Parameter Studies

Sean Gasiorowski
Neutrino ML Meeting

September 27th, 2022

SLAC NATIONAL ACCELERATOR LABORATORY

Stanford University | U.S. DEPARTMENT OF ENERGY

Introduction

Points from last week:

- We need to better understand which parameters matter
    - What constraints are there from physics/the physical detector? (cf. work by Yifan)
    - Empirically, how much impact do parameters have relative to each other/electronics noise
- On right, plot from last week:
    - Scale of diffusion much smaller than pixel size — might explain some of the issues we've been having

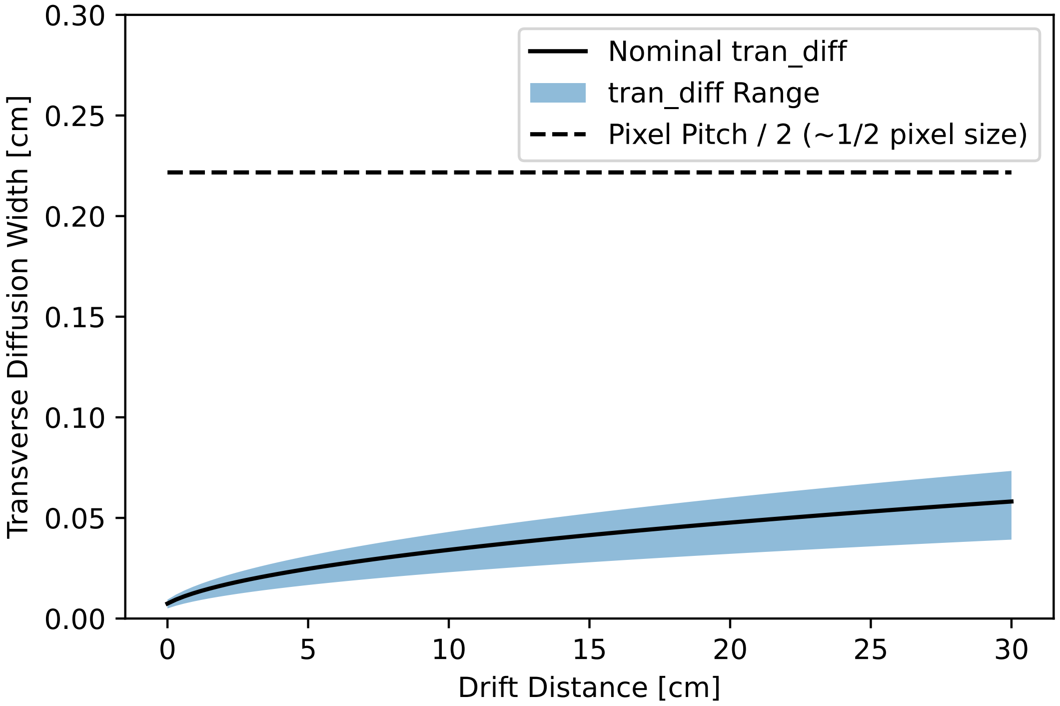

Empirical Checks

Procedure:

- Baseline: nominal parameter values (from larnd-sim/Yifan)
- Vary each parameter individually by a set amount
- Look at impact of parameter on (separately) ADC, x, y, and t output
    - "impact" defined using Dynamic Time Warping (DTW) — different from our loss!
        - Way of comparing two variable length sequences
        - Chosen because current loss definition couples these outputs via a spatial matching
        - Note: not differentiable, but there is a differentiable version (Soft-DTW) — maybe worth exploring as an alternative loss!

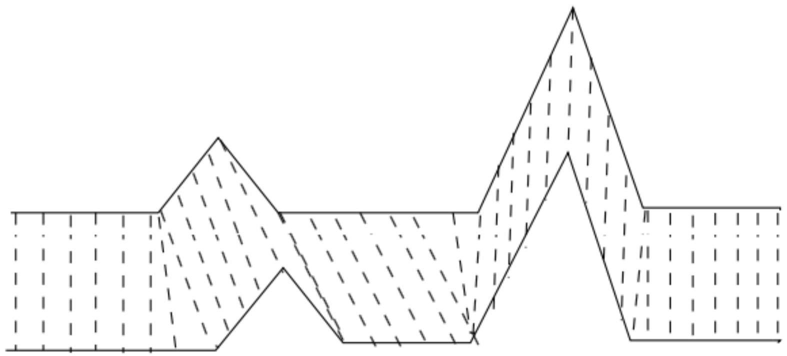

Schematic image of DTW (Wikipedia)

Empirical Checks

Procedure:

- To get a **percent difference**, calculate DTW between nominal output and output shifted by a known amount
  - For us, e.g. for ADC, we calculate DTW between:
    - Nominal ADC and Nominal ADC + 0.01*(Nominal ADC)
    - Nominal ADC and Nominal ADC - 0.01*(Nominal ADC)
    - Average these DTW values => “this how big a 1% shift in output is”
  - Changes due to changing parameters can then be written as multiples of this 1% shift
- Can use this baseline to assess noise level
  - Keep parameters at nominal
  - Simulate with noise (here 10 times)
  - Take DTW between no noise and each noisy simulation, use mean of those values as noise baseline
- Here: use same 10 tracks as have been studying (first 10 in sample, no z length selection)

Empirical Checks: ADC

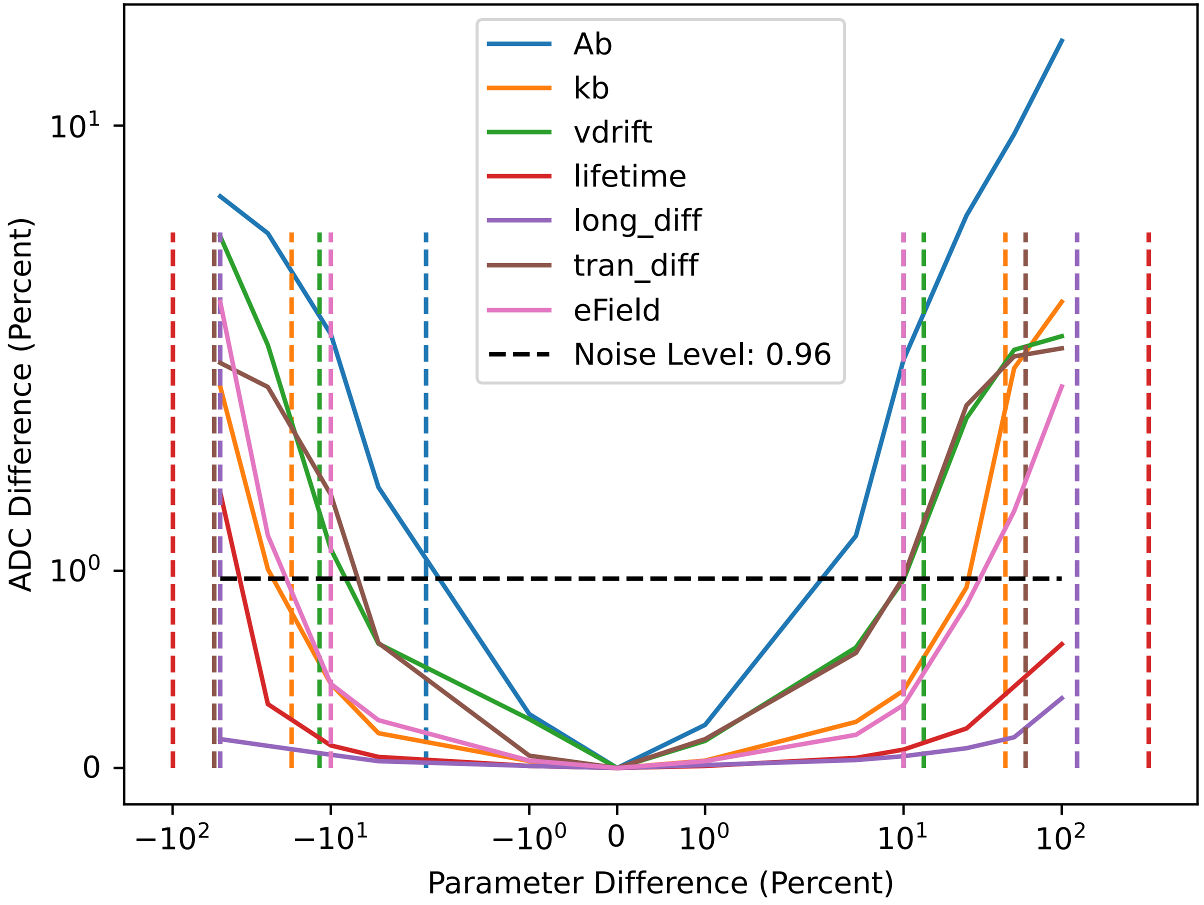

- Solid lines: percent change in ADC as a function of percent change in parameter value
- Vertical dashed lines: parameter bounds (from Yifan) — matching colors => same parameter as solid lines
- Horizontal line: Noise level (as discussed on previous slide)
- Axes are in symlog scale
  - Only go down to -50% parameter value to avoid 0's (with -100%)

Empirical Checks: ADC

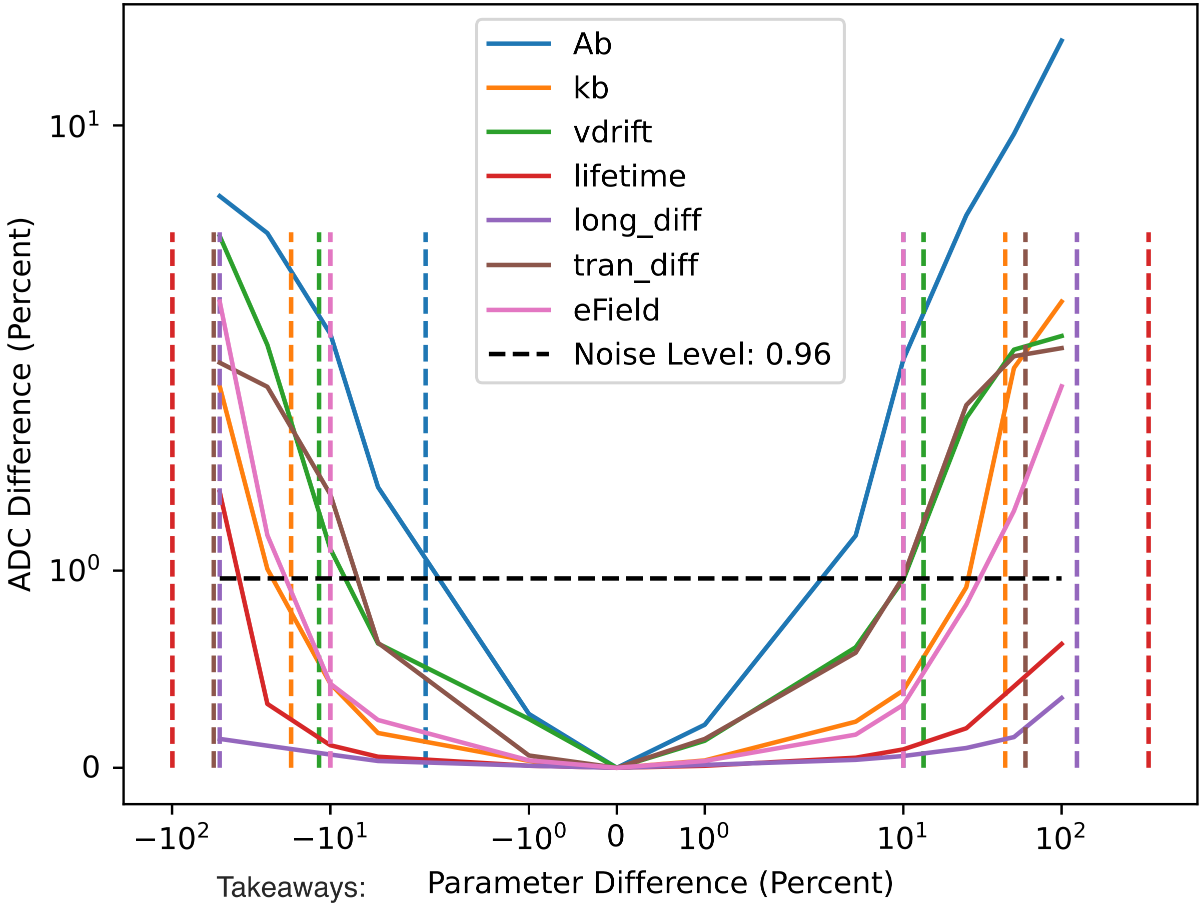

Takeaways:

- long_diff and eField have little impact
- Ab has largest impact with smallest parameter difference
- tran_diff, lifetime, vdrift, kb relevant for large % changes, low lifetime => stronger impact

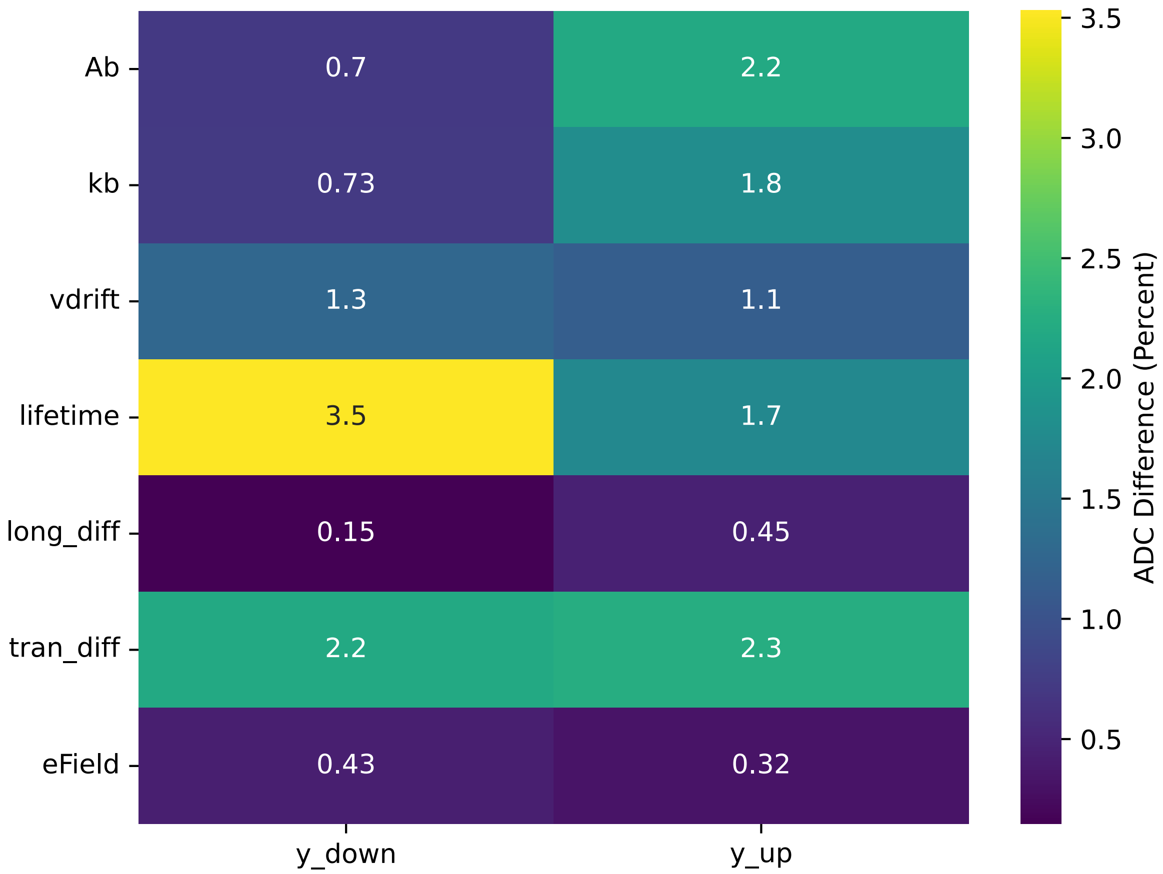

- For a sense of “maximal impact” — table has interpolated y-axis values at bounds of parameter range (dashed lines)

SLAC

Empirical Checks: t

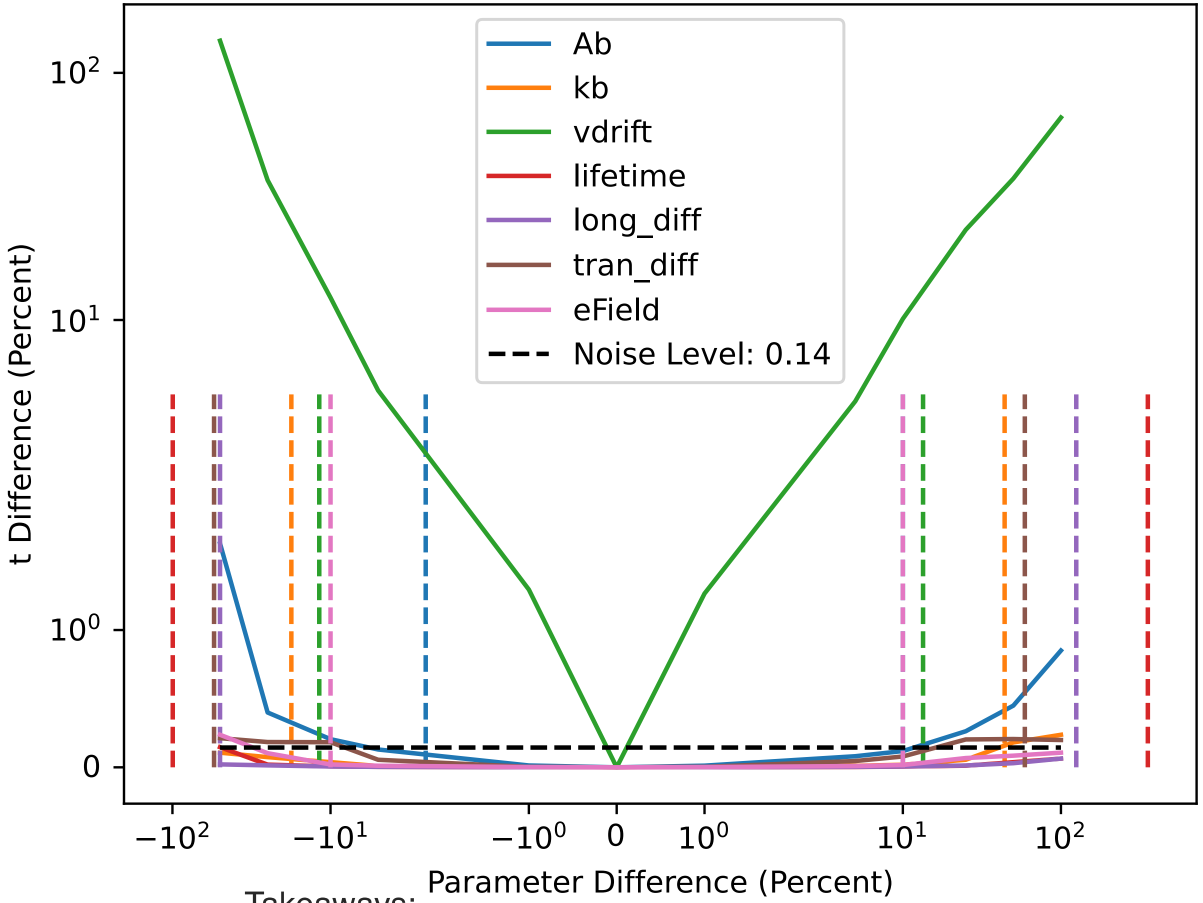

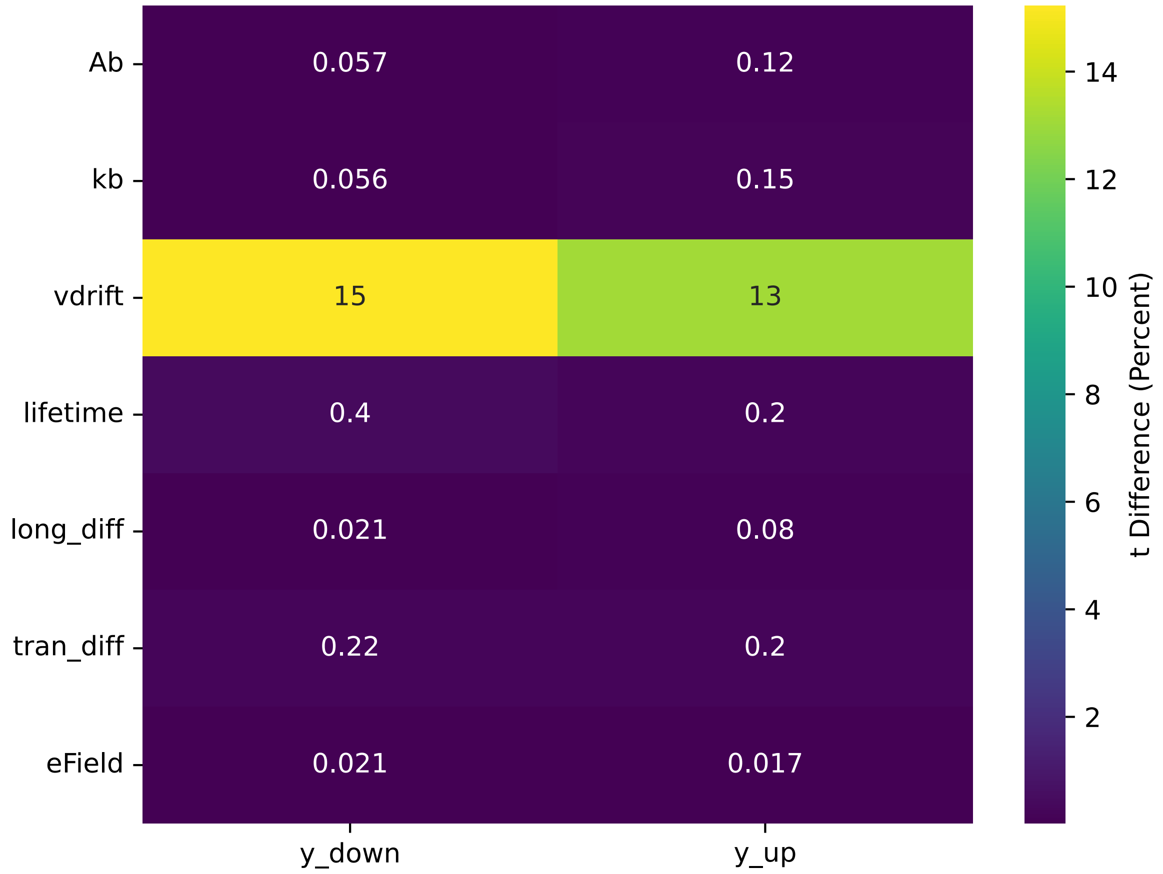

Takeaways:

- vdrift only relevant parameter for time (and has huge impact, unsurprisingly)
  - Might be convenient for fitting independently!

SLAC

Empirical Checks: x

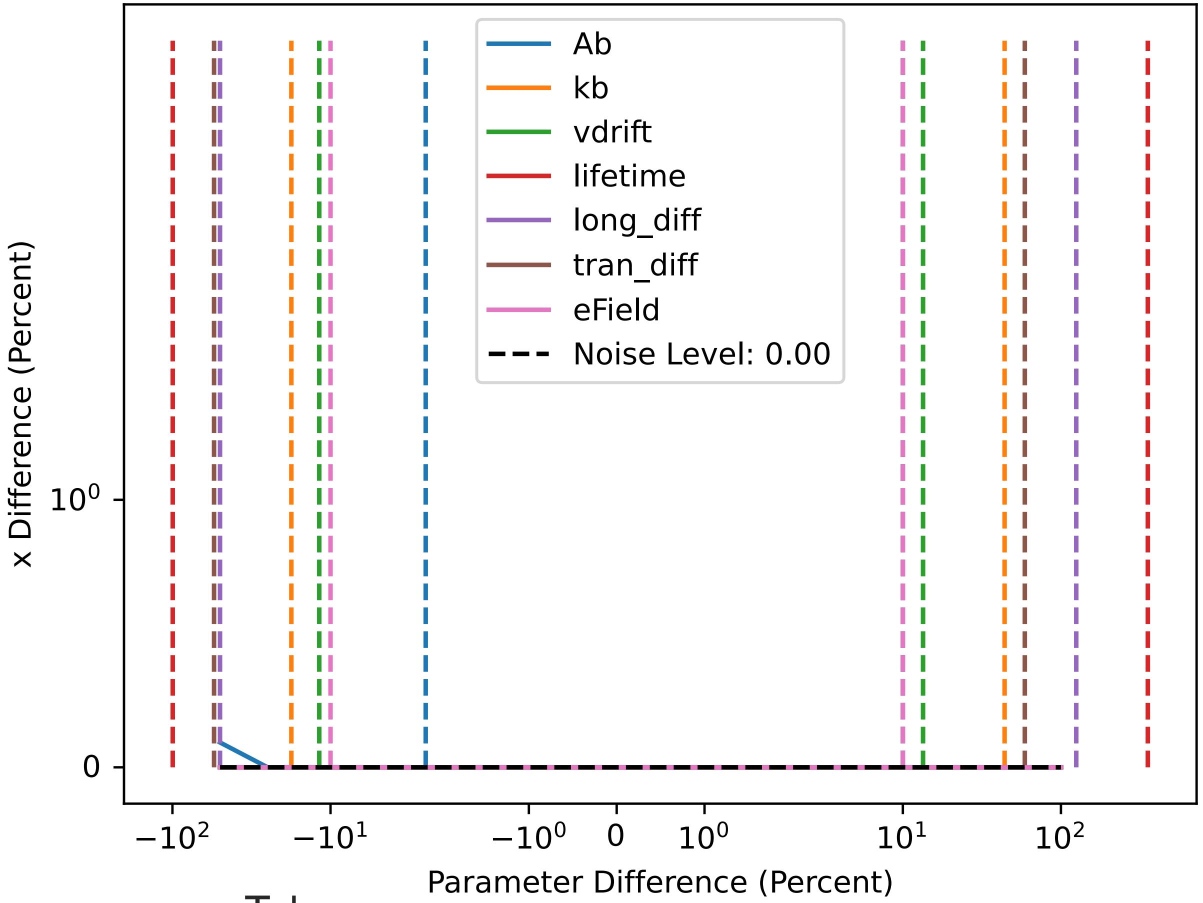

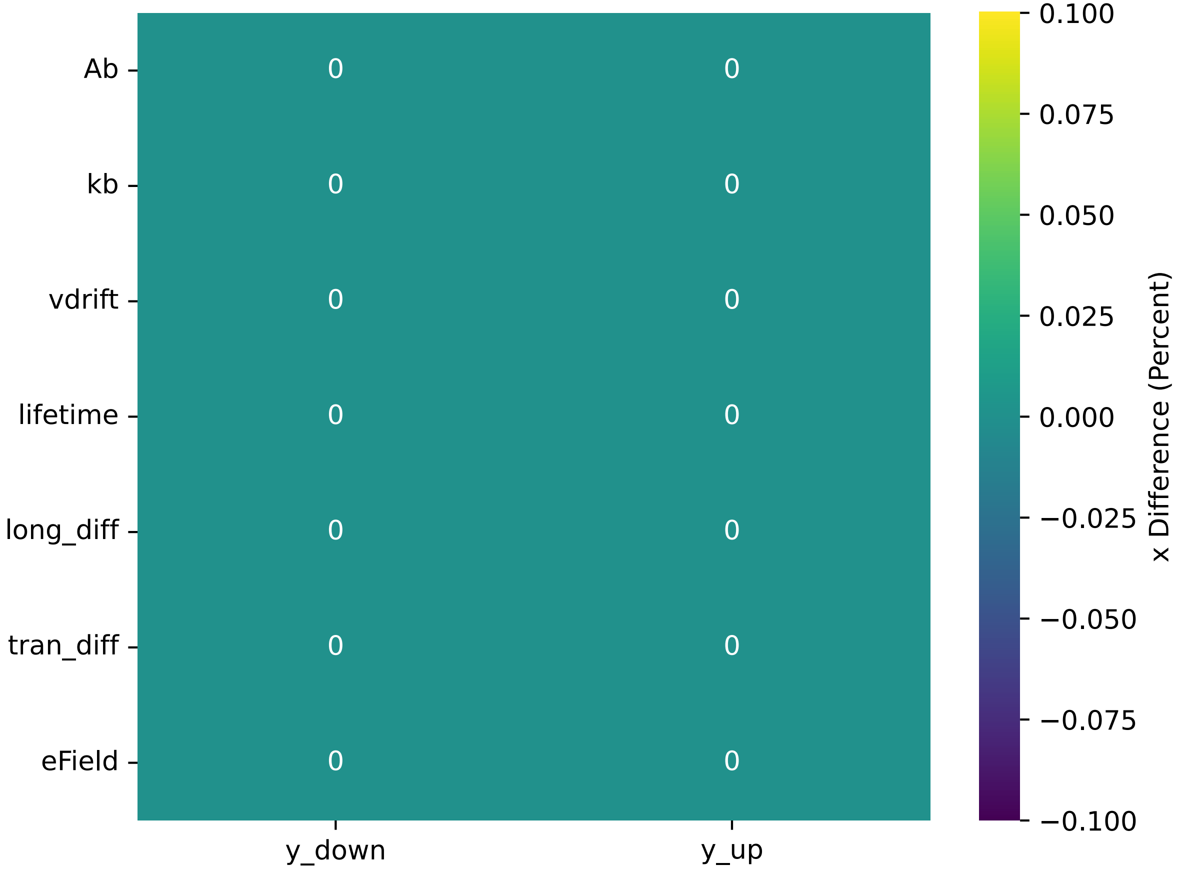

Takeaways:

- Nothing really changes x

Empirical Checks: y

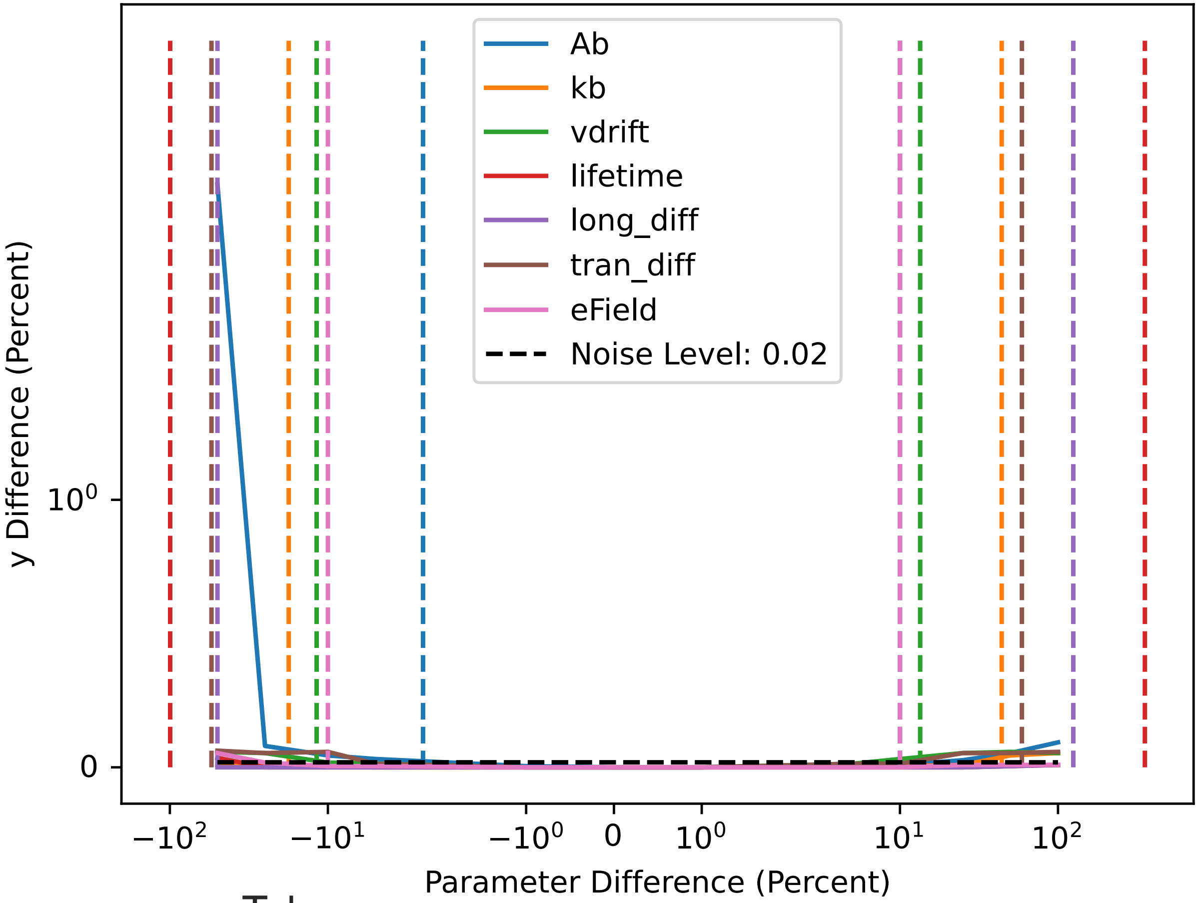

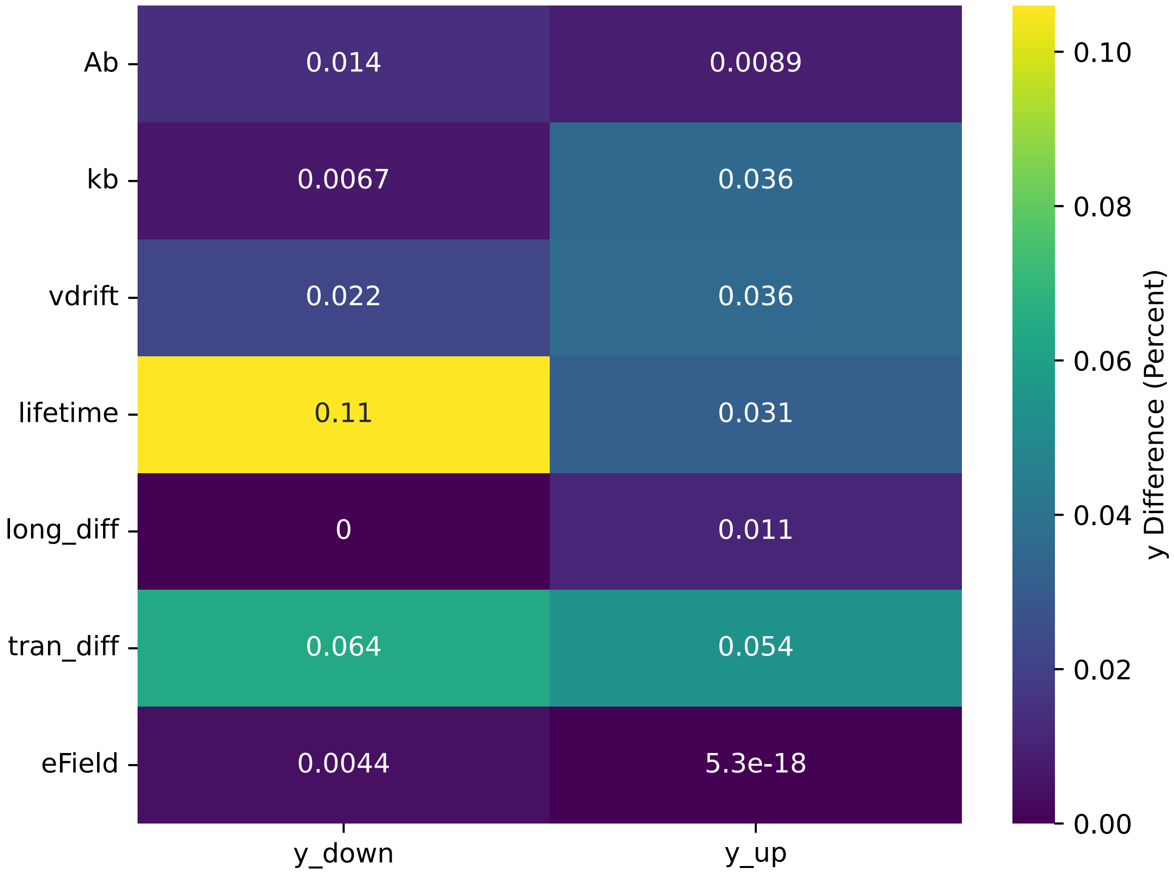

Takeaways:

- Changes to y are quite small

Takeaways

For this set of tracks (at least)

- **Longitudinal diffusion** and **eField** are the two most irrelevant parameters in their ranges, changes fall below noise level
- **Lifetime** requires very large changes for notable impact on ADC output, low values are more impactful
- **Transverse diffusion** is relevant!
    - Most open question — why do we have convergence troubles?
        - Might come from physics intuition — impact is from tail/edge effects => loss landscape isn't smooth/nice
        - Maybe something like a DTW loss captures this better?
- **kb** has maybe comparable to/smaller impact than tran_diff, but seems to be nicer in optimization
- **vdrift** has a massively dominant impact on the timing — we can probably get away with fitting this on its own, just using that info
- Some major scale differences => not unexpected that multi-parameter has some trouble

Let's look at some fit results with the same tracks to see what things look like.

How do we look in 1D?: Ab (no noise)

Converges well! ✅

10 tracks, 2 batches
lr=1e-2
Range=[0.78, 0.88]
Nom = 0.8

Large impact,
1D looks great

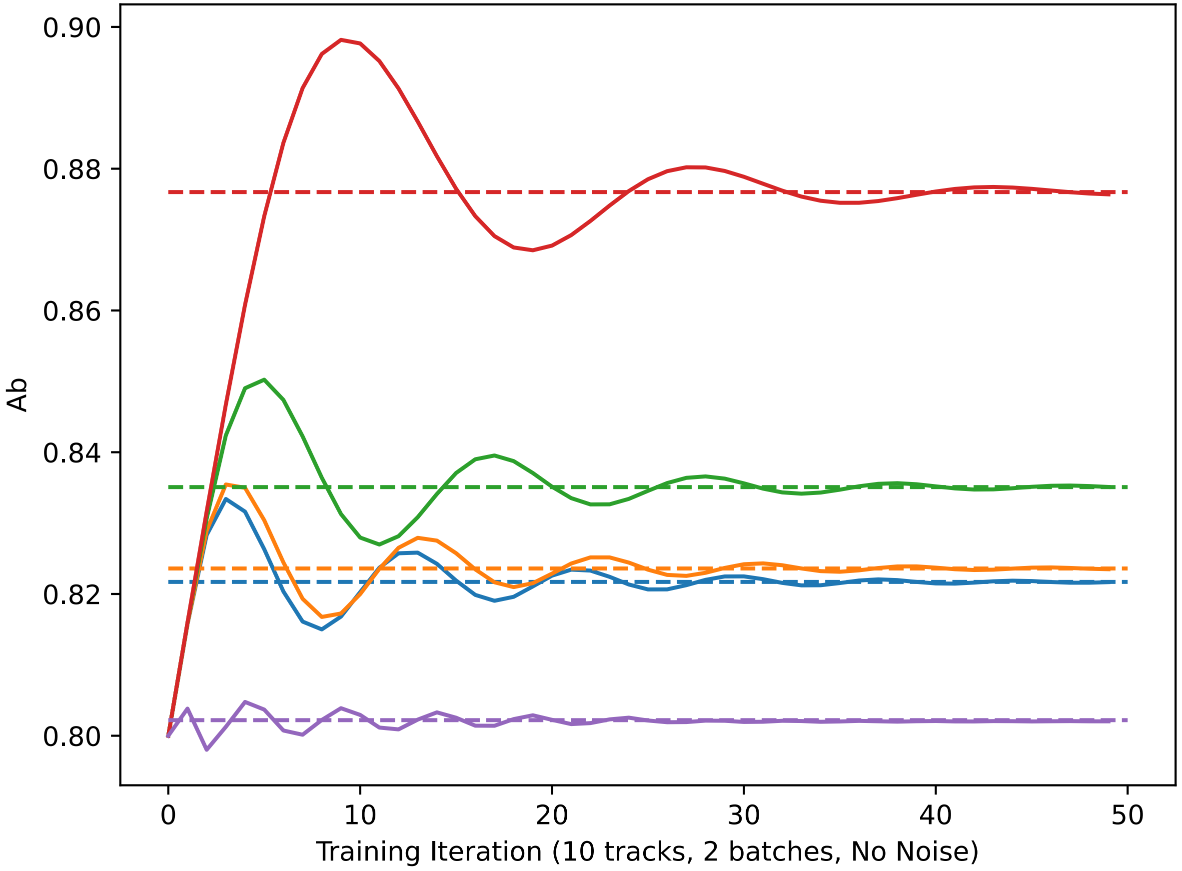

How do we look in 1D?: kb (no noise)

Converges well! ✅

10 tracks, 2 batches
lr=1e1
Range=[0.04, 0.07]
Nom= 0.0486

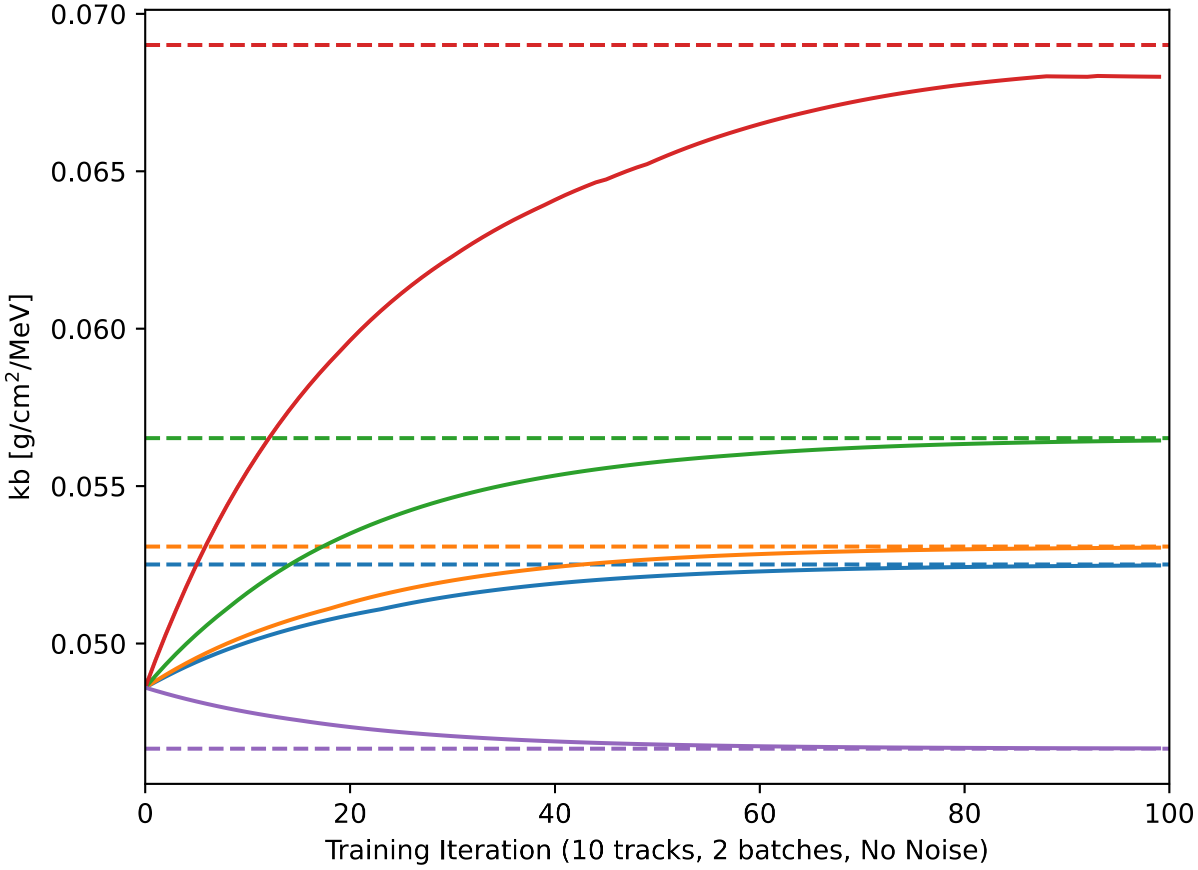

Some impact,
1D looks good

How do we look in 1D?: eField (no noise)

Converges well! ✅

10 tracks, 2 batches
lr=1e-2
Range=[0.45, 0.55]
Nom=0.5

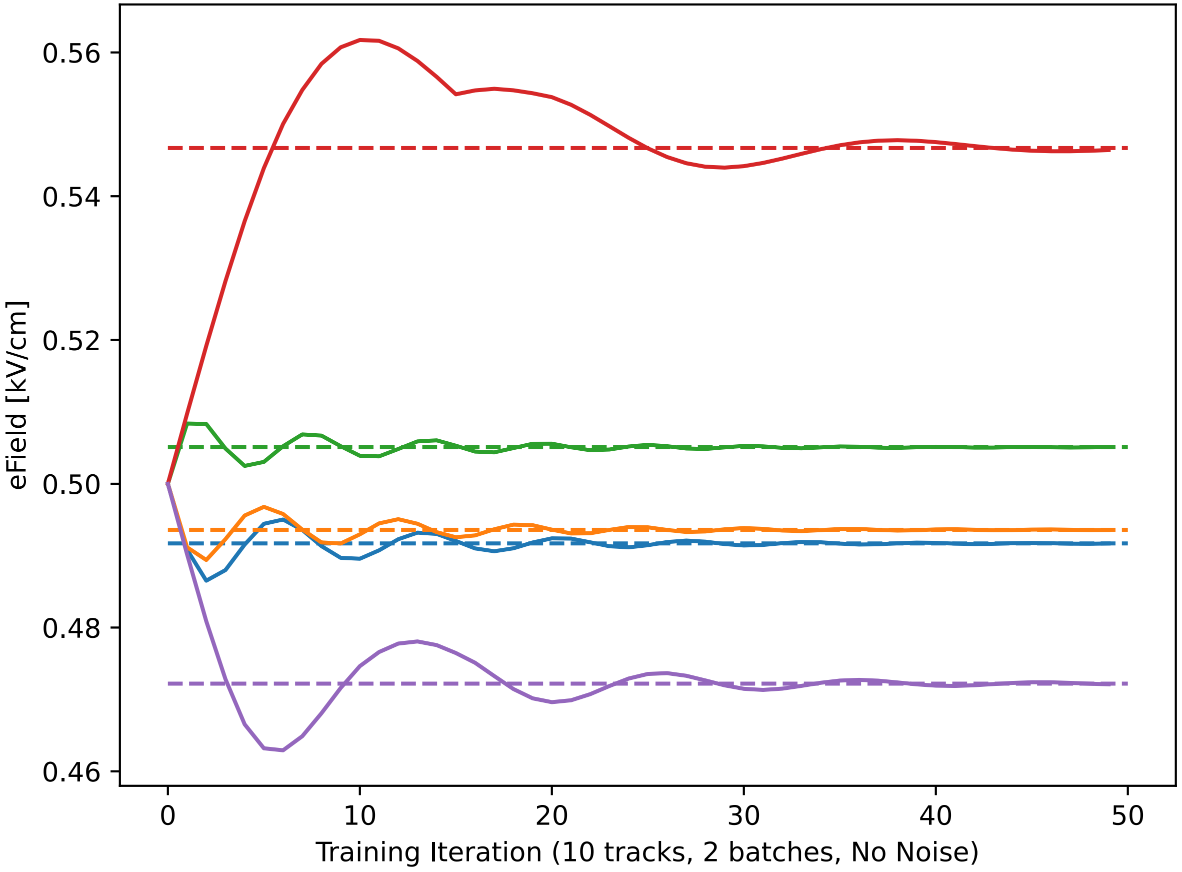

Small impact,
but 1D looks
good

How do we look in 1D?: vdrift (no noise)

With ADC in loss: red diverges ❌

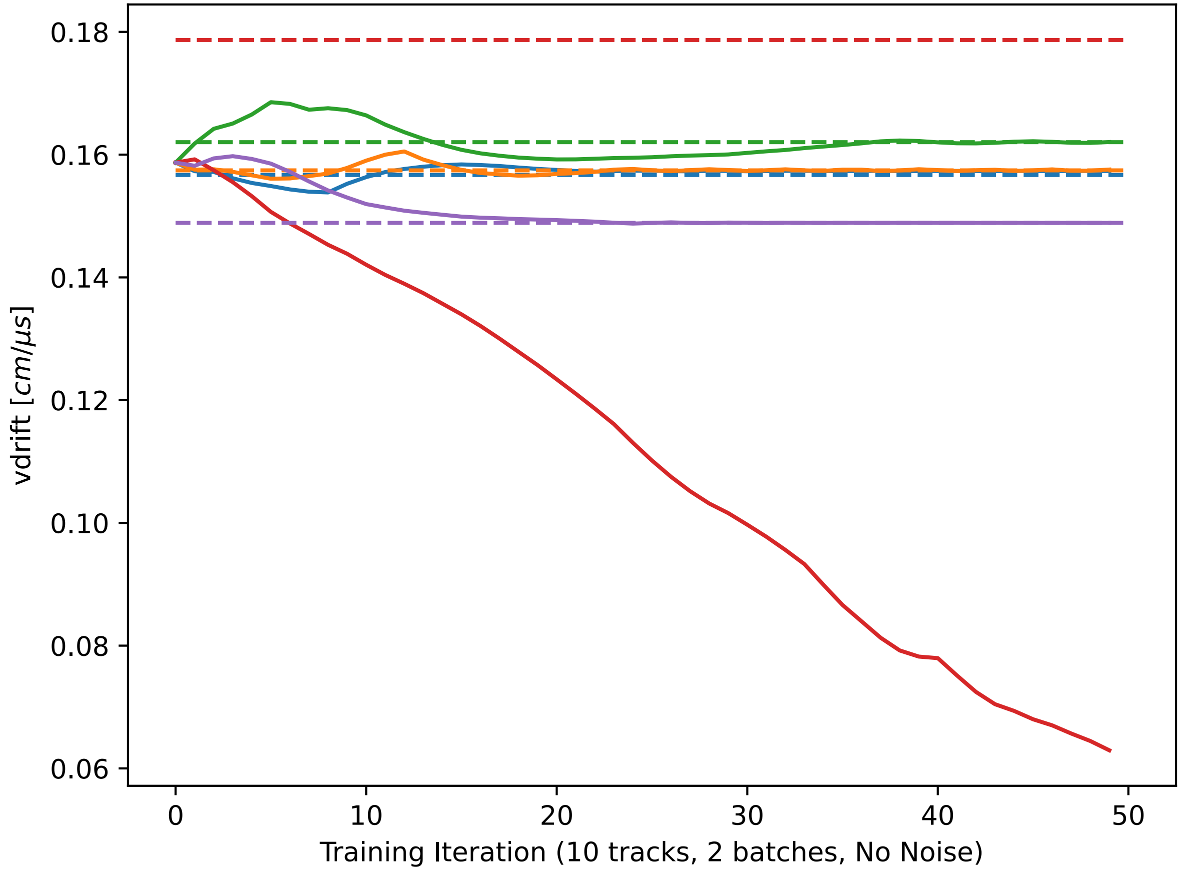

10 tracks, 2 batches
lr=1e-2
Range=[0.14, 0.18]
Nom=0.1587

With only x, y, t: Converges well! ✅

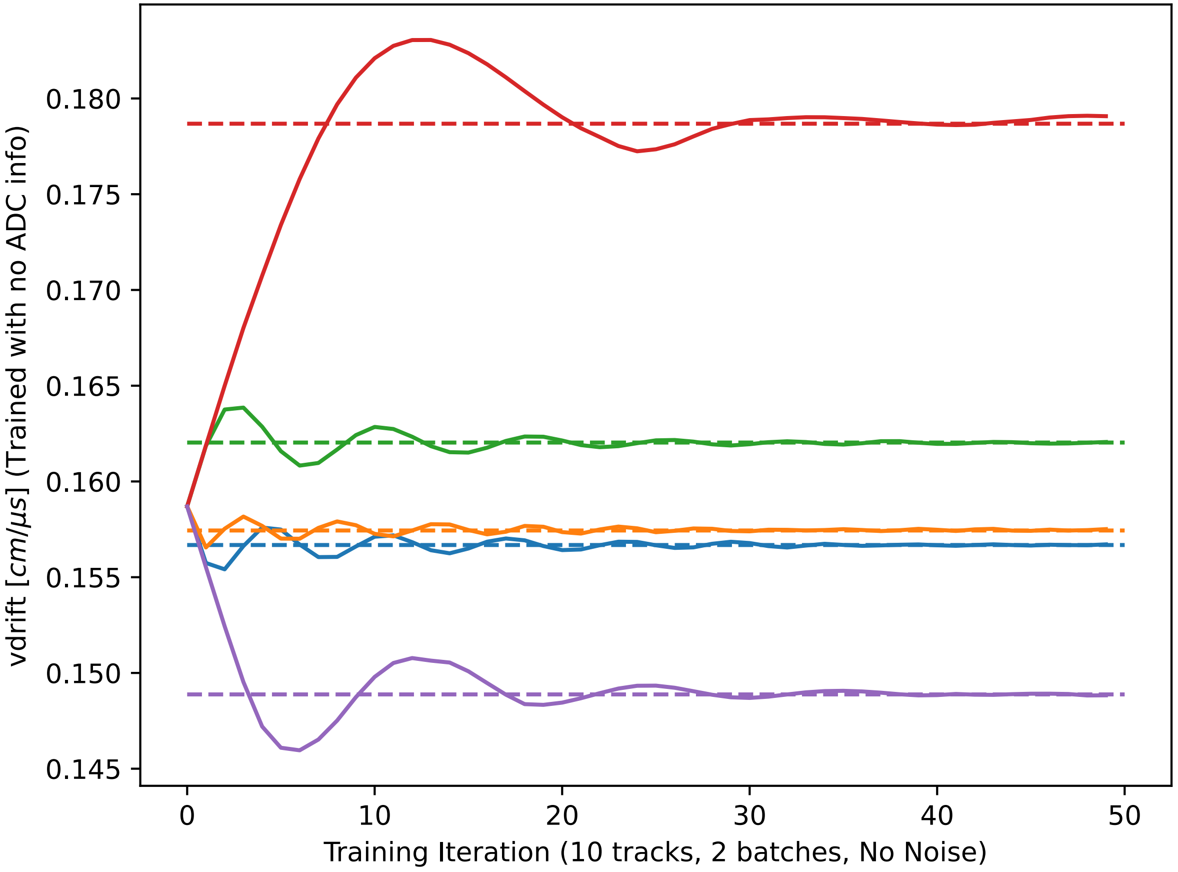

Large impact, 1D focused on timing looks good

SLAC

How do we look in 1D?: tran_diff (no noise)

Does not converge (unless we have a very good initial guess) ❌

10 tracks, 2 batches
lr=1e1
Range=[4e-6, 14e-6]
Nom=8.8e-6

Relevant impact, 1D looks bad :(

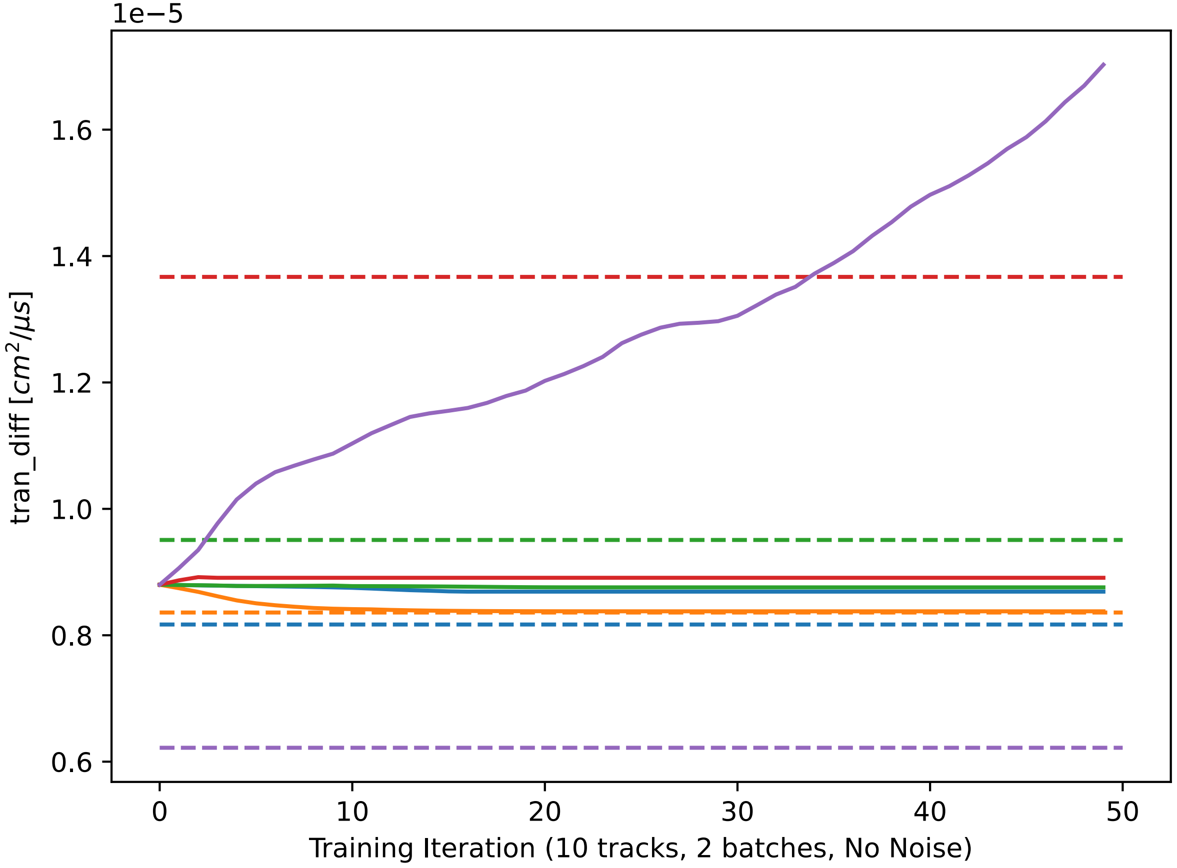

How do we look in 1D?: long_diff (no noise)

Converges ok 🆗

10 tracks, 2 batches
lr=1e4
Range=[2e-6, 9e-6]
Nom=4e-6

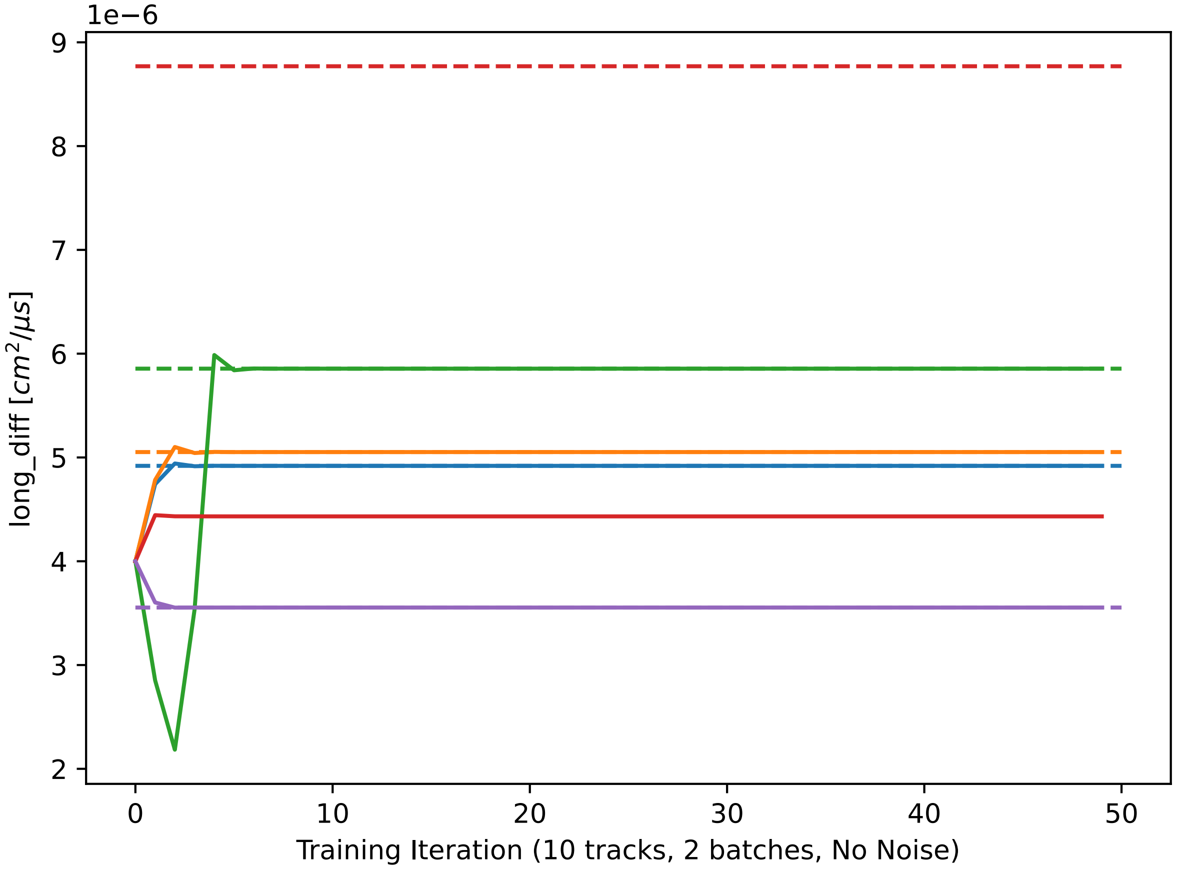

Small impact, 1D looks fine

How do we look in 1D?: lifetime (no noise)

Converges ok 🆗

10 tracks, 2 batches
lr=1e4
Range=[10, 1e4]
Nom=2.2e3

Impact for large changes, 1D looks fine – but need to understand where we have trouble!

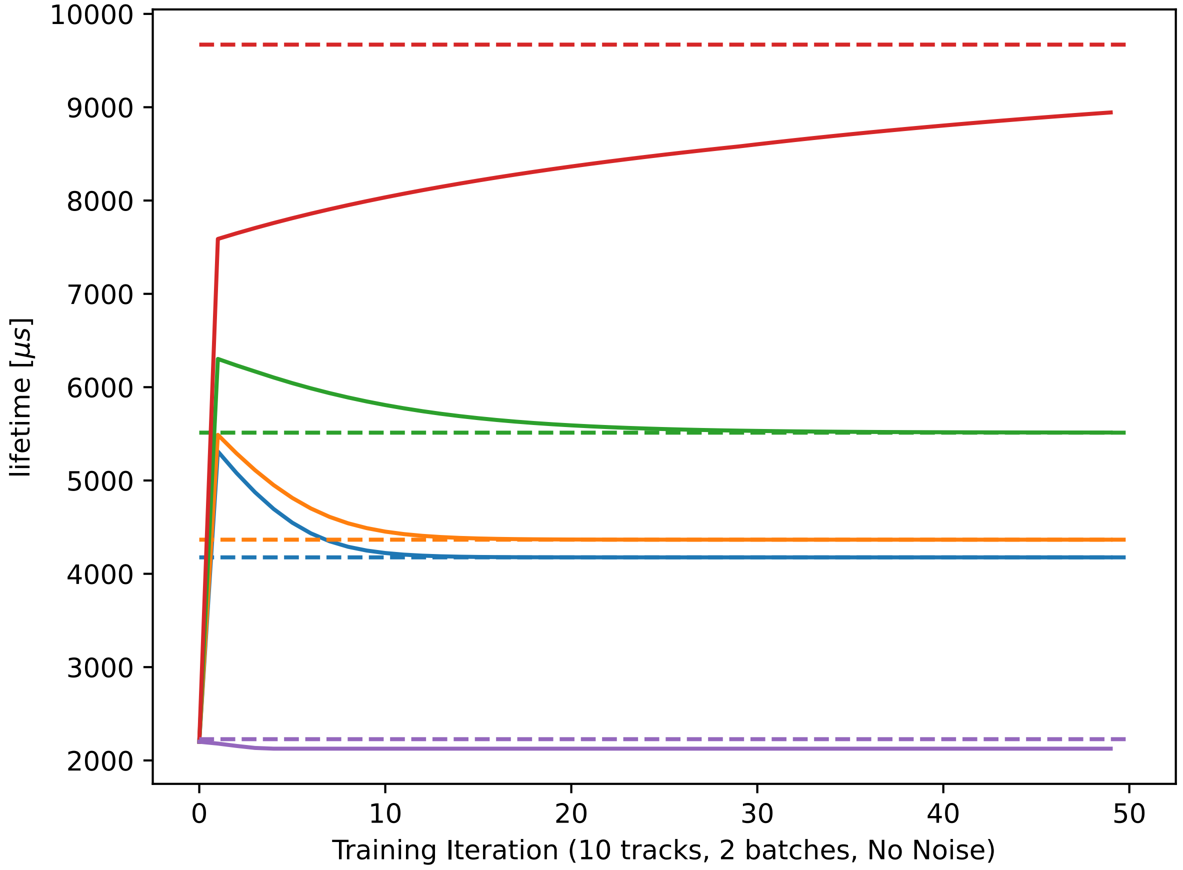

2D Fit: Ab + kb

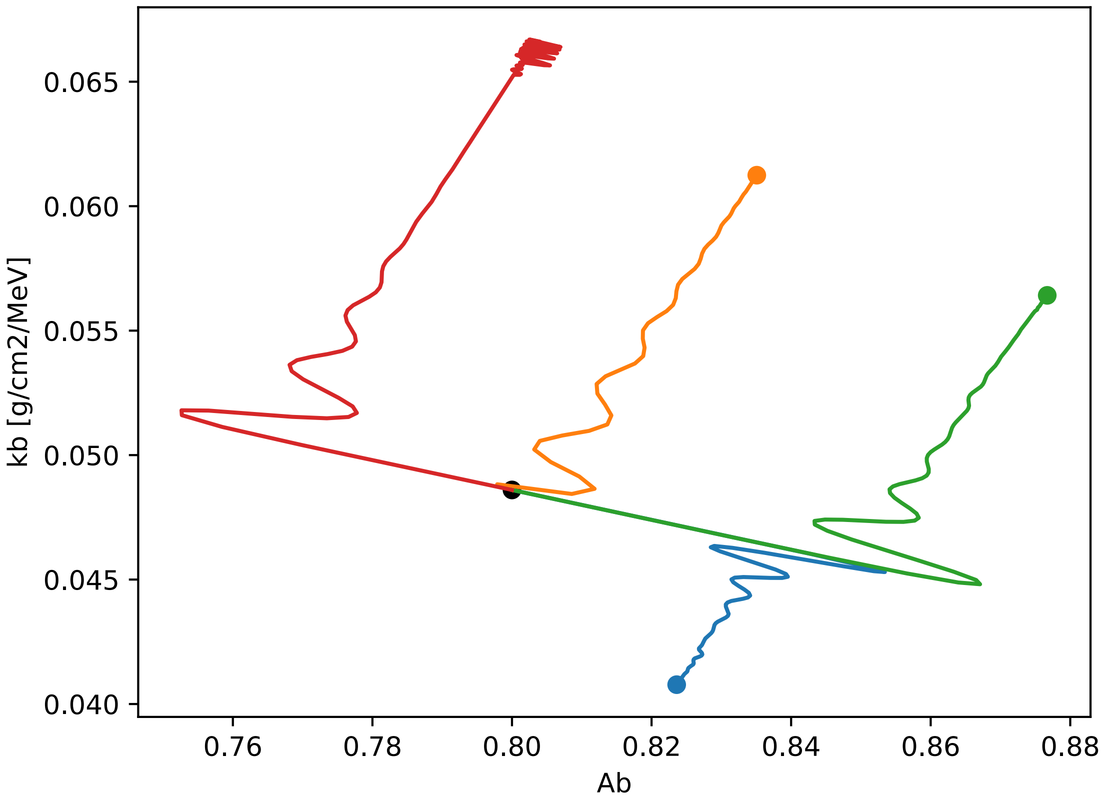

2D fit in Ab and kb, lr=1e-2 for both

- Same tracks as above studies
- Here fitting differences from nominal values instead of values themselves, but point being, fits converge (though red oscillates a bit around target)

SLAC

Closing Questions/Next Steps

What's going on with transverse diffusion?

- Can we understand this impact better? Are there ways to make the loss nicer (e.g. DTW)?

How do we square these results with the physics intuition constraints?

- Last week we said tran_diff was irrelevant. Seems like not so!
  - Check that impact is actually happening from edge/tail effects
  - Try ridiculously large values (regime where spread ~ pixel size) to see if there's a "nice" regime
- Studies from Yifan provide some constraints on lifetime
  - Check optimization in relevant (low) lifetime range

How do these results impact our scope?

- With noise, broad region of parameter space is washed out, some parameters (eField, long_diff) entirely below noise level, some parameters (e.g. lifetime) only have notable impact for very large changes
  - Maybe demo is Ab + kb (+ tran_diff?) and vdrift is fit independently
  - Can also include analysis of e.g. lifetime in region of expected sensitivity