

BIB mitigation with AI

Simons progress report
Sep 1, 2026

John Dervan¹, Kiley Kennedy², Ryan Roberts³

¹*University of Tennessee, Knoxville*

²*Princeton University*

³*University of Chicago*



Quick Updates

Meetings:

*Meetings are back to the
Wednesday 10a CT slot*

Update from UChicago

- Wrapped up initial smartpixels for muon collider effort
- Preprint: <https://arxiv.org/abs/2608.21327>
- Dataset will appear here: <https://doi.org/10.5281/zenodo.22016562> (waiting on approval from smartpixels collaboration)
- **Thanks to everyone who provided comments on the draft and the work in general.**
- Study of LUT+FF as a metric converged on LUT + 0.5 *FF as more representative, but concluded
- Enforcing that weights be programmable – main effect is on the latency and initiation interval, but these are likely implementation dependent
- Strange “arcs” in dataset conclusively tied to phi rotation – believed not to be a significant barrier for these studies

Update: August 5, 2026

Quick Updates

Meetings:

*Meetings are back to the
Wednesday 10a CT slot*

Update from UChicago

- FPGA optimization: turning off DSPs was expected to decrease latency, further study required
- Targeting a more realistic 10–20ns (vs 5 ns) produces sensible results with lower latency + resource usage for Model 1
- SmartPixels for μ C paper draft posted on Slack; **open for comments**

Update from Princeton

- Bit flip studies for LGNs presented at last meeting
- Training with bit flips during training improves robustness to bit flips at inference time, though expected real-world bit flip rate much lower than tested
 - May provide robustness to other distortions of input data, including detector noise

Update: July 14, 2026




Quick Updates

Meetings:

Need to reschedule new meeting time—when2meet forthcoming

Update from UChicago

- Showed design space exploration plots from HLS4ML with trial configs sorted by performance rank vs. LUTs, areaScore, and designLatency
- Plotted bkg rejection vs Catapult area score → knee occurs at ~200k μm^2




Update: June 9, 2026

Major Updates


Report on SmartPixels / LGN – Princeton

- Studying the impact of sensor irradiation using SmartPixels irradiated datasets (0, 370, 1100 fb)
- Observe significantly diminished performance when only training over non-irradiated data
- Improved performance when training includes irradiated data
- Open questions/ongoing work:
 - Why asymmetry in left plot? Why 370 fb and 1100 fb behavior “flipped”?
 - Extend training scheme studies to LGNs
 - Synthesis in FPGAs vs ASICs

LGN Discrete: Train non-irradiated, test irradiated



Unquantized Torch: Different training schemes



Major Updates

DNN Model Update – Chicago (Daniel)

- Input scaling significantly improved QKeras / HLS agreement
- Open questions/ongoing work: understanding and improving agreement


SmartPixel ASIC – Chicago (Daniel)

- Plans:
 - Run cold tests without the chip this week
 - Run cold test with the chip next week

General Project Direction

- Brief discussion about various project distinction / cohesion – plan to revisit in a few more weeks

Current Roadblocks




Update: May 8, 2026

Major Updates


Report on SmartPixels / LGN – Princeton

- Studying the impact of sensor irradiation using SmartPixels irradiated datasets (0, 370, 1100 fb)
- Observe significantly diminished performance when only training over non-irradiated data
- Improved performance when training includes irradiated data
- Open questions/ongoing work:
 - Why asymmetry in left plot? Why 370 fb and 1100 fb behavior “flipped”?
 - Extend training scheme studies to LGNs
 - Synthesis in FPGAs vs ASICs

LGN Discrete: Train non-irradiated, test irradiated



Unquantized Torch: Different training schemes



Major Updates

DNN Model Update – Chicago (Daniel)

- Input scaling significantly improved QKeras / HLS agreement
- Open questions/ongoing work: understanding and improving agreement

SmartPixel ASIC – Chicago (Daniel)


- Plans:
 - Run cold tests without the chip this week
 - Run cold test with the chip next week

General Project Direction

- Brief discussion about various project distinction / cohesion – plan to revisit in a few more weeks

Current Roadblocks

Neuromorphic networks (UTK) – some delays in communication with collaborators




Update: April 24, 2026

Major updates since last time


- **Report on SEE Impact Studies – Chicago**

- General degradation of model performance
 - Models with high background rejection generally see significant degradation when subjected to radiation effects
 - Possible mitigations under study
- Susceptibility of individual bits
 - Individual bits that are particularly important can cause temporary module failures
 - Should be able to identify and reset modules on time scales $\ll$ interfill (Tova)



with both external likelihood
timing information

ral networks on Sr
ized NNs for Model
ith simulated sens
level of damage



3



son for LGN

Major updates since last time

Meetings:

Wednesdays at 11 am ET

[[Indico](#)] — all welcome!



*Incomplete LGN

²Training⁴

- **Report on LGN synthesis — mlla**

- LGN synthesis now possible with both external library ([da4ml](#)) and in-house script
 - FPGA resource usage and timing information produced for Models 2 (cluster profiles) and 3 (convolutions)
- Quantized torch training of neural networks on Smartpixels dataset to provide fair comparison for LGN
 - LGNs outperforming quantized NNs for Model 2
- Initial studies using datasets with simulated sensor damage from radiation
 - Testing possibility of providing level of damage as model input

Update: April 10, 2026

Major updates since last time

- **Report on SEE Impact Studies – Chicago**

- General degradation of model performance
 - Models with high background rejection generally see significant degradation when subjected to radiation effects
 - Possible mitigations under study
- Susceptibility of individual bits
 - Individual bits that are particularly important can cause temporary module failures
 - Should be able to identify and reset modules on time scales $\ll$ interfill (Tova)

- **Report on LGN synthesis – Mila**

- LGN synthesis now possible with both external library ([da4ml](#)) and in-house script
 - FPGA resource usage and timing information produced for Models 2 (cluster profiles) and 3 (convolutions)
- Quantized torch training of neural networks on Smartpixels dataset to provide fair comparison for LGN
 - LGNs outperforming quantized NNs for Model 2
- Initial studies using datasets with simulated sensor damage from radiation
 - Testing possibility of providing level of damage as model input

Current Roadblocks

- **LGN:** Availability of muon collider dataset

Ongoing Work

- **LGNs:** Further investigation of radiation damage effects.
- **SNNs:** Evaluating personpower.
- **DNNs:** Understanding optimal choice for input quantization
- **Smartpix Lab Testing:** Calibrating cold box temperature measurement

Open Tasks

- **SNN:** Iterate on dataset, want to first conclude ECAL study before moving to smartpixels
- **Tracking:** Muon collider dataset incorporating merged clusters with detailed digitization
 - Common harness
- **General:** Gain knowledge about ASIC synthesis

Update: March 27, 2026

Major updates since last time

- *Meetings: Wednesdays at 11 am ET [[Indico](#)] — all are welcome!*
- **BDTs for BIB rejection**
 - Heard update from Jordan + nice discussion
- **SNNs (UTK)**
 - Applications to ECAL
 - Previous work: developing a training dataset
- **Outlining our overarching goals/pillars—subject to shuffling + redefinition!**
 - Defining operation of ML-enabled FEs
 - Towards a public dataset for BIB rejection
 - More advanced synthesis of BIB rejection models
 - Benchmark models using current technologies under consideration (DNN, BDT, SNN, LGN)

Current roadblocks

- **LGN** – unclear why we cannot exactly reproduce public results with SmartPixels dataset
- **Smartpix Training** – Messy environment for quantized training and synthesis – QKeras out of date, hls4ml differences for Vitis and Catapult
- **Smartpix Lab Testing** – Noisy filter outputs make validation harder


Open tasks

- **LGNs**: proceed with synthesis
- **SNNs**: iterate on dataset, want to first conclude ECAL study before moving to smartpixels
- **DNNs**: Want to refine estimates of chip area, data rate. Eventually validate synthesis on FPGA.
 - Longer term: consider differential distributions, digitization implementation.

Update: March 13, 2026

Major updates since last time (3/13/2026)

- **Meetings:** Wednesdays at 11 am ET [[Indico](#)]
 - All are welcome to join!
 - First meeting on 3/11 [[Indico](#)] – several contributions (not enough time to cover all!)
- **LGN Updates:**
 - Very nice introduction to LGNs (Lino)
 - Partially reproduced SmartPixels public results w/ conventional NNs (Mila)
 - LGNs (**relaxed** and **discretized**) perform nearly as well as **conventional NNs** on SmartPixels dataset
- **UTK ML for clustering**
 - ML applications to cluster shape to be discussed at next meeting
 - SNN for smartpixels study in revival/literature phase; training to be described
- **Chicago Smartpix NNs for BIB**
 - Understanding details of hyperparameter tuning.
 - Lab testing of filtering ASIC. (lower priority recently)
 - Rough investigation of SEE effects in weights.



Current roadblocks (3/13/2026)

- No significant roadblocks
- **LGN** – unclear why we cannot exactly reproduce public results with SmartPixels dataset
- **Smartpix Training** – Messy environment for quantized training and synthesis – QKeras out of date, hls4ml differences for Vitis and Catapult
- **Smartpix Lab Testing** – Noisy filter outputs make validation harder

Open tasks

- **LGNs**: proceed with synthesis
- **SNNs**: tasks TBD following literature conclusion
- **DNNs**: Want to refine estimates of chip area, data rate. Eventually validate synthesis on FPGA.
 - Longer term: consider differential distributions, digitization implementation.