometry

260 cm tall, 130 cm radius. XENONnT field cage scaled up.

Real field, solved

Laplace's equation in FEniCS with a linear PTFE wall-charge model.

Full electron transport

Drift, diffusion, electron lifetime, survival probability maps.

Light & charge yields

NEST for S1/S2; BACCARAT/Geant4 optical maps.

Key point: the model only ever sees the hit pattern and the drift time — exactly what a real experiment sees.

RESULT · 1

An order of magnitude less calibration data

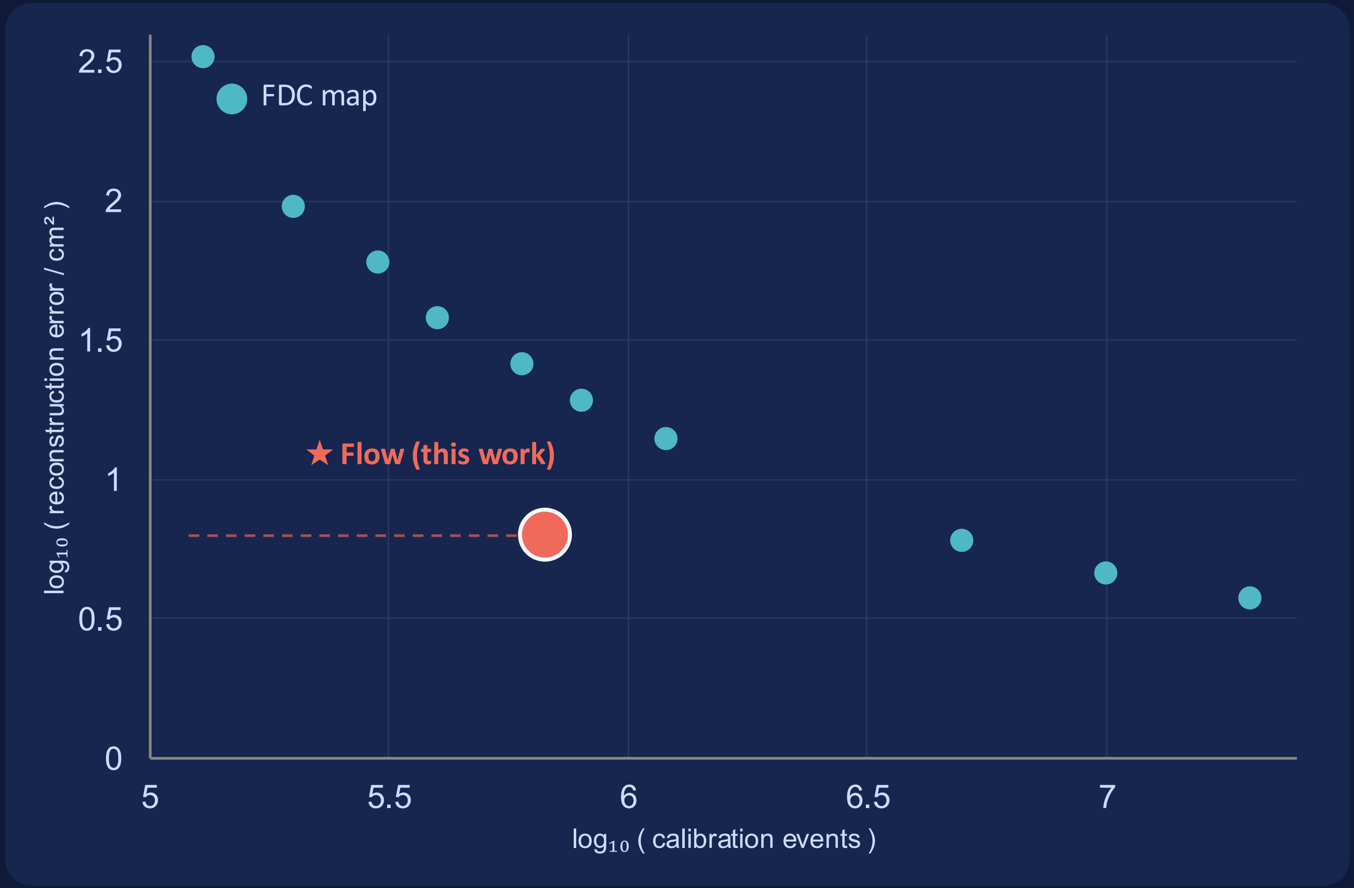

Same accuracy. ~10× less data — and at fixed data the flow is more accurate.

MSE ≈ 6 cm² on a held-out test set of 5×10⁵ events.

Better near the wall — where rare-event searches need it most

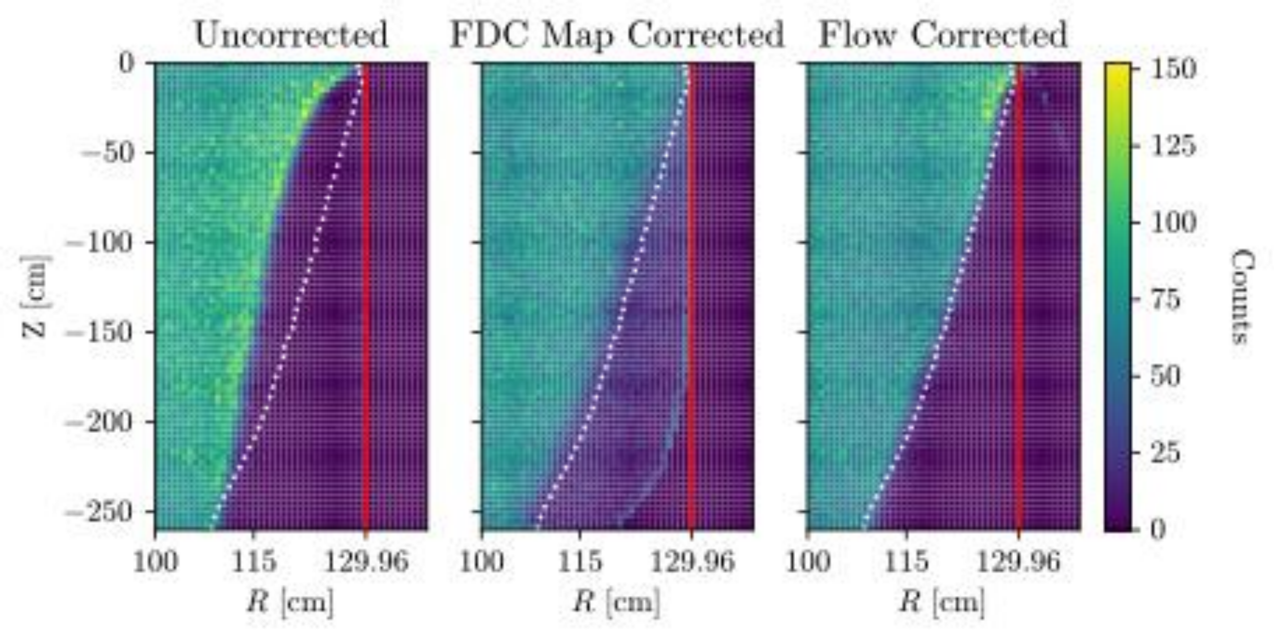

Event density near the wall: uncorrected → FDC → flow. White dotted = edge of charge-insensitive volume.

Why the edge matters

FDC starves at the edge

Limited wall statistics → it can't reconstruct the boundary cleanly.

Flow tracks the contour

Corrects events right up to the charge-insensitive edge — smoothly.

RESULT · 3

A bonus the FDC map can never give: the field itself

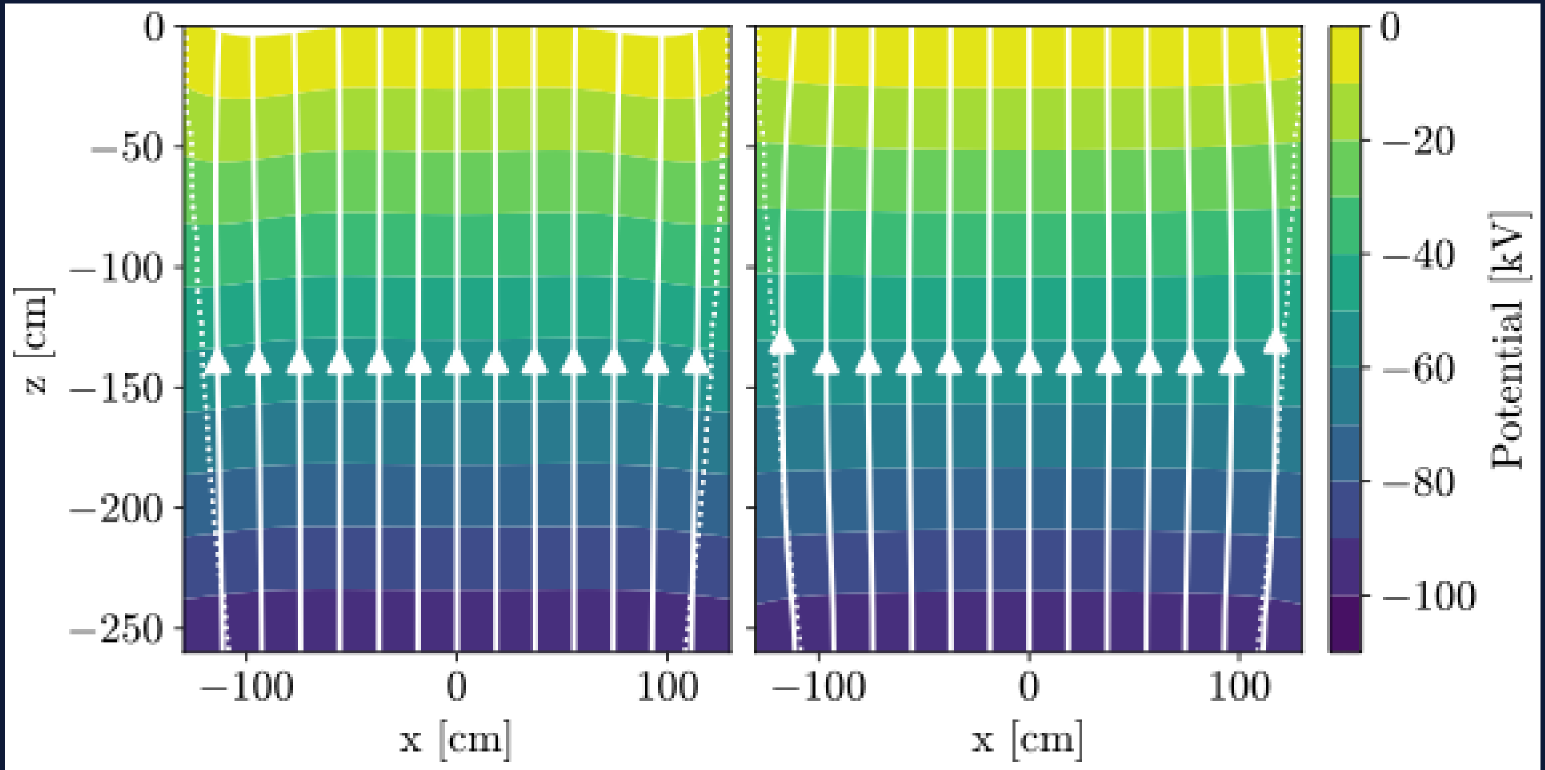

Learned scalar potential & field lines (left) vs. simulation ground truth (right).

Because g_φ is a real potential…

- We read out the scalar potential and the electric field lines directly.
- They match the ground-truth field.
- An interpretable, physical field model — learned purely from calibration data.

TAKEAWAYS

Put the physics in the architecture, and the data does more.

Conservative by construction

A continuous normalizing flow over a learned scalar potential — curl-free, smooth, differentiable. No penalty term, no balancing knob.

10× less calibration data

Same accuracy as an FDC map with an order of magnitude fewer events → monthly/weekly field monitoring instead of yearly.

Better where it counts

Sharper reconstruction near the wall and preserved local structure — directly improving fiducialization & background rejection.

An interpretable field

Recovers the actual potential and field lines from data; uncertainties propagate through the differentiable transform.

Thank you

Physics-informed continuous normalizing flows for TPC electric fields

Looking ahead: · deployment on current experiments (LZ, XENONnT)
Apply similar methods to measure liquid flow

Questions?

arXiv:2511.01897 · *pgaemers@slac.Stanford.edu*

https://github.com/RiceAstroparticleLab/fieldflow