

Neutrino Weekly Status

Sean Gasiorowski

On behalf of the neutrino ML team

March 9th, 2022



U.S. DEPARTMENT OF
ENERGY

Stanford
University

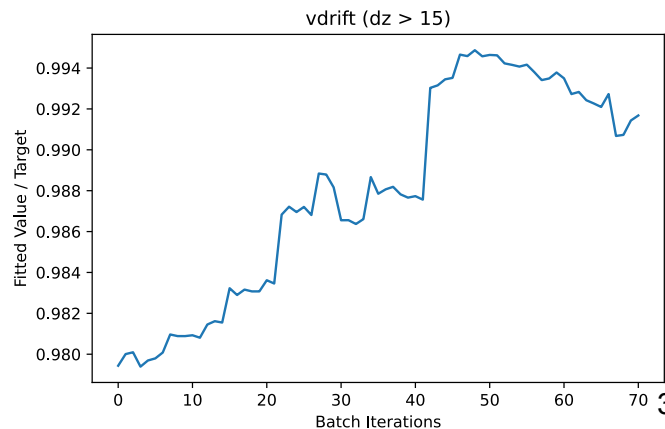
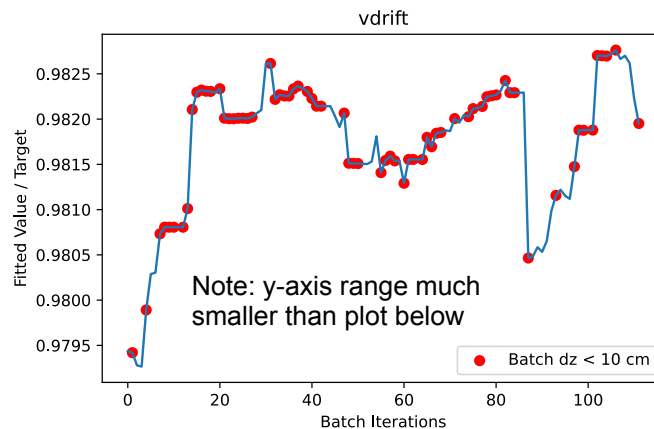


NATIONAL
ACCELERATOR
LABORATORY

Broad areas of work

- Debugging (differentiable TPC sim)
- Results: parameters and metrics (differentiable TPC sim)
- Speed/performance (differentiable TPC sim)
- Inverse solver
- SIREN PMT

- Last week: discussion of dataset coverage/convergence tests
 - Yifan's talk discussion:
 - Checks of convergence as function of e.g., track angle
 - Can also find an expected width for inputs (before the detector simulation)
- Know there are some noise dependent effects for eField (nice [summary slides](#) from Yifan)
- Quick vdrift results:
 - Jumpy convergence path even with no readout noise
 - Flat bits seem to be tied to batch z coverage, basic requirement on this => less jumpy (though still not perfect)
- Follow-ups (discussions with Youssef/Yifan): try a large run (big batches, let run for a while) — in progress (Sean)



Results: parameters and metrics

- Results here still dependent on optimization concerns
- **Thoughts towards the future** (discussions with Daniel/Kazu):
 - How do we tell how good the converged optimization is?
 - Do we have a baseline to compare to?
 - Standard calibration procedures, e.g.
 - If not, what statements can we make
 - How good is 5% different from the target parameter? Need a way of characterizing what this means (correspondence to some percent variation in ADC counts, e.g.)
 - How do we define an uncertainty on the fitted parameters?
 - Use loss/local landscape in some way? Ensembling could also work, but is expensive

- Code is being de-notebookified (Sean/Yifan/Youssef)
- Loops + checkpointing keep memory under control but lead to longer runtime
- Youssef is working on multi-GPU setup
 - Should make it more feasible to run with larger batches

- Presentation today from Kazu with data setup
- First step: train discriminative model
 - Gives a baseline that we can discuss with Stefano Ermon's group (re: uncertainties, etc)

- Patrick has been updating some plots
- Draft for group to read has been incoming for a few weeks
 - Patrick's goal is draft for general comments this week
- **Thoughts towards the future** (discussions with Kazu):
 - Is it enough to store means for each voxel, or is some distributional info needed?
 - Check: run slow simulation many times, look at these distributions