

# On Estimating Gradients of Discrete Stochastic Programs for Detector Design Optimization

CHEP 2023

Lukas Heinrich, TUM  
Michael Kagan SLAC

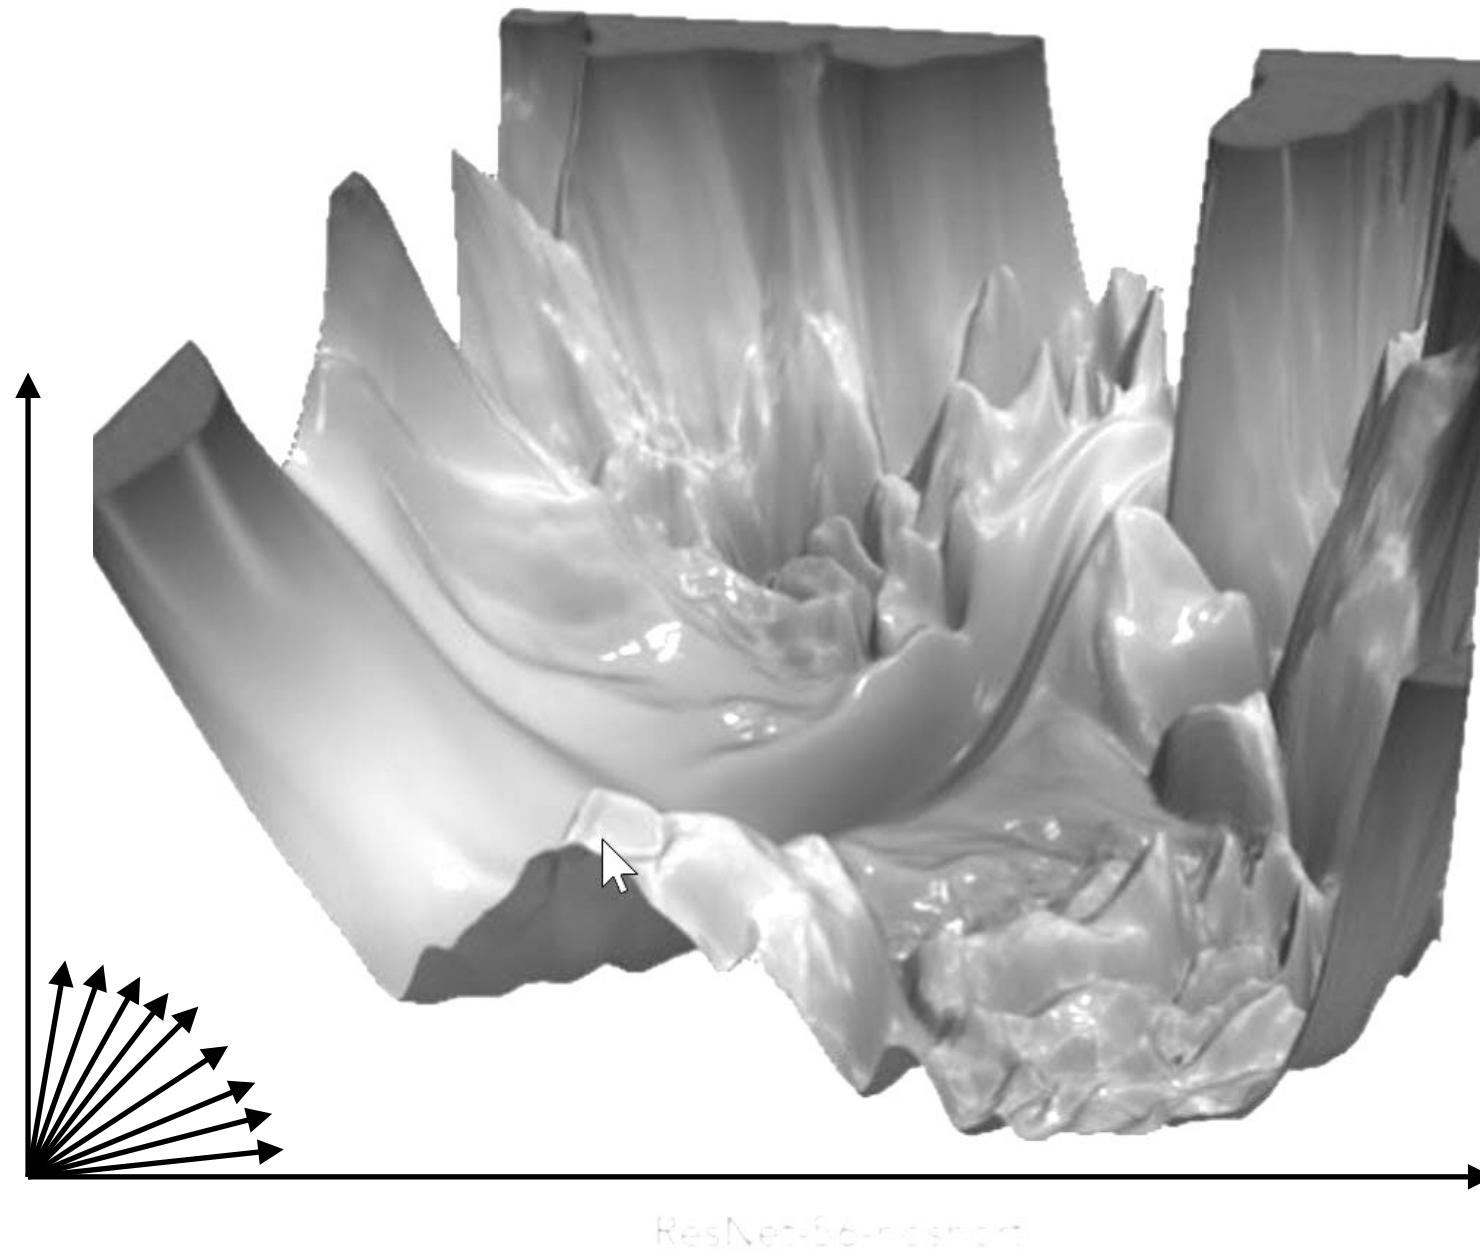
Technische  
Universität  
München



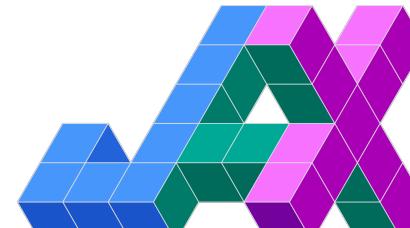
# Differentiable Programming

A major enabler of large-scale Machine Learning:

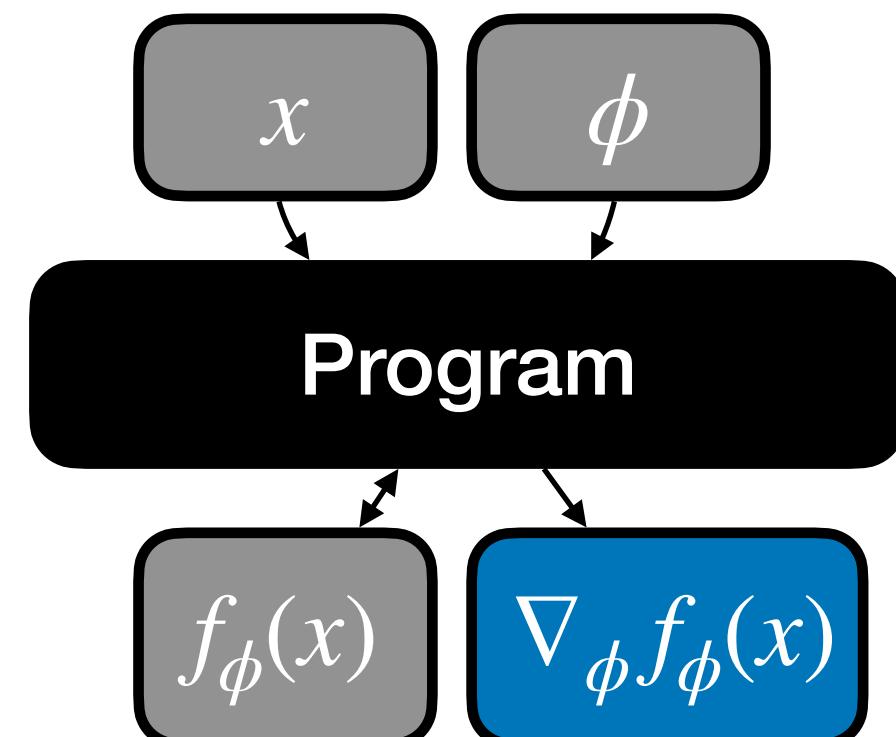
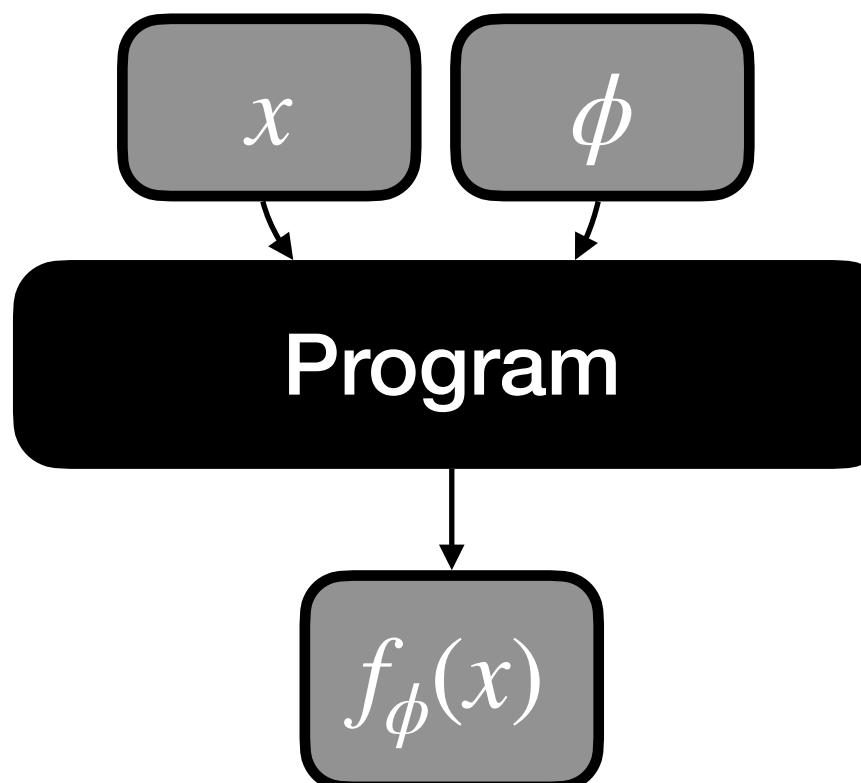
*Ability to compute gradients of numeric programs*



To deal with hyper-planes in a 14-dimensional space,  
visualize a 3D space and say 'fourteen' to yourself very loudly.  
-G. Hinton



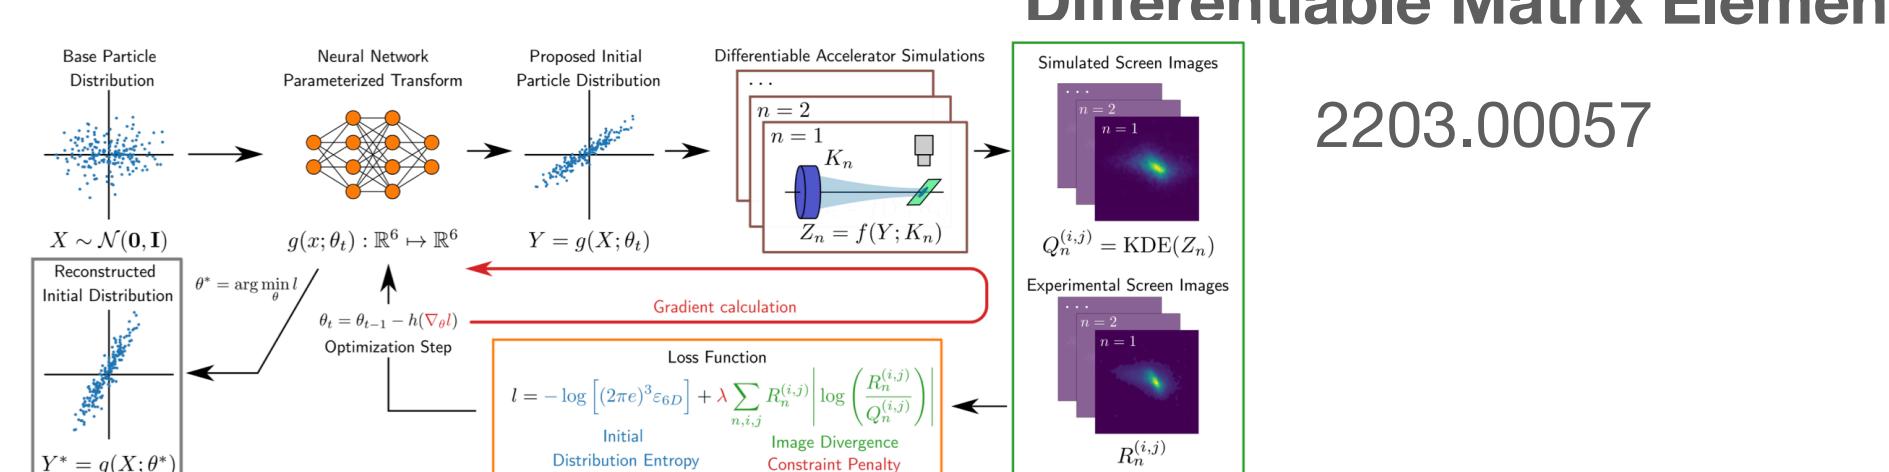
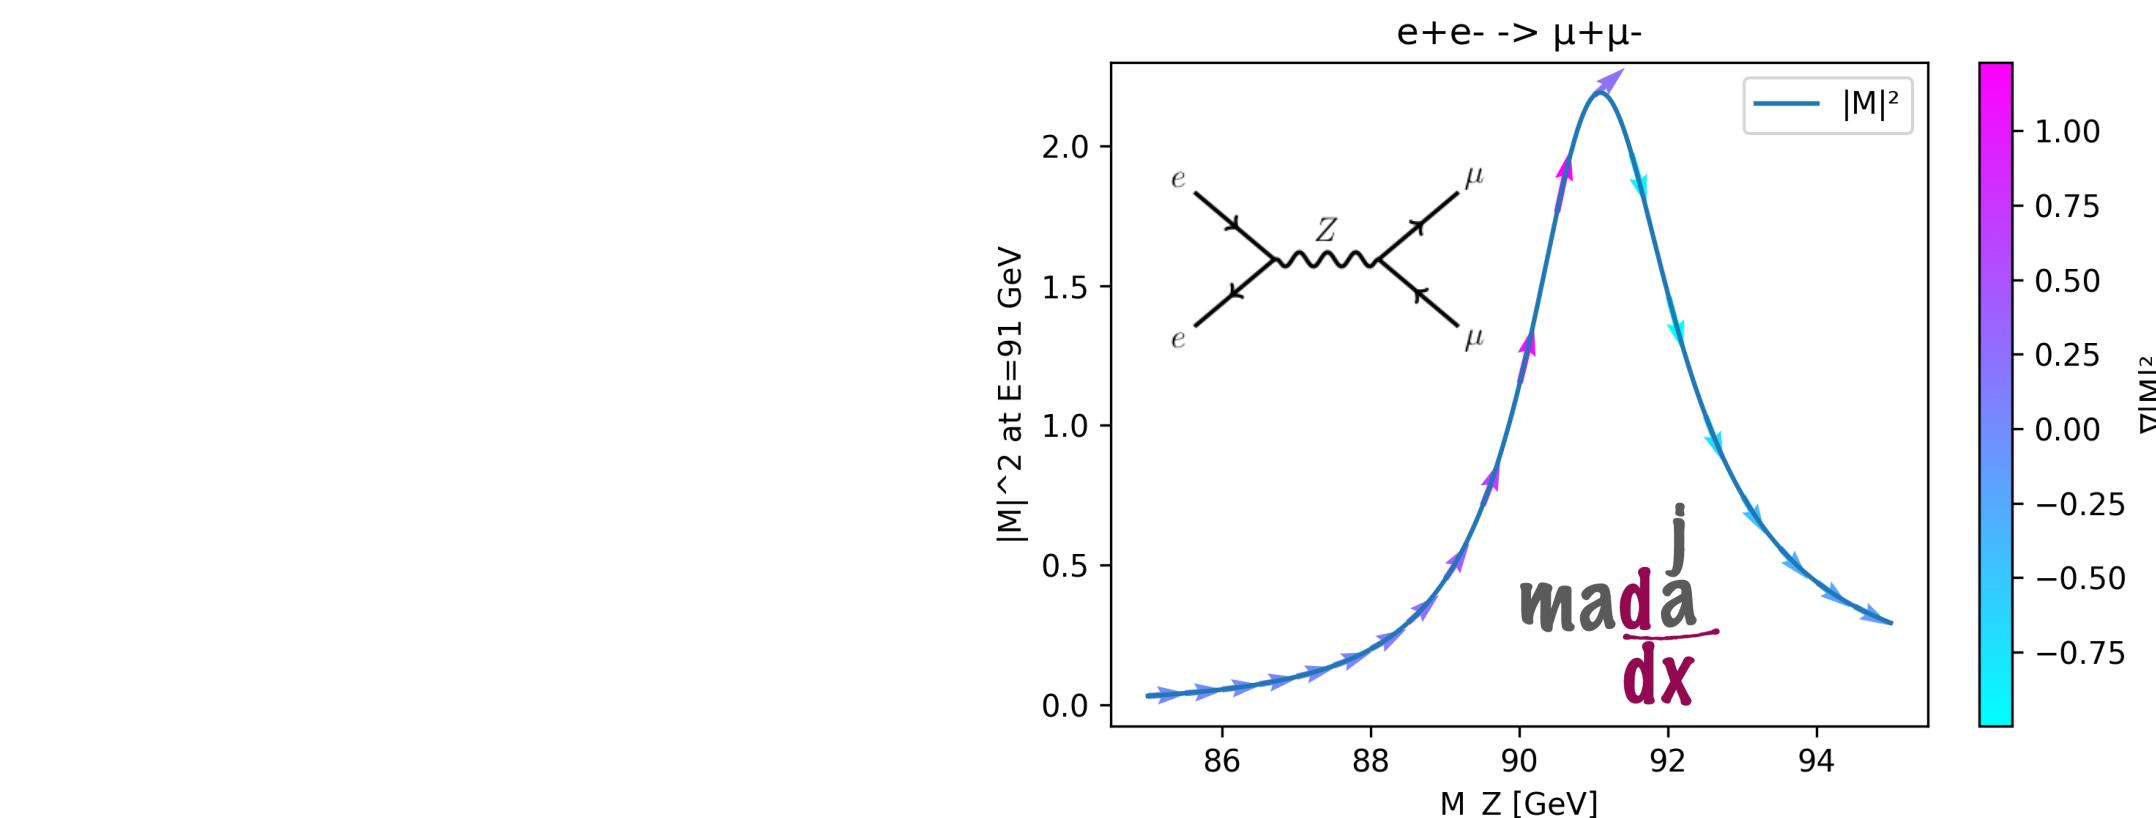
P Y T O R C H



# Gradients Beyond ML

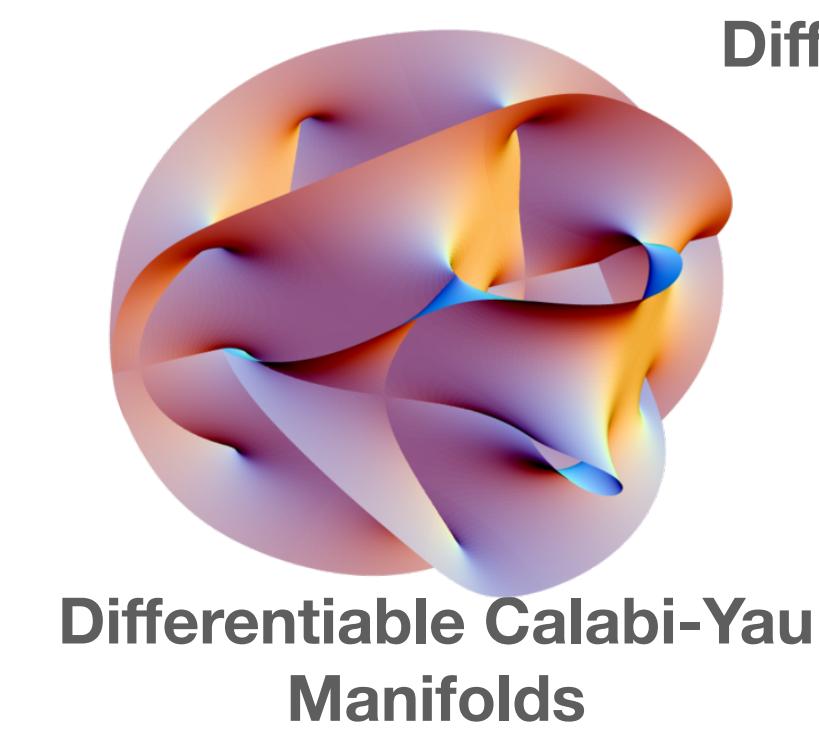
Interest in applying DP to programs that are not NNs:

*powerful method to combined physics + ML*



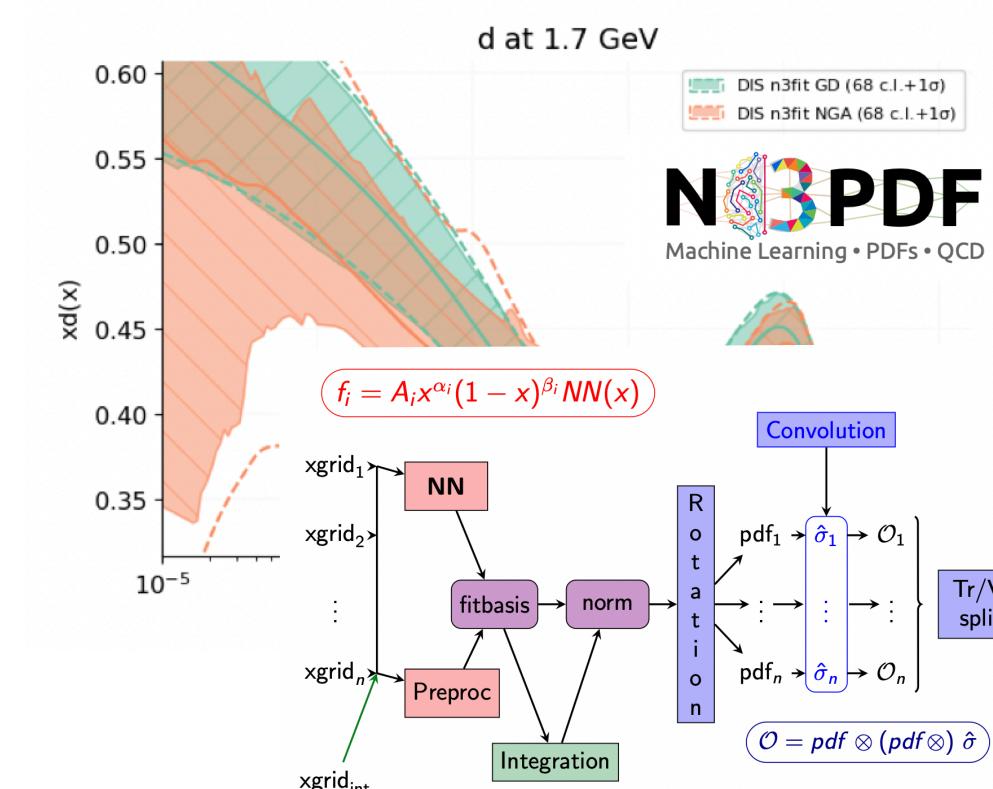
Differentiable Accelerator Simulation

2211.09077



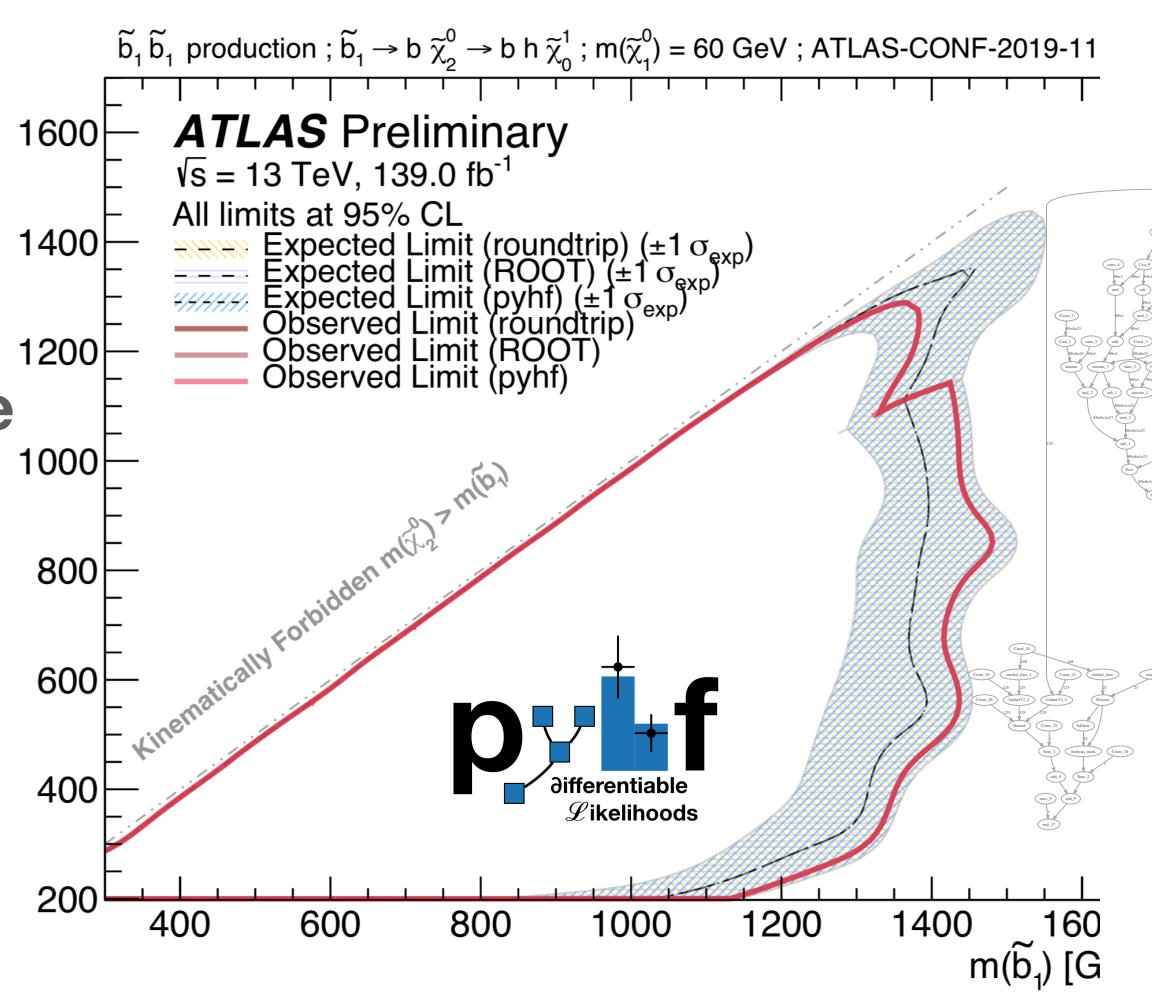
Differentiable Calabi-Yau Manifolds

2211.12520



Differentiable Proton Structure

1907.05075



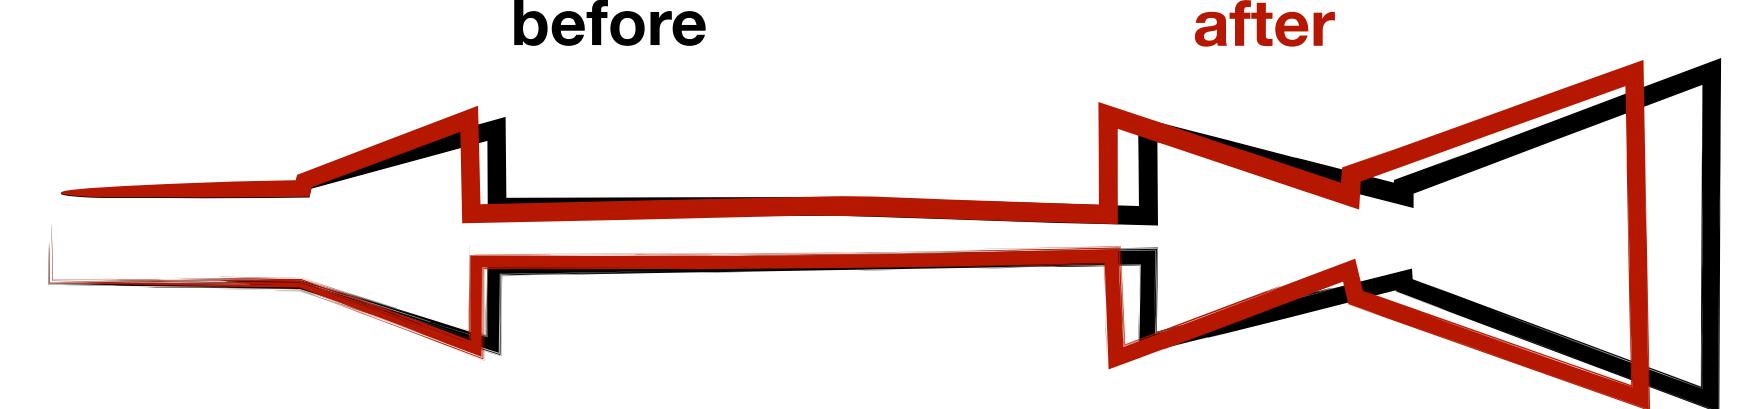
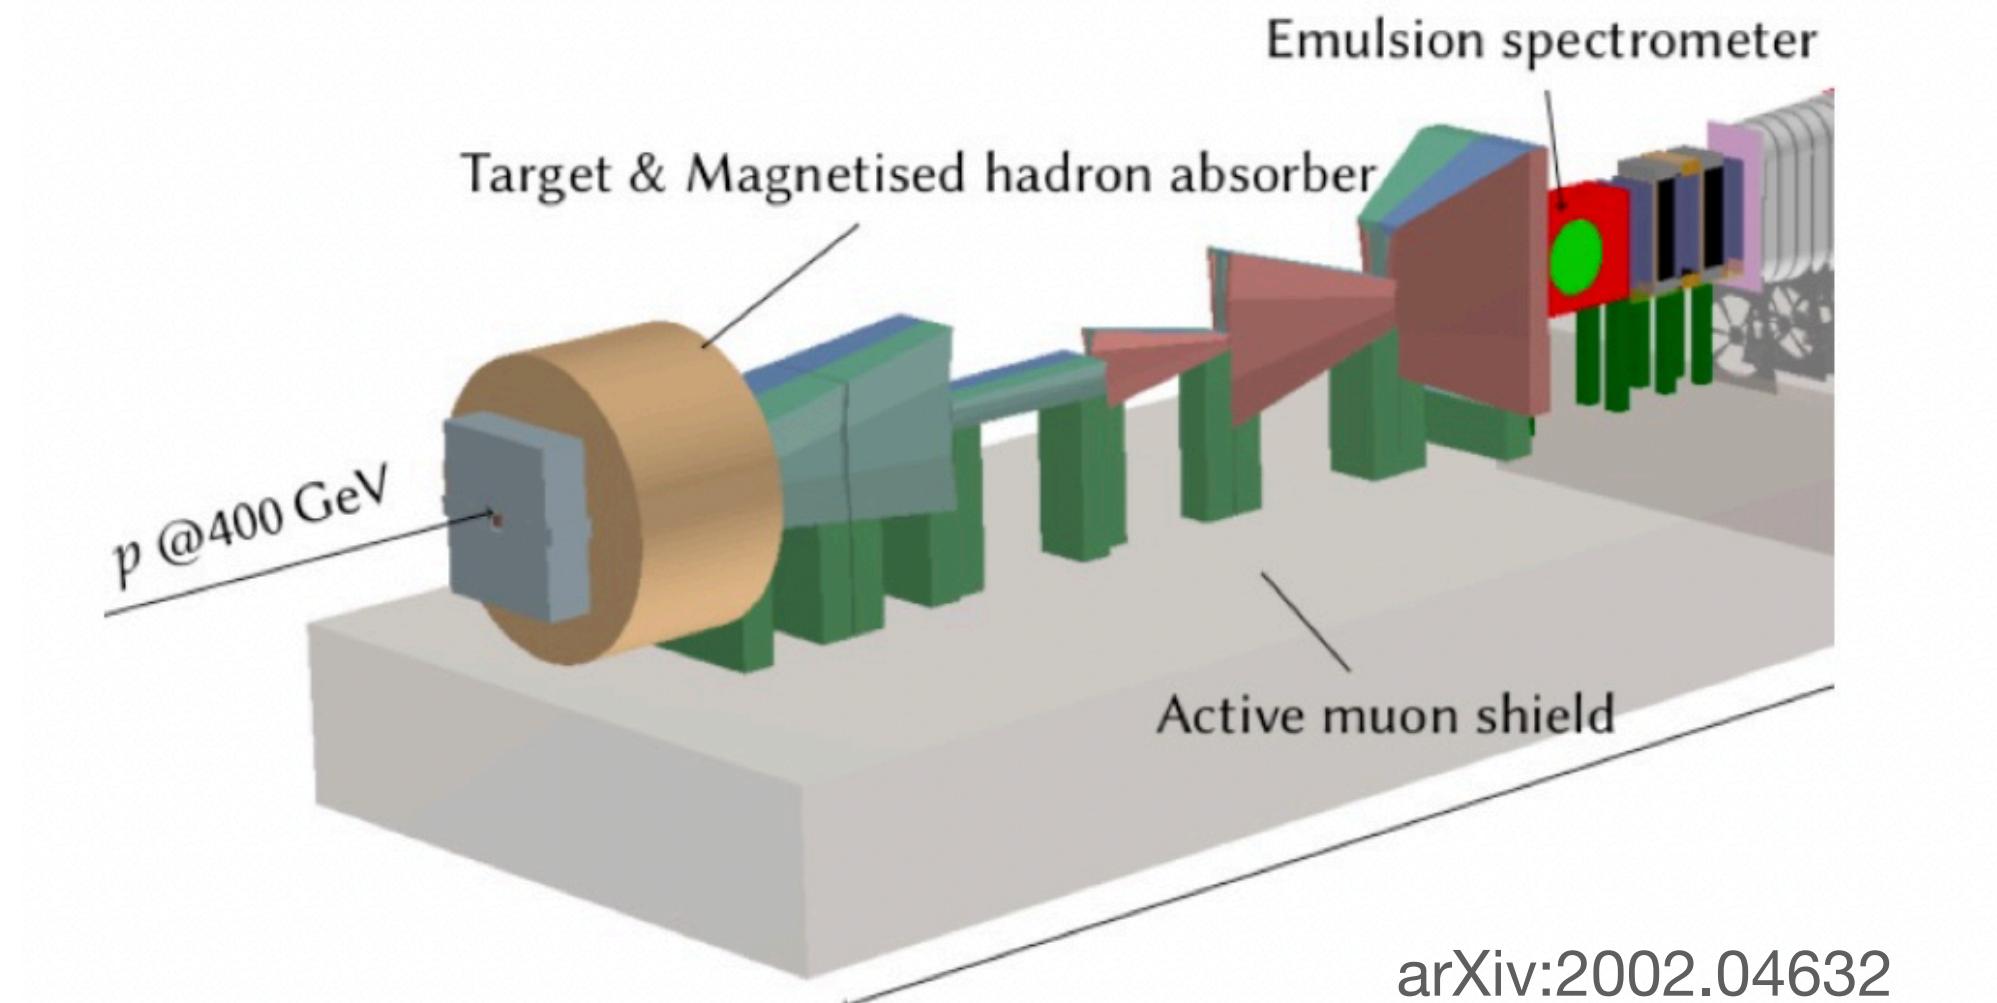
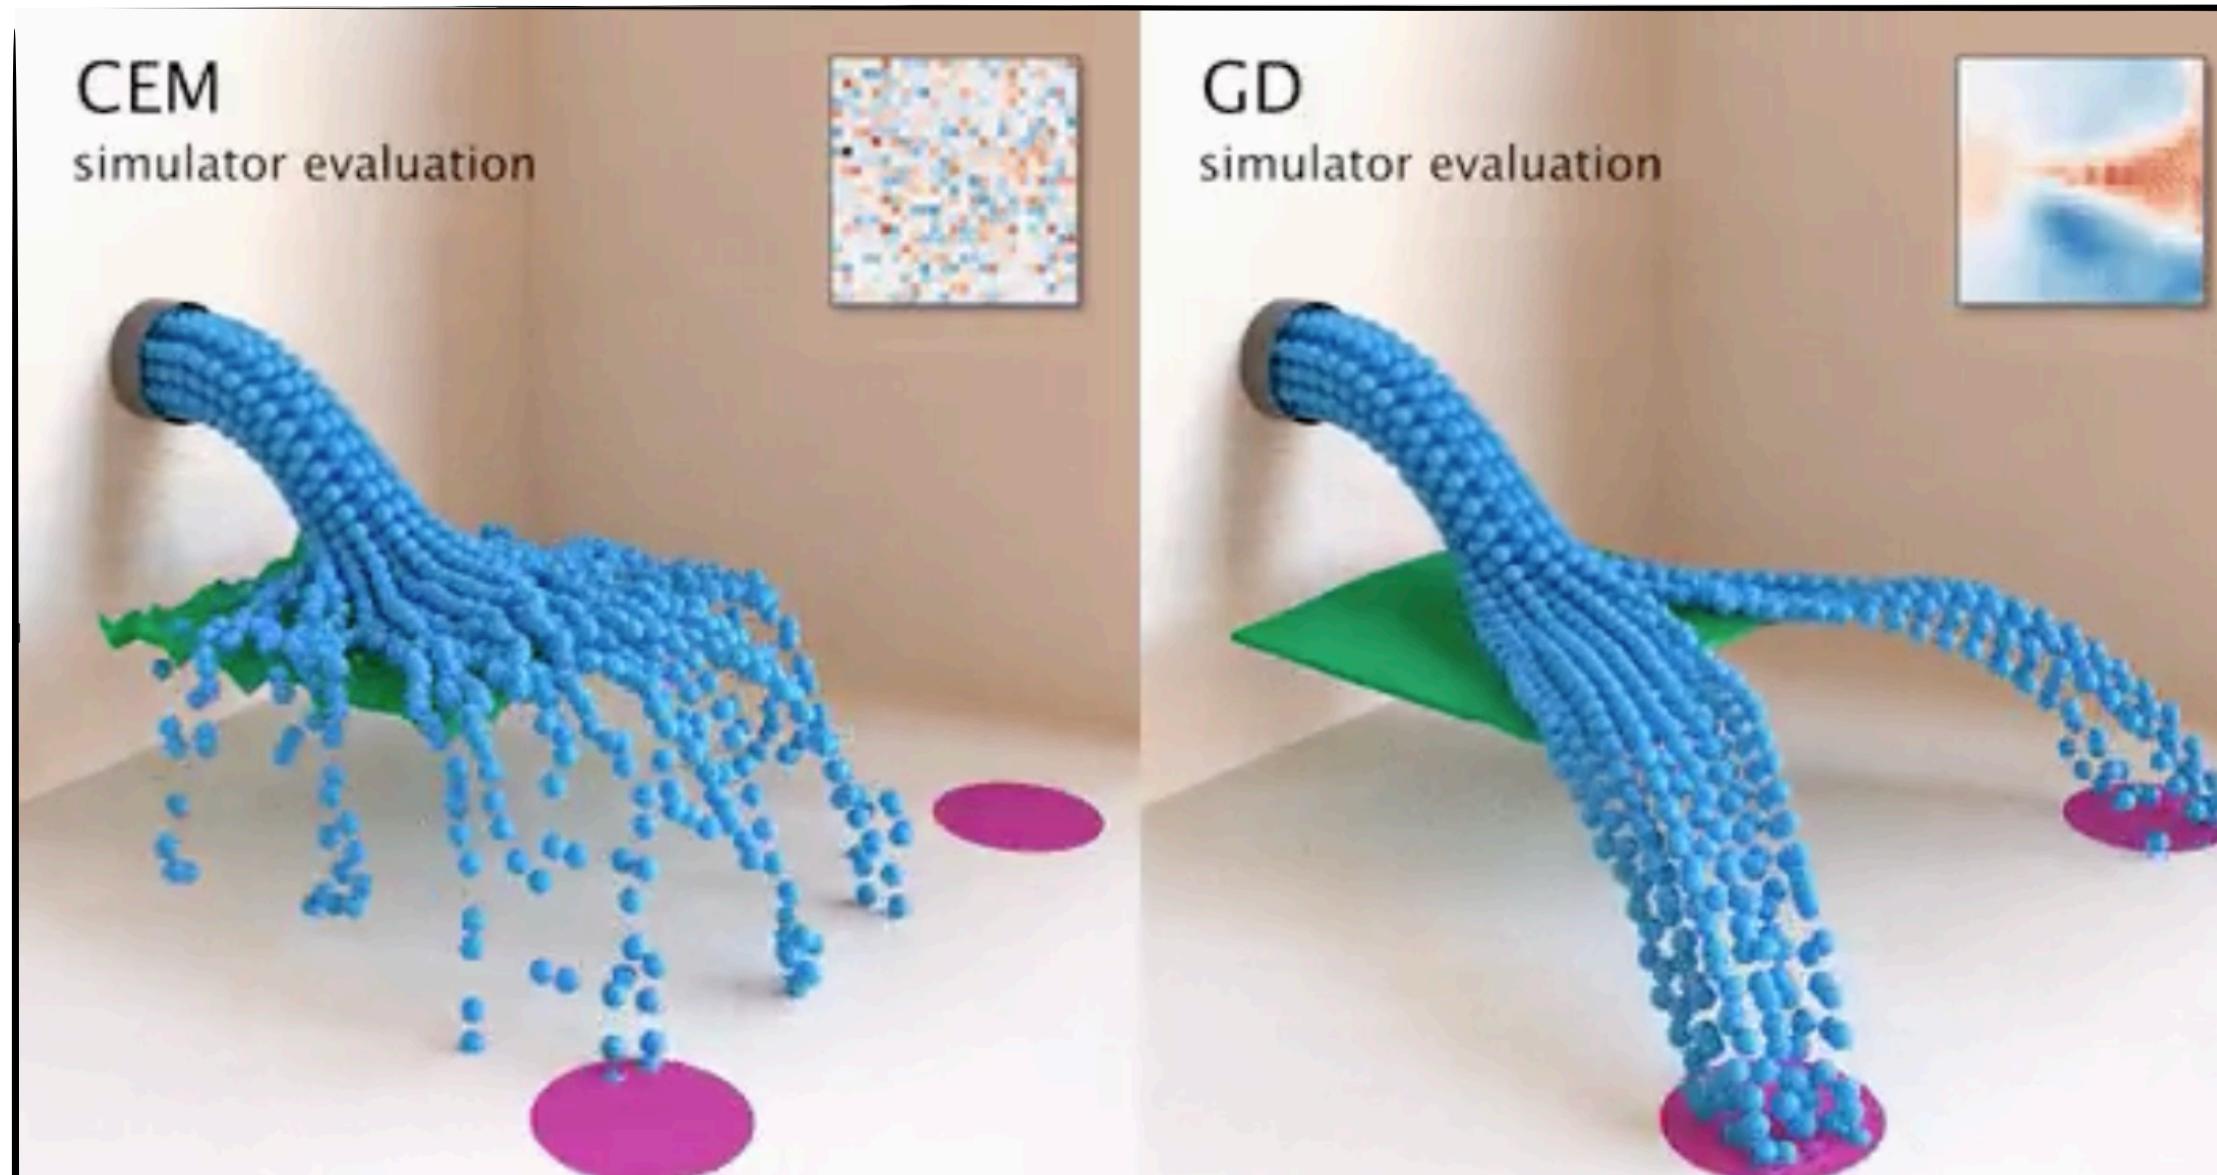
Differentiable Inference

10.21105/joss.02823

# Emerging Usecase: Design Optimization

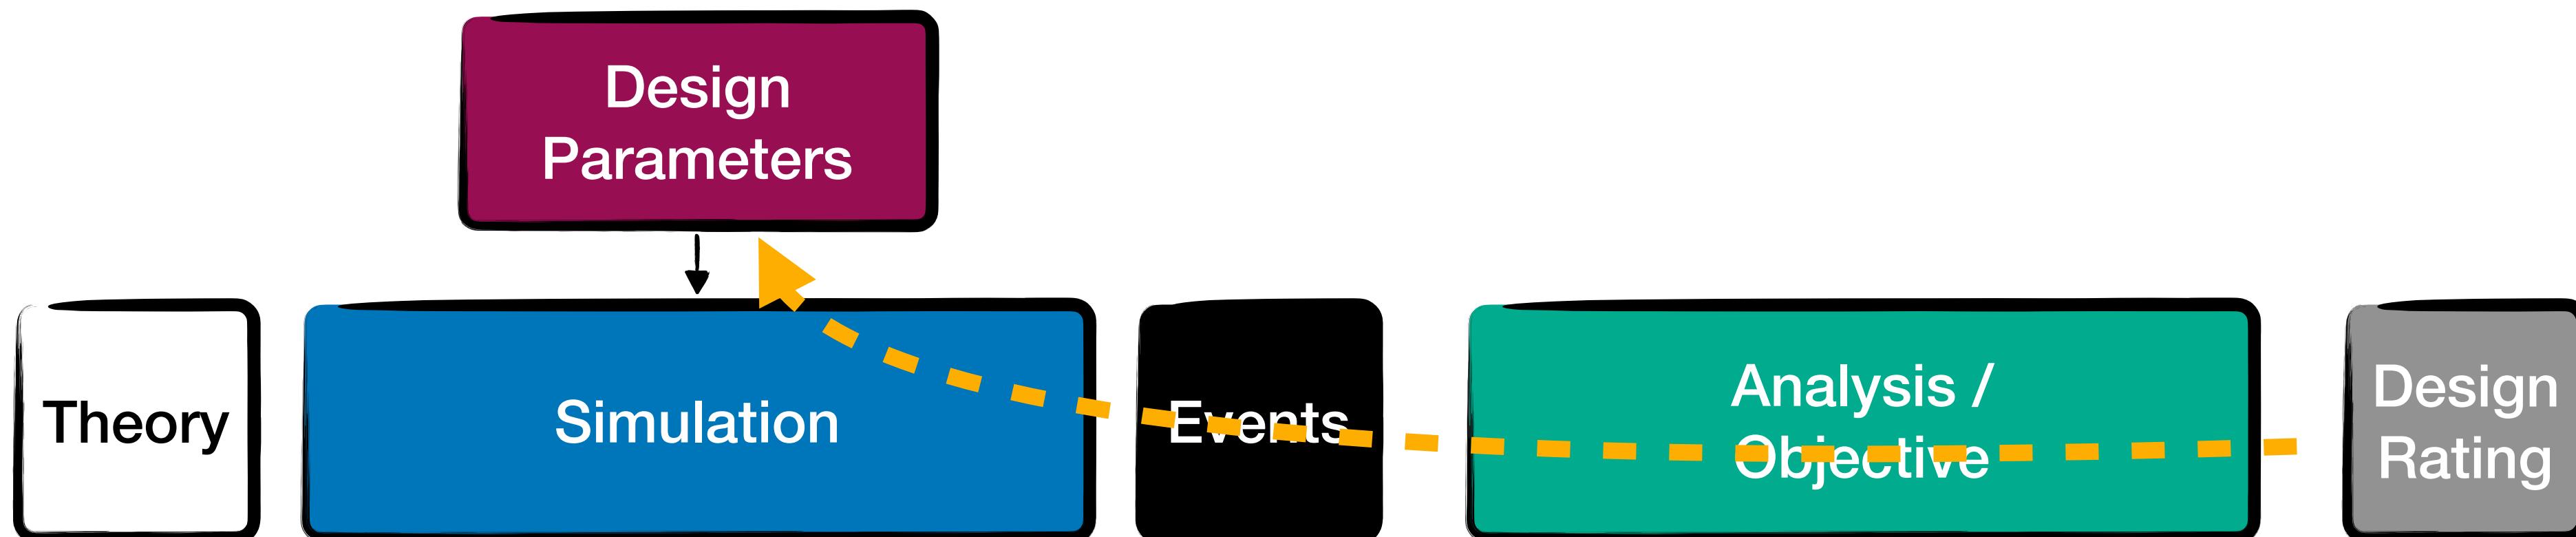
One of the most consequential high-dim. optimizations:

*find a good detector design*



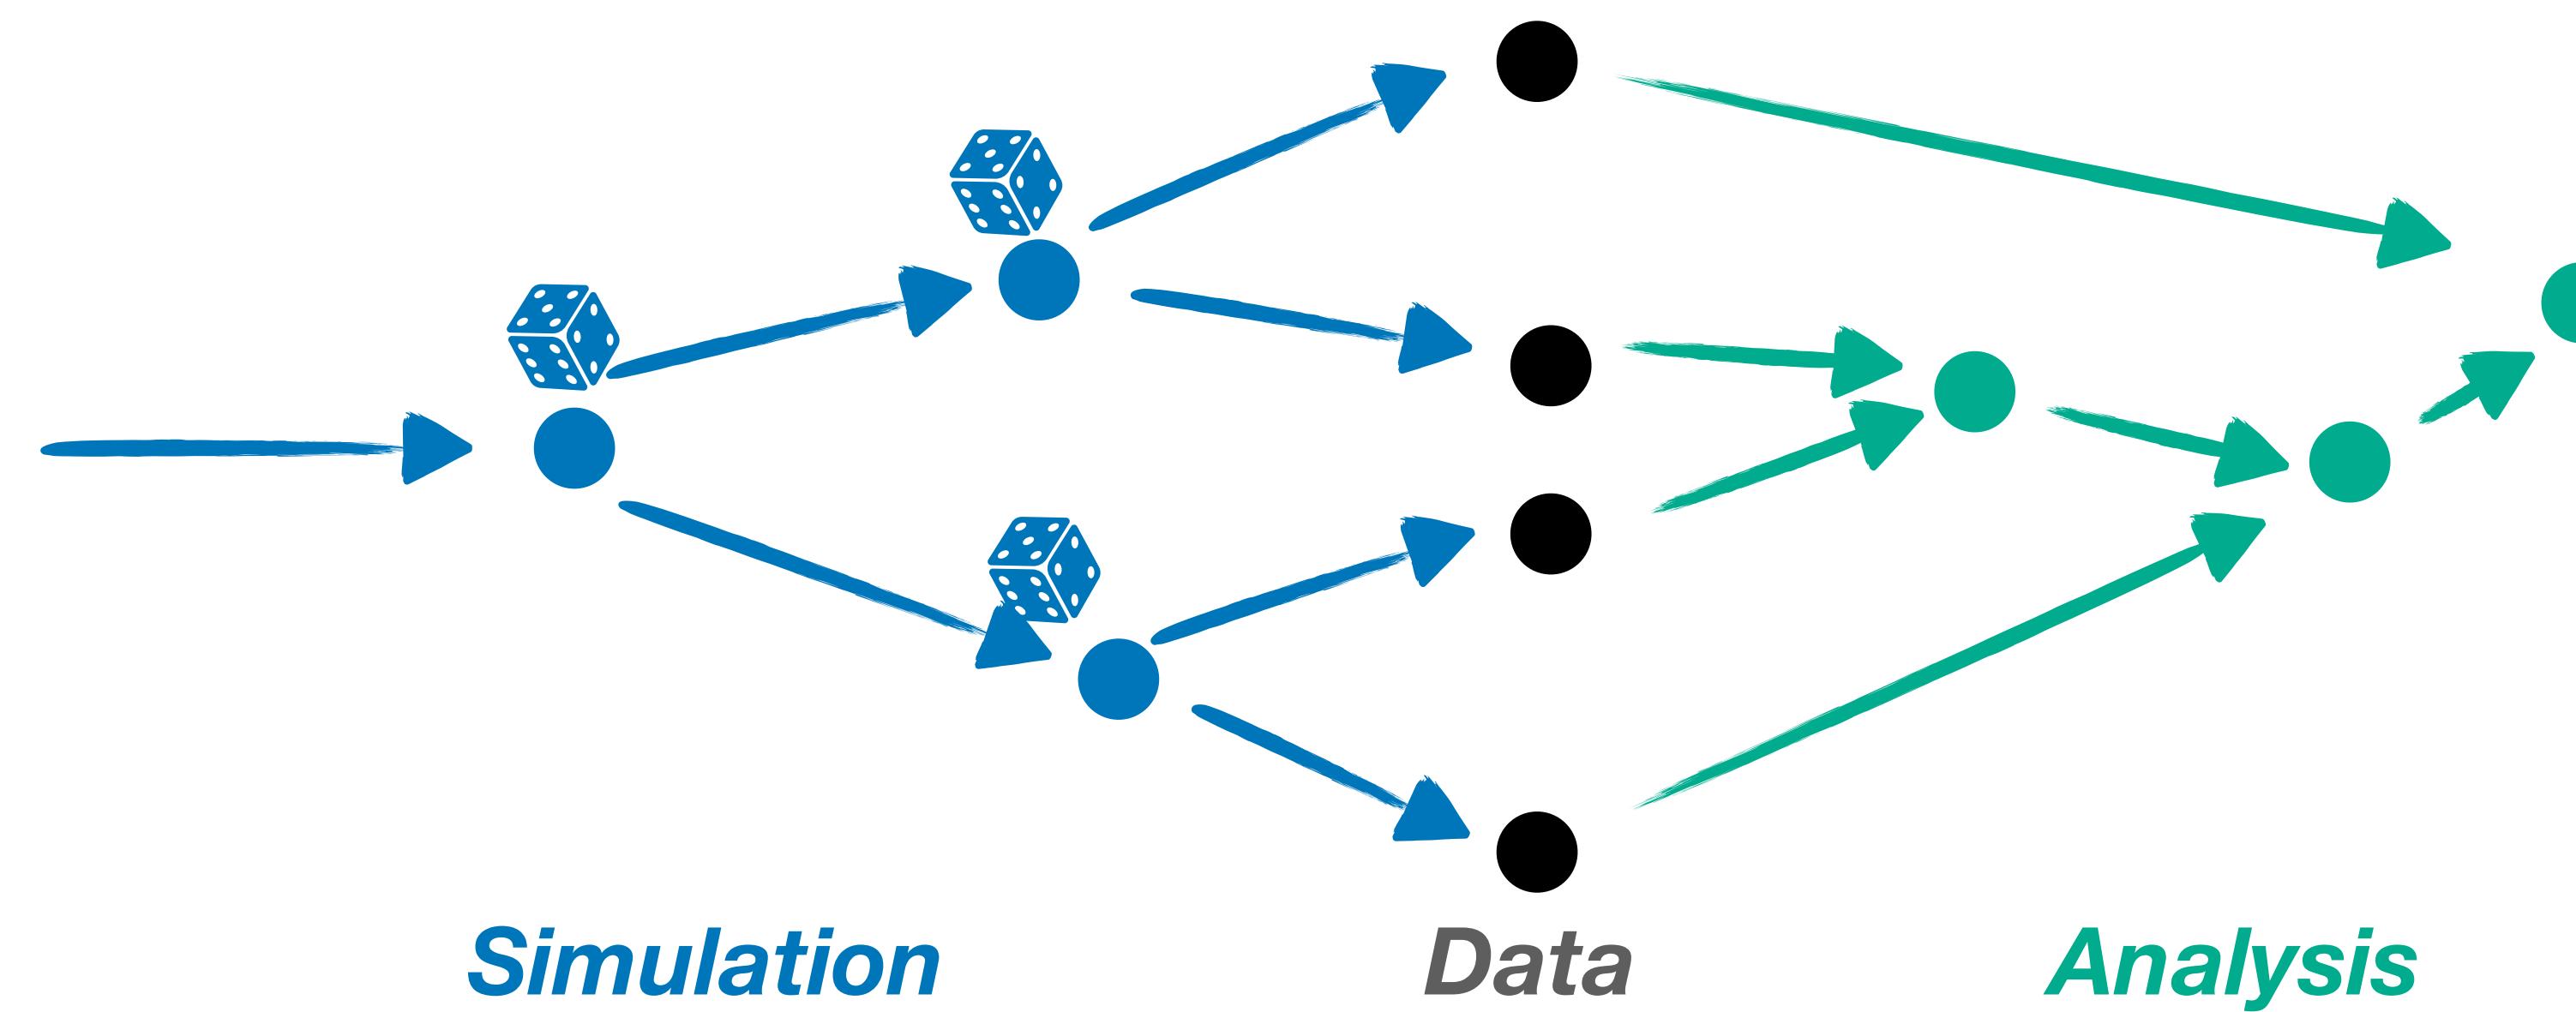
# A Requirement:

For a gradient-based design optimization, one needs both:  
*a differentiable simulation and reco/analysis*



# A Small Problem

To first order, most of simulation is **stochastic branching**, while most of analysis is **iterative clustering** with many discrete choices



*HEP: fundamentally incompatible with differentiable programming?*

# Not necessarily

The gradients we need are gradients of **expectation values** over probability distributions & observables with discrete aspects

$$\nabla_{\phi} \bar{f}(\phi) = \nabla_{\phi} \mathbb{E}_{p_{\phi}} [f] = \nabla_{\phi} \int dx \ f(x, \phi) p_{\phi}(x)$$

The diagram shows the mathematical expression for the expectation value gradient. A green arrow points from the term  $f(x, \phi)$  to a teal box labeled "non-differentiable evaluation function". A blue arrow points from the term  $p_{\phi}(x)$  to a pink box labeled "discrete data  $x$ ". A blue curved arrow points from the pink box back to the integral, labeled "probability density".

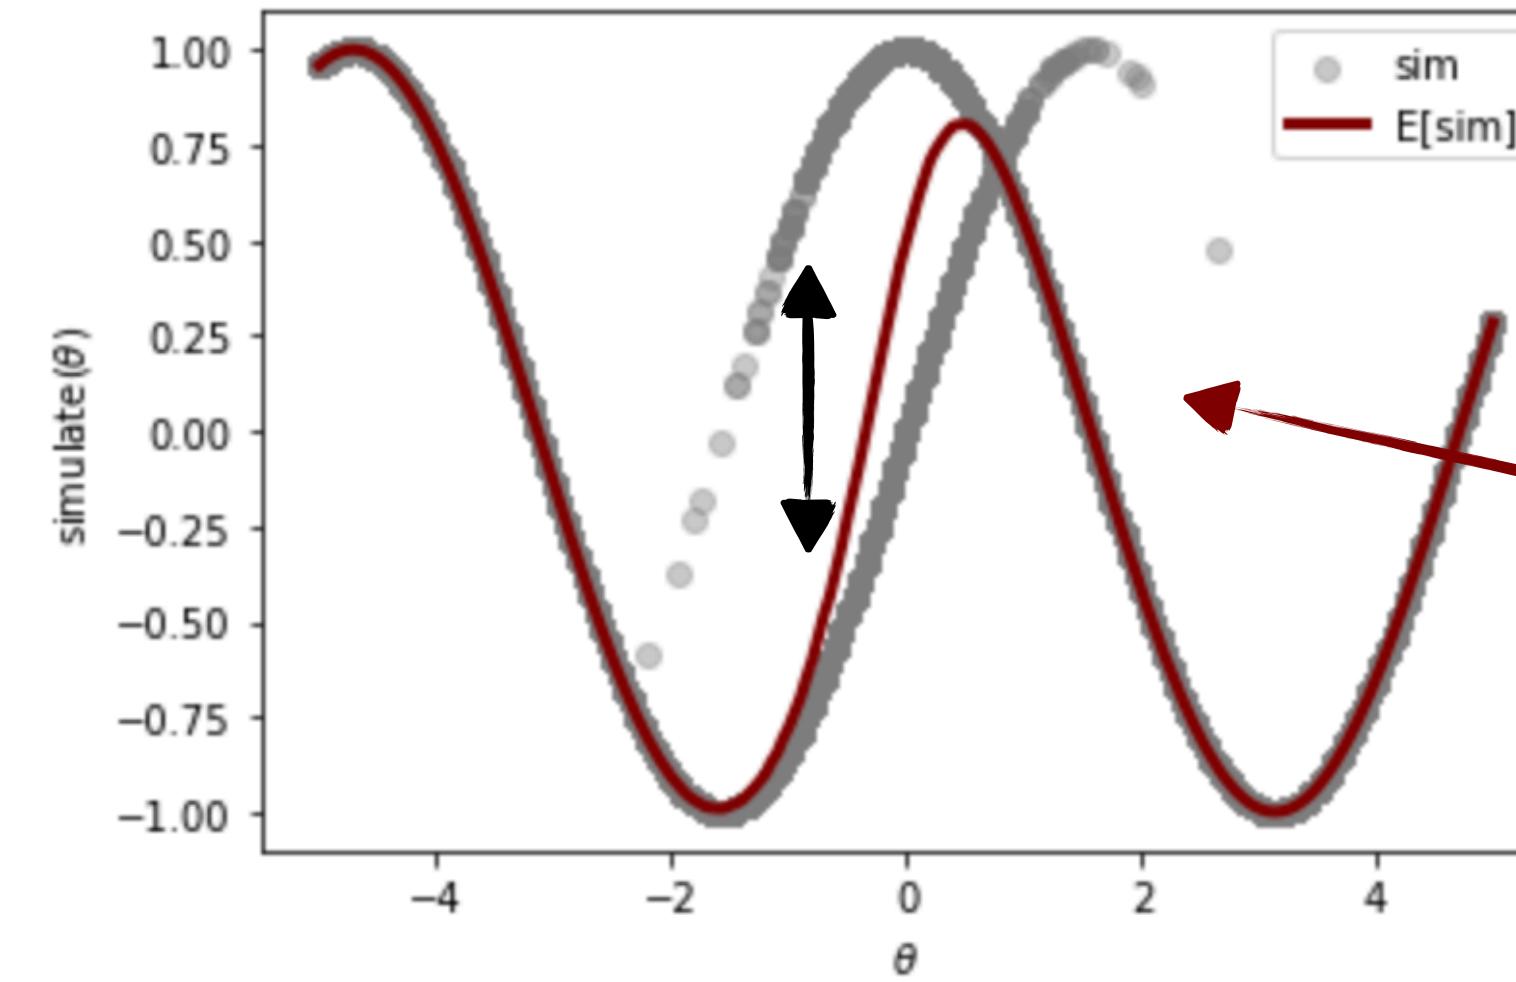
In many cases, the expectation value *is* differentiable  
(just not via standard automatic differentiation)

# Simple Examples

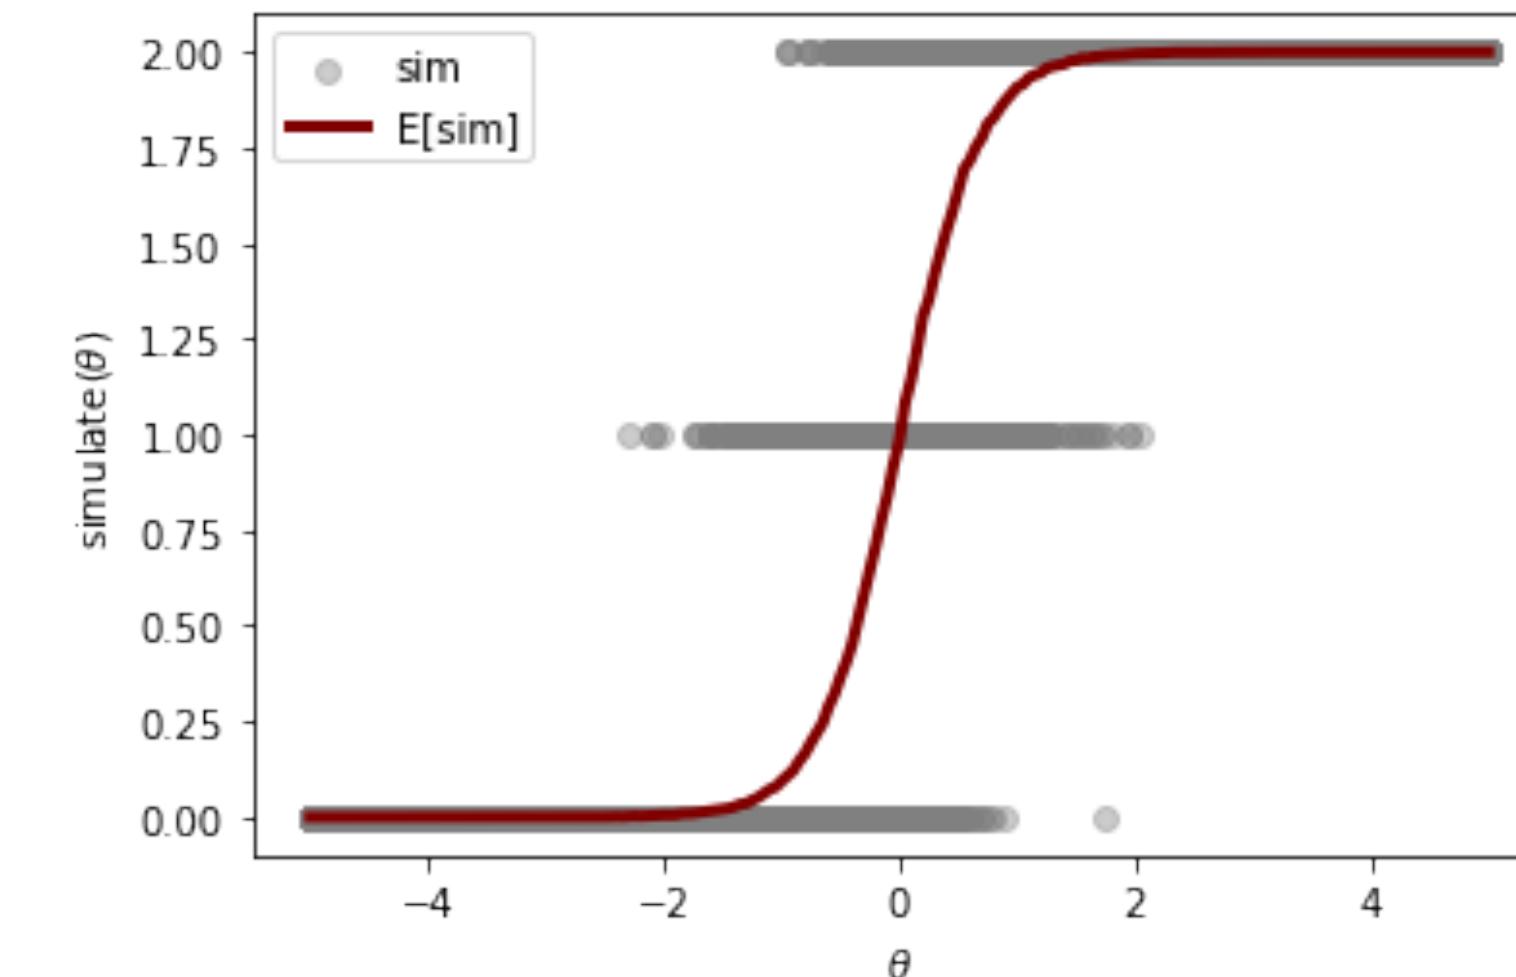
```
def simulate(theta):
    p = sigmoid(theta)
    x = bernoulli(p) #0 or 1
    if x == 0:
        eval = sin(theta)
    else:
        eval = cos(theta)
    return eval
```

```
def program(theta):
    p = sigmoid(theta)
    x1 = bernoulli(p) #0/1
    x2 = bernoulli(p) #0/1
    eval = x1 + x2 # 0, 1 or 2
    return eval
```

Discrete Jumps



Smooth Expectation Value

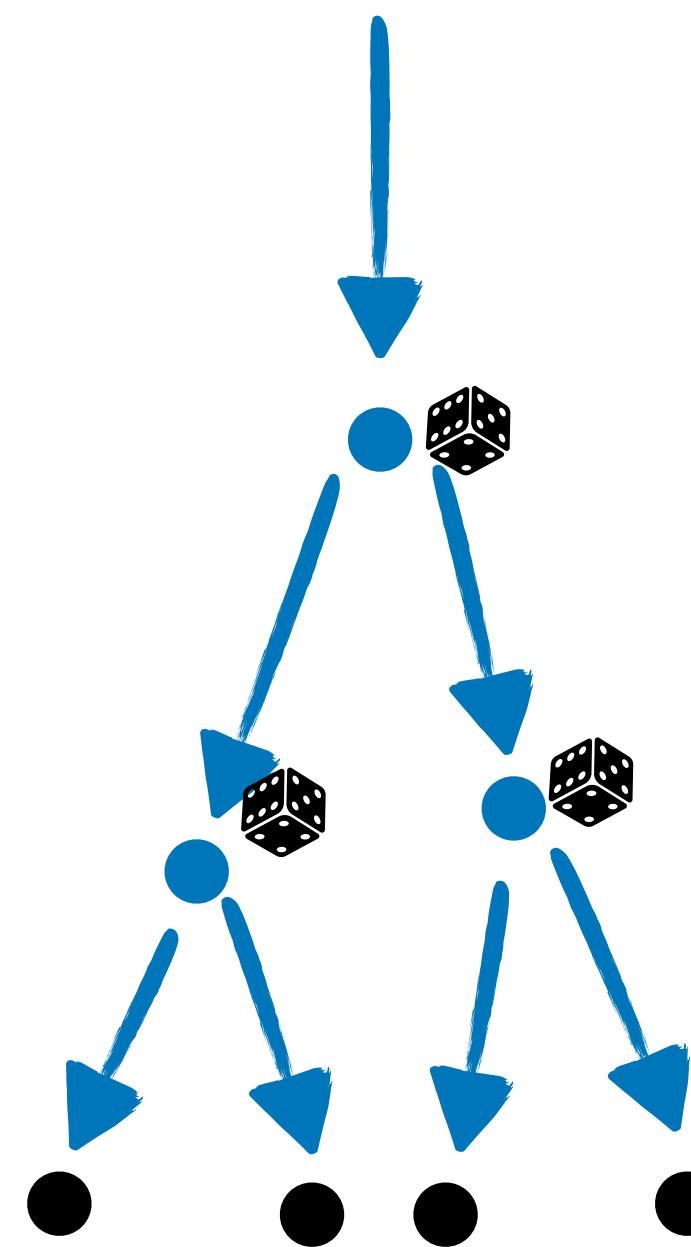


# Differentiating Expectation Values

The best known algorithm for gradient estimation is used a lot in *Reinforcement Learning - score function estimation*

$$\nabla \bar{f}(\phi) = \mathbb{E}_{p_\phi}[f(x) \nabla_\phi \log p_\phi(x)]$$

requires tracking probabilities (and their gradients) while running code  
→ probabilistic programming



HEP Simulator: discrete processes



Atari Games discrete actions



# Differentiating Expectation Values

More recently, renewed interest in approaches based on  
“smoothed perturbation analysis”

Extend fwd-mode AD to discrete-stochastic environments

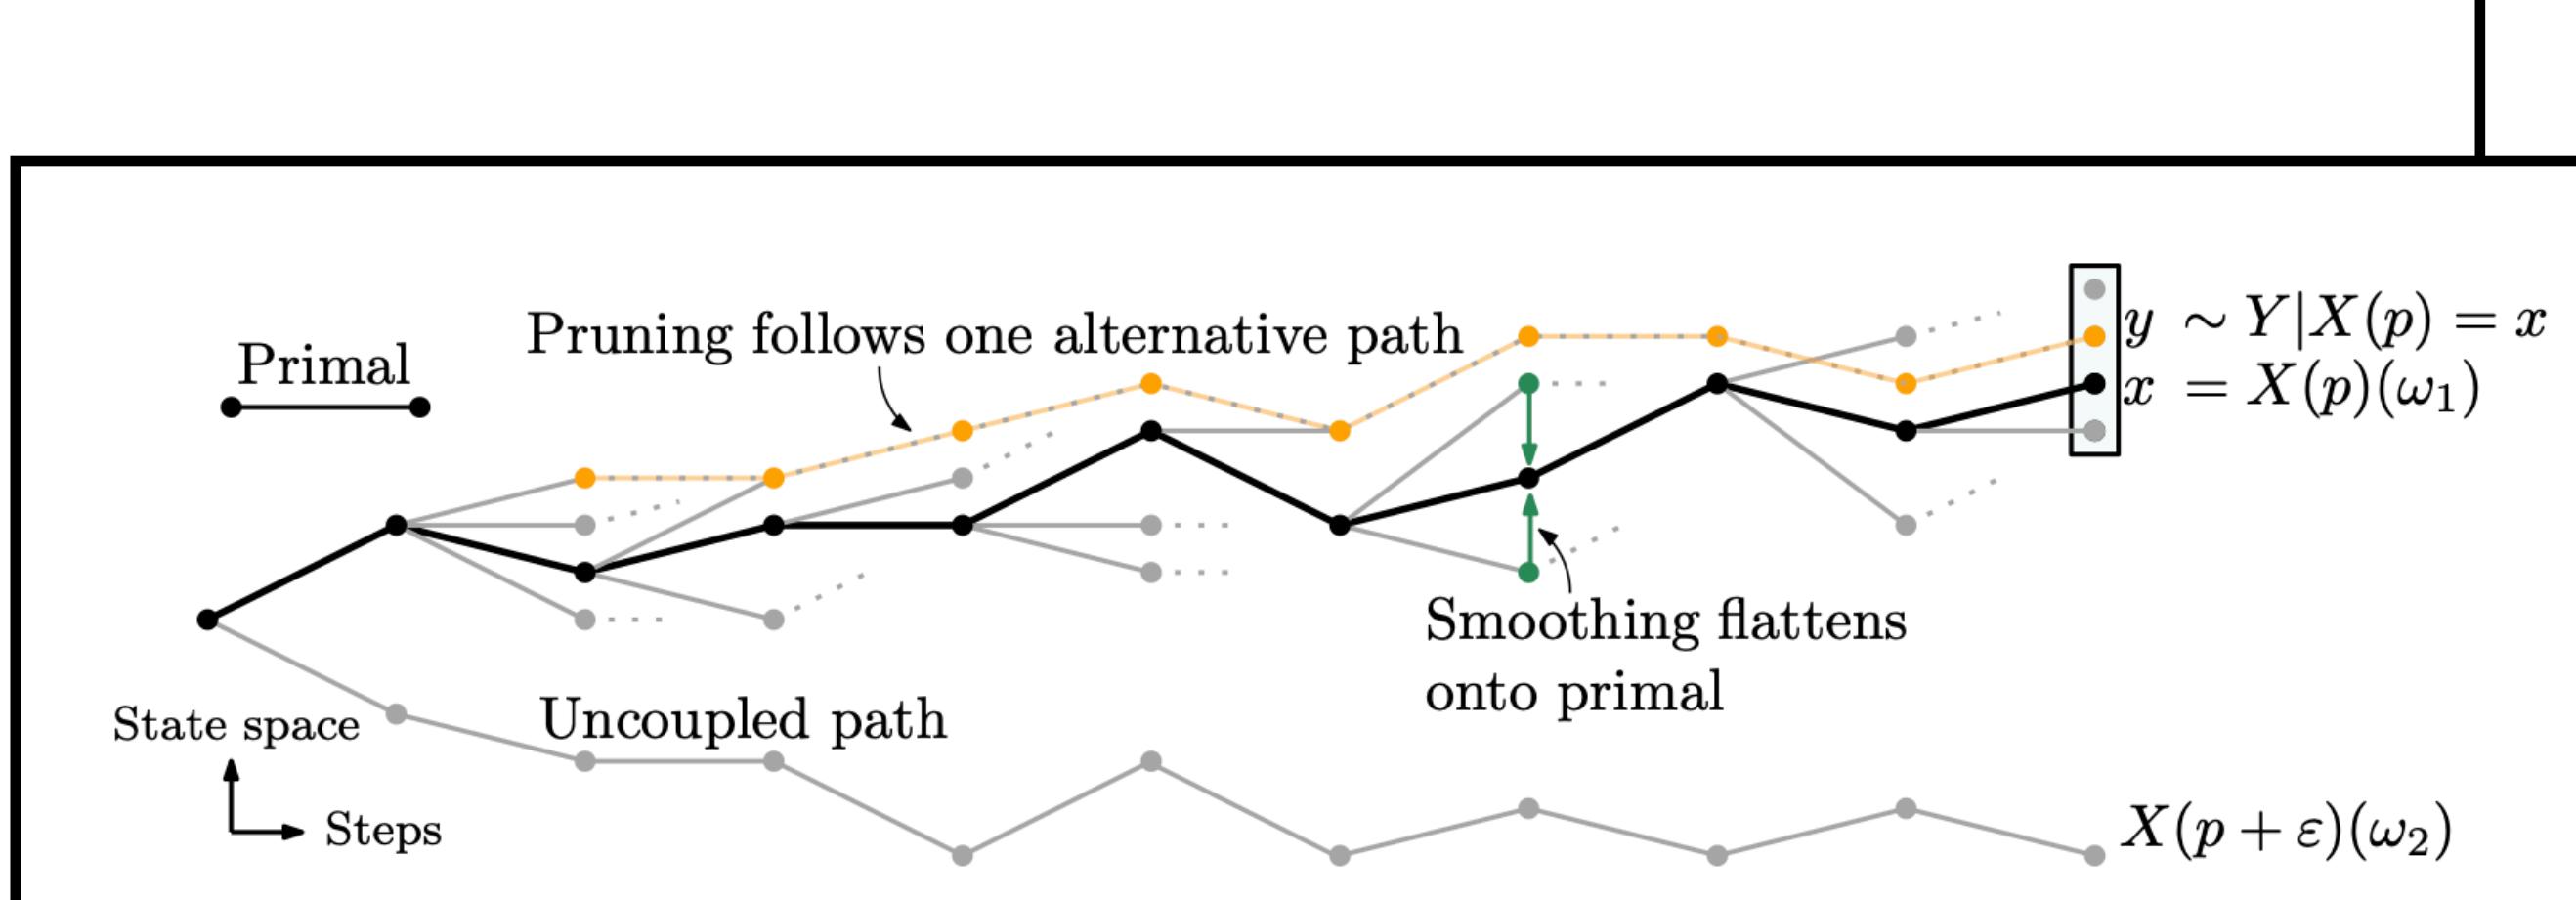
$$(x, \delta) \rightarrow (x, \delta, w, y)$$

Dual Numbers  
in Standard AD

Quadruplets  
in Stochastic AD

$$\nabla_{\phi} \bar{f}(\phi) = \mathbb{E}_{p_{\phi}} [\delta + w(x - y)]$$

arxiv: 2210.08572



## Automatic Differentiation of Programs with Discrete Randomness

Gaurav Arya  
Massachusetts Institute of Technology, USA  
aryag@mit.edu

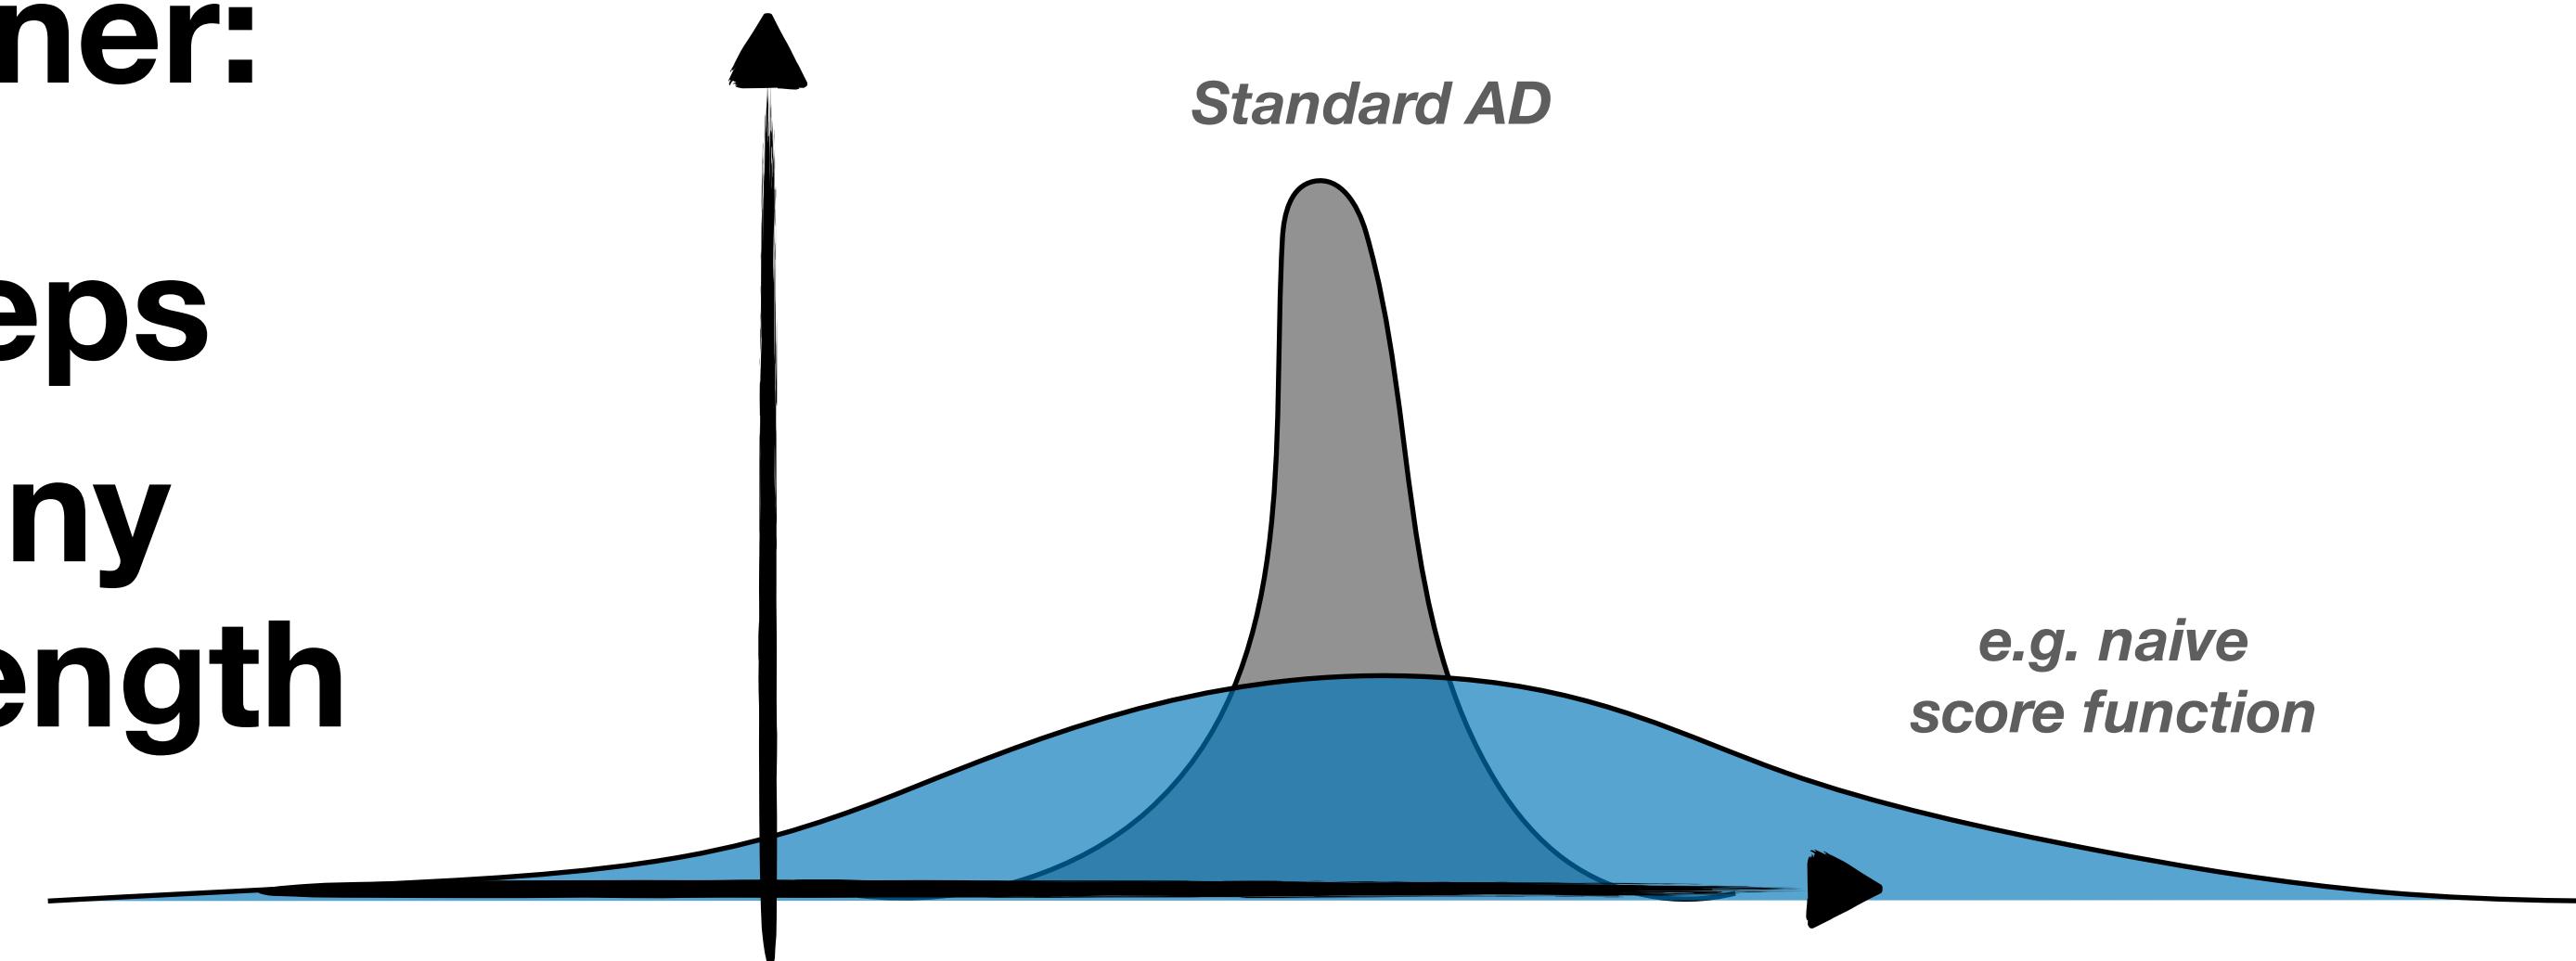
Moritz Schauer  
Chalmers University of Technology, Sweden  
University of Gothenburg, Sweden  
smoritz@chalmers.se

# Key Questions in HEP

Gradients Estimation in Stochastic Programs can suffer from very high variance → unstable optimization

HEP is in an interesting corner:

- a lot of tiny stochastic steps
- branching processes, many independently, variable-length processes



**Q: Will these methods work in a HEP setting?**

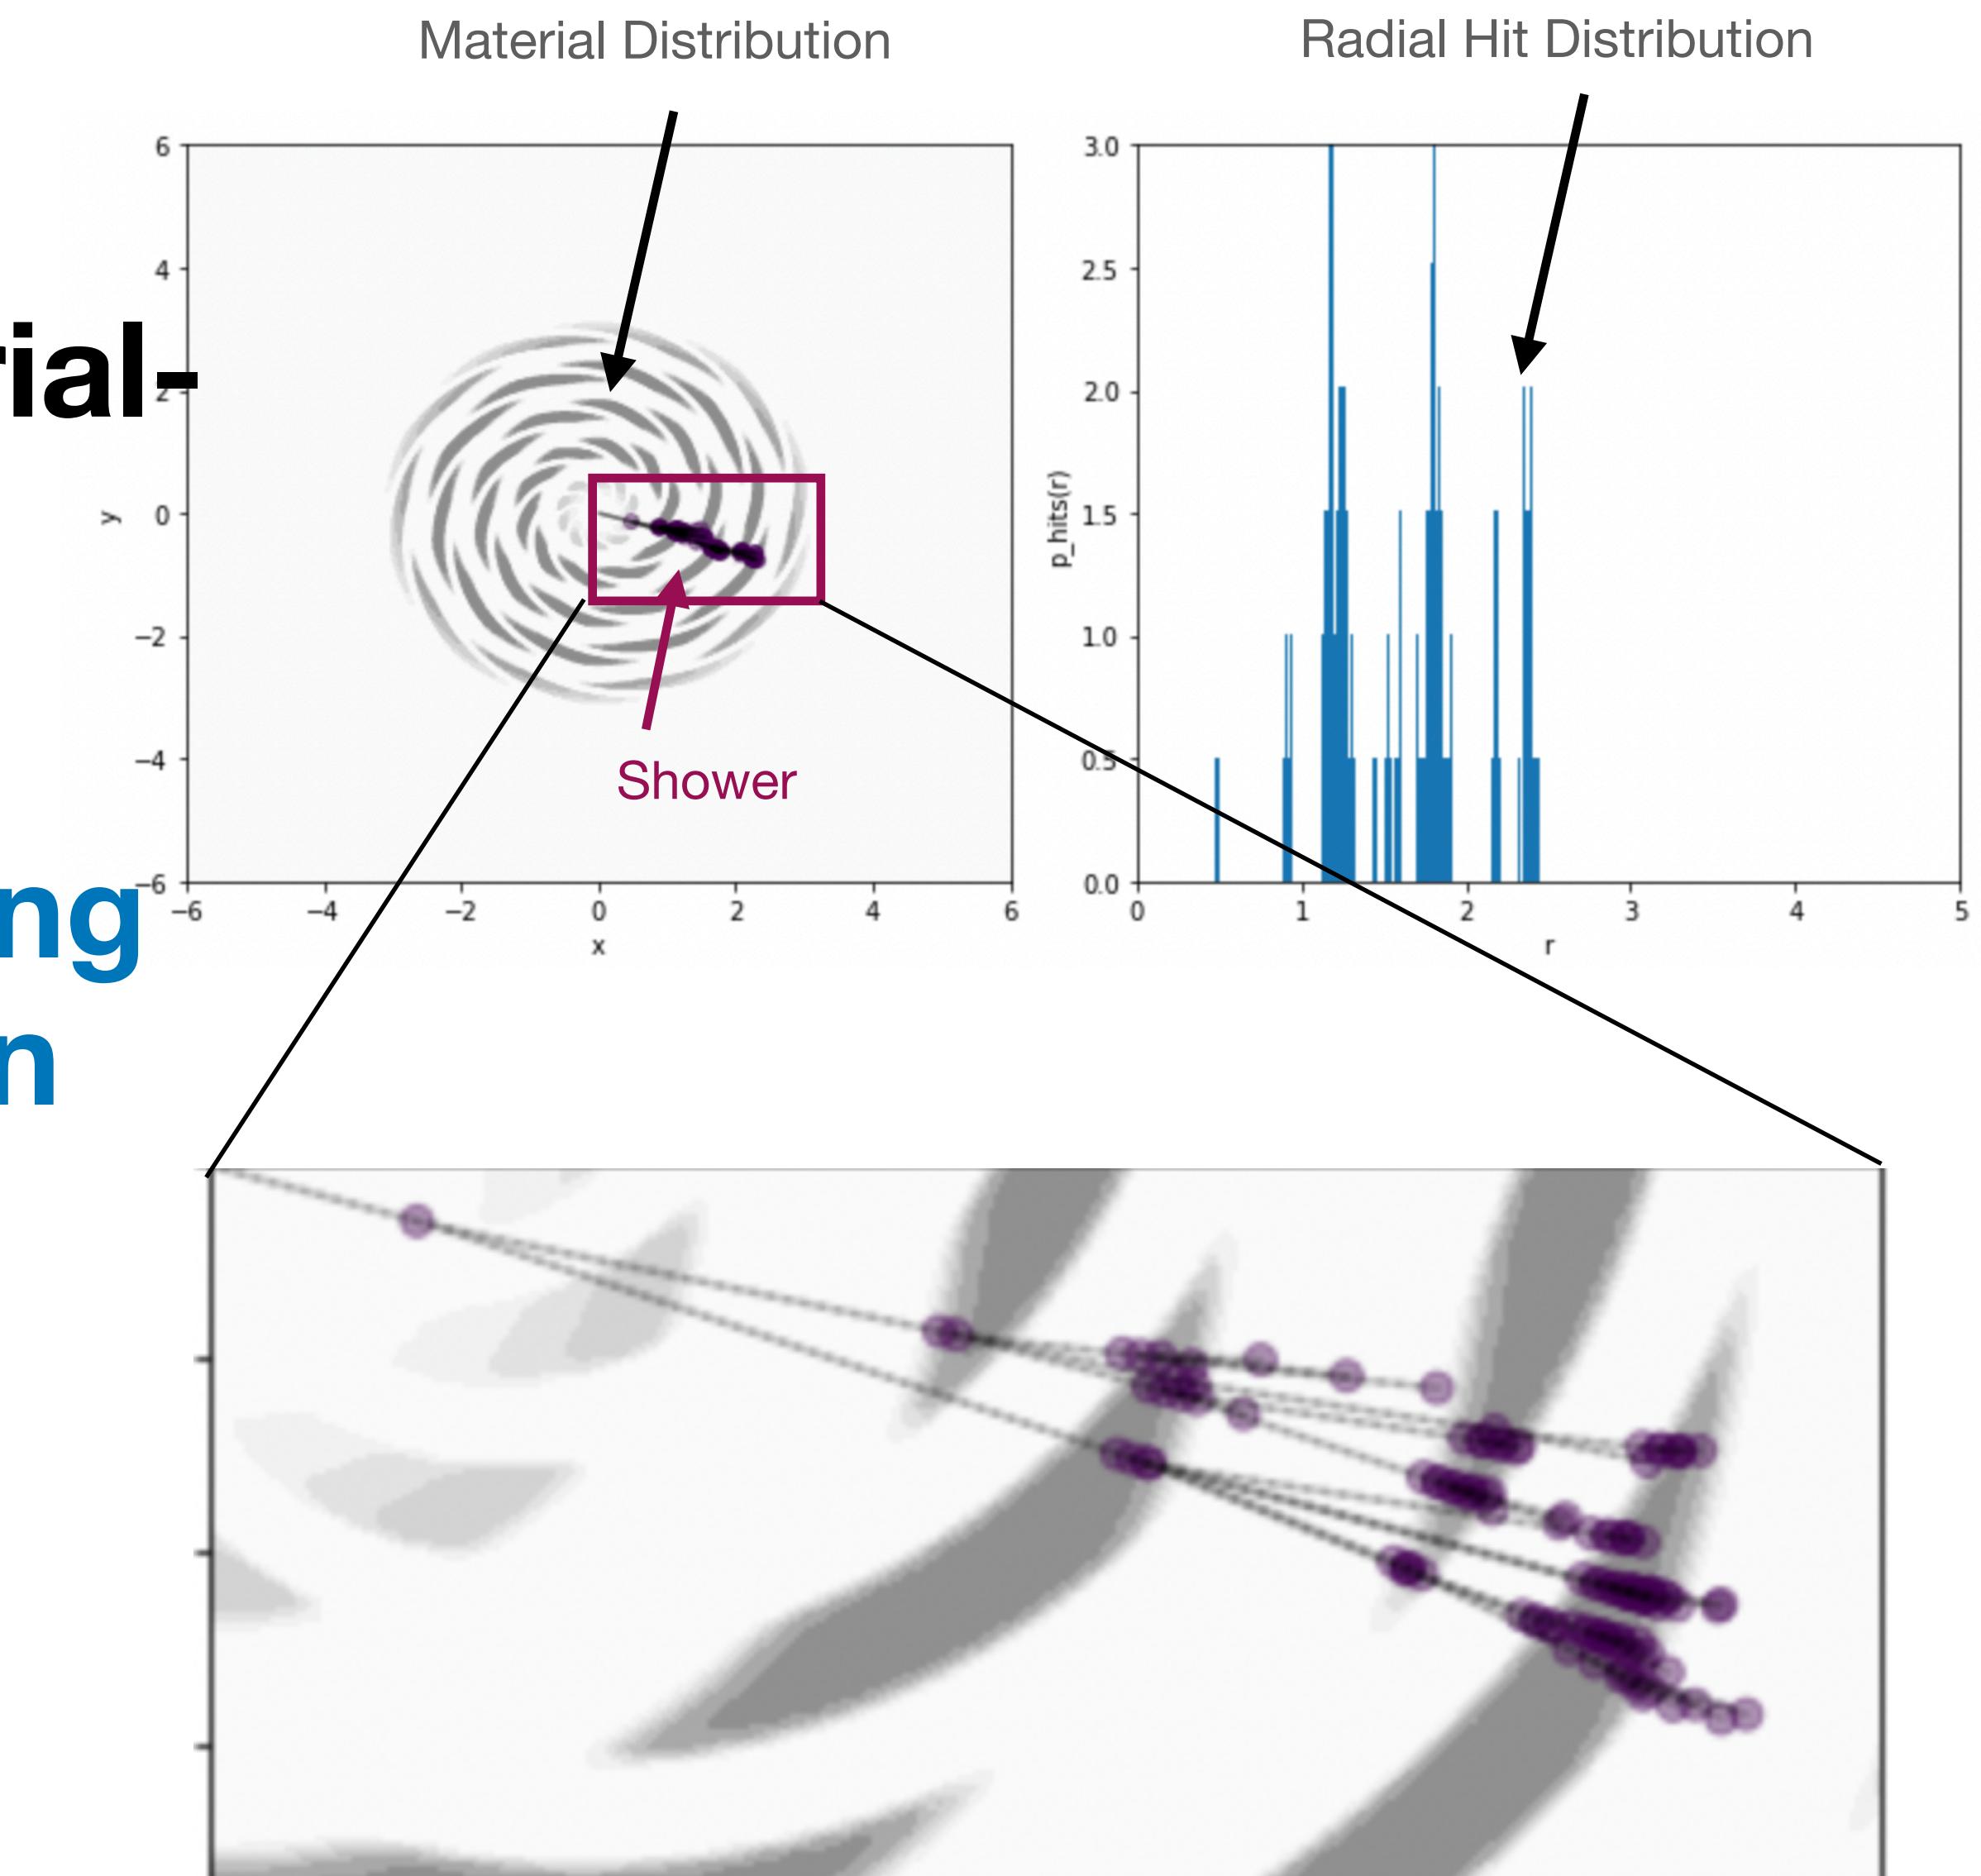
# Example Application

We consider a simplified material-induced particle shower

High Density: E-loss and splitting  
Low Density: linear propagation

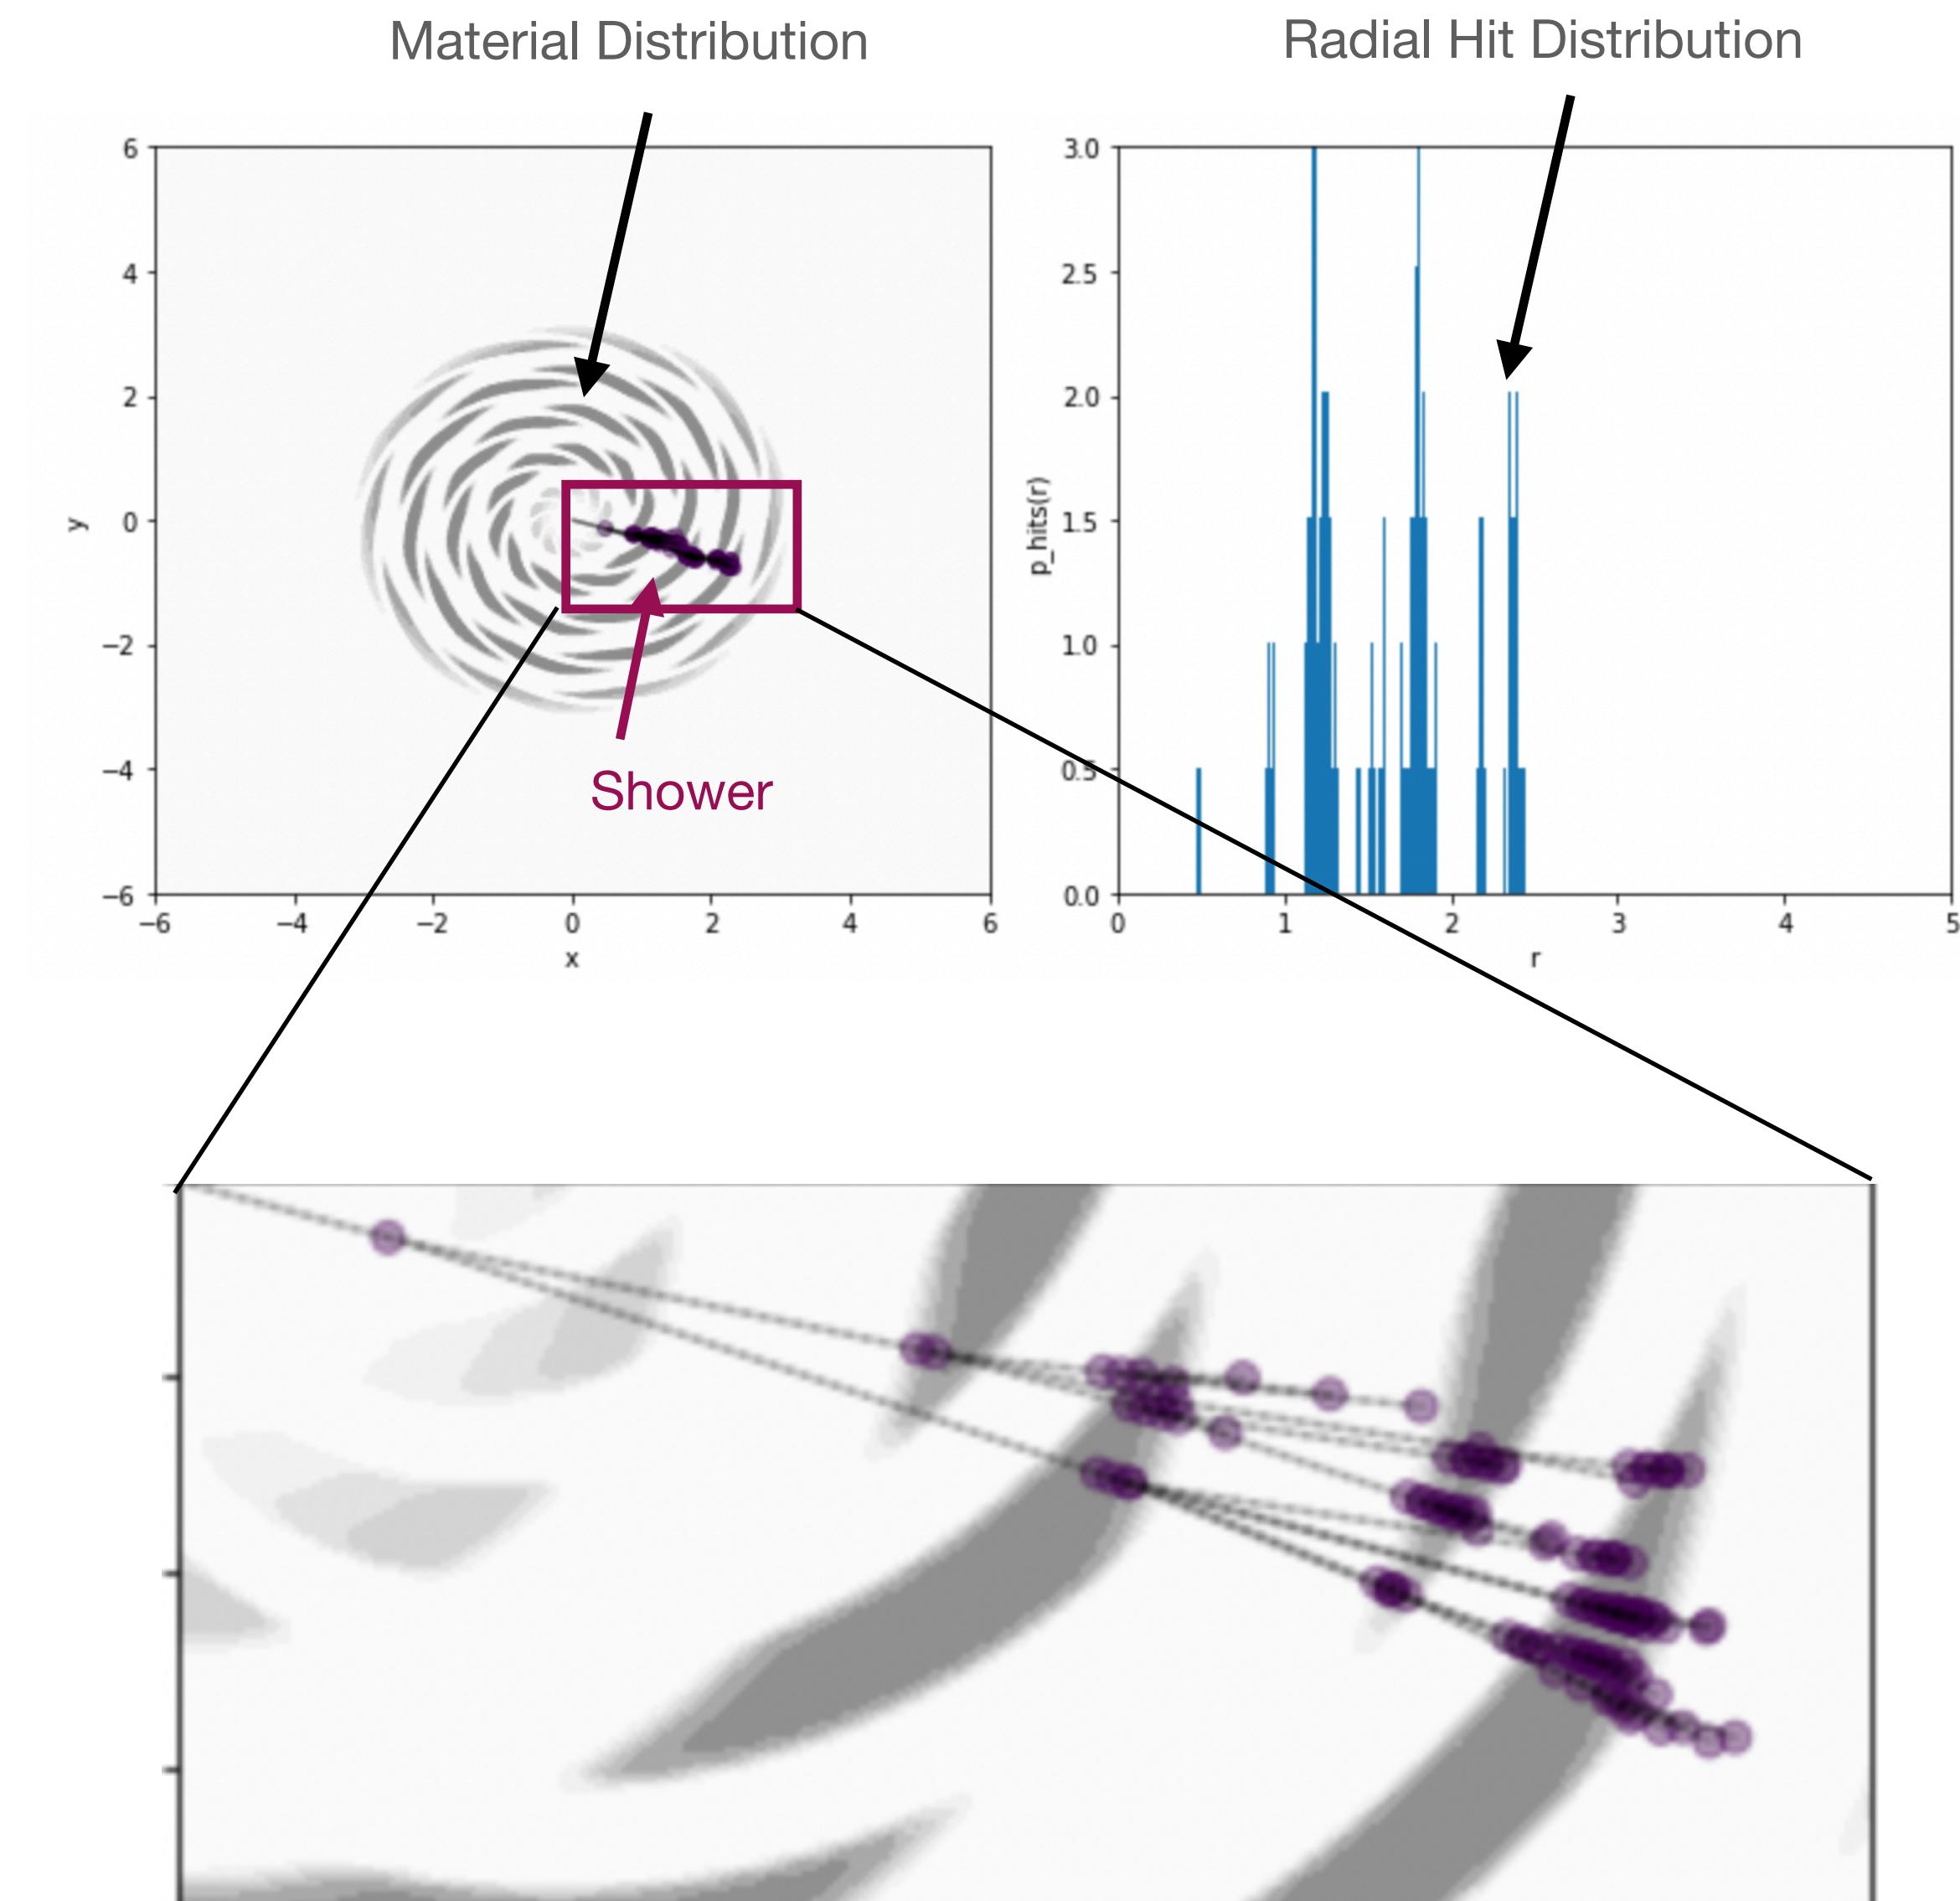
Design Parameter:  
Radial Distance of Material

Design Goal: Shower Depth



# Example Application

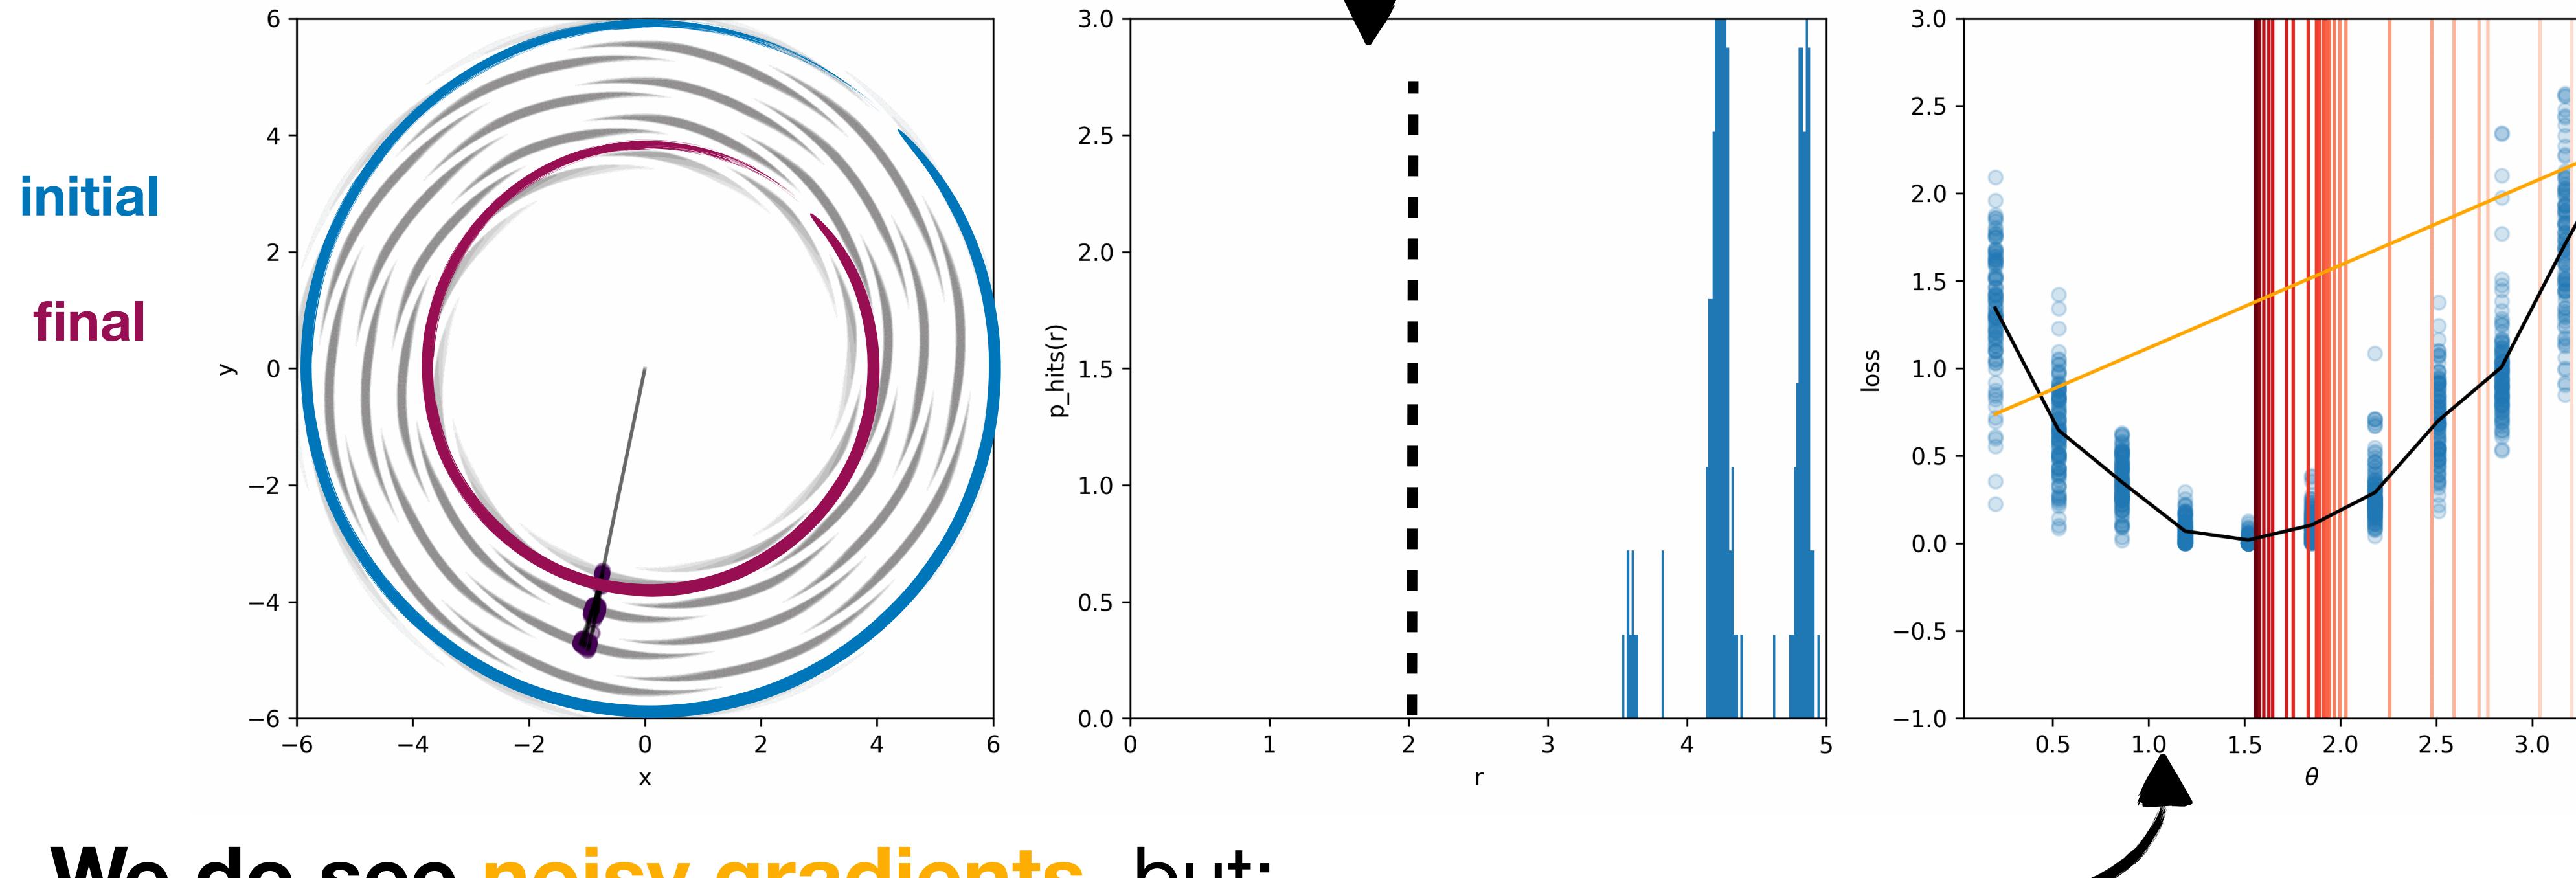
```
● ● ●  
  
def simulate(t,ϕ){  
    t' = propagate(t)  
    dointeract ~ p(interact|material(x,y,z| ϕ))  
    if dointeract:  
        hits.append([x,y,z])  
        dosplit ~ p(split|t')  
        if dosplit:  
            t1,t2 = split(t')  
            simulate(t1), simulate(t2)  
        else:  
            E = (1-α) E // energy loss  
            simulate(t')  
    }  
}
```



# Optimizing Material Distribution

At least for this toy model, optimizing material for hits properties works

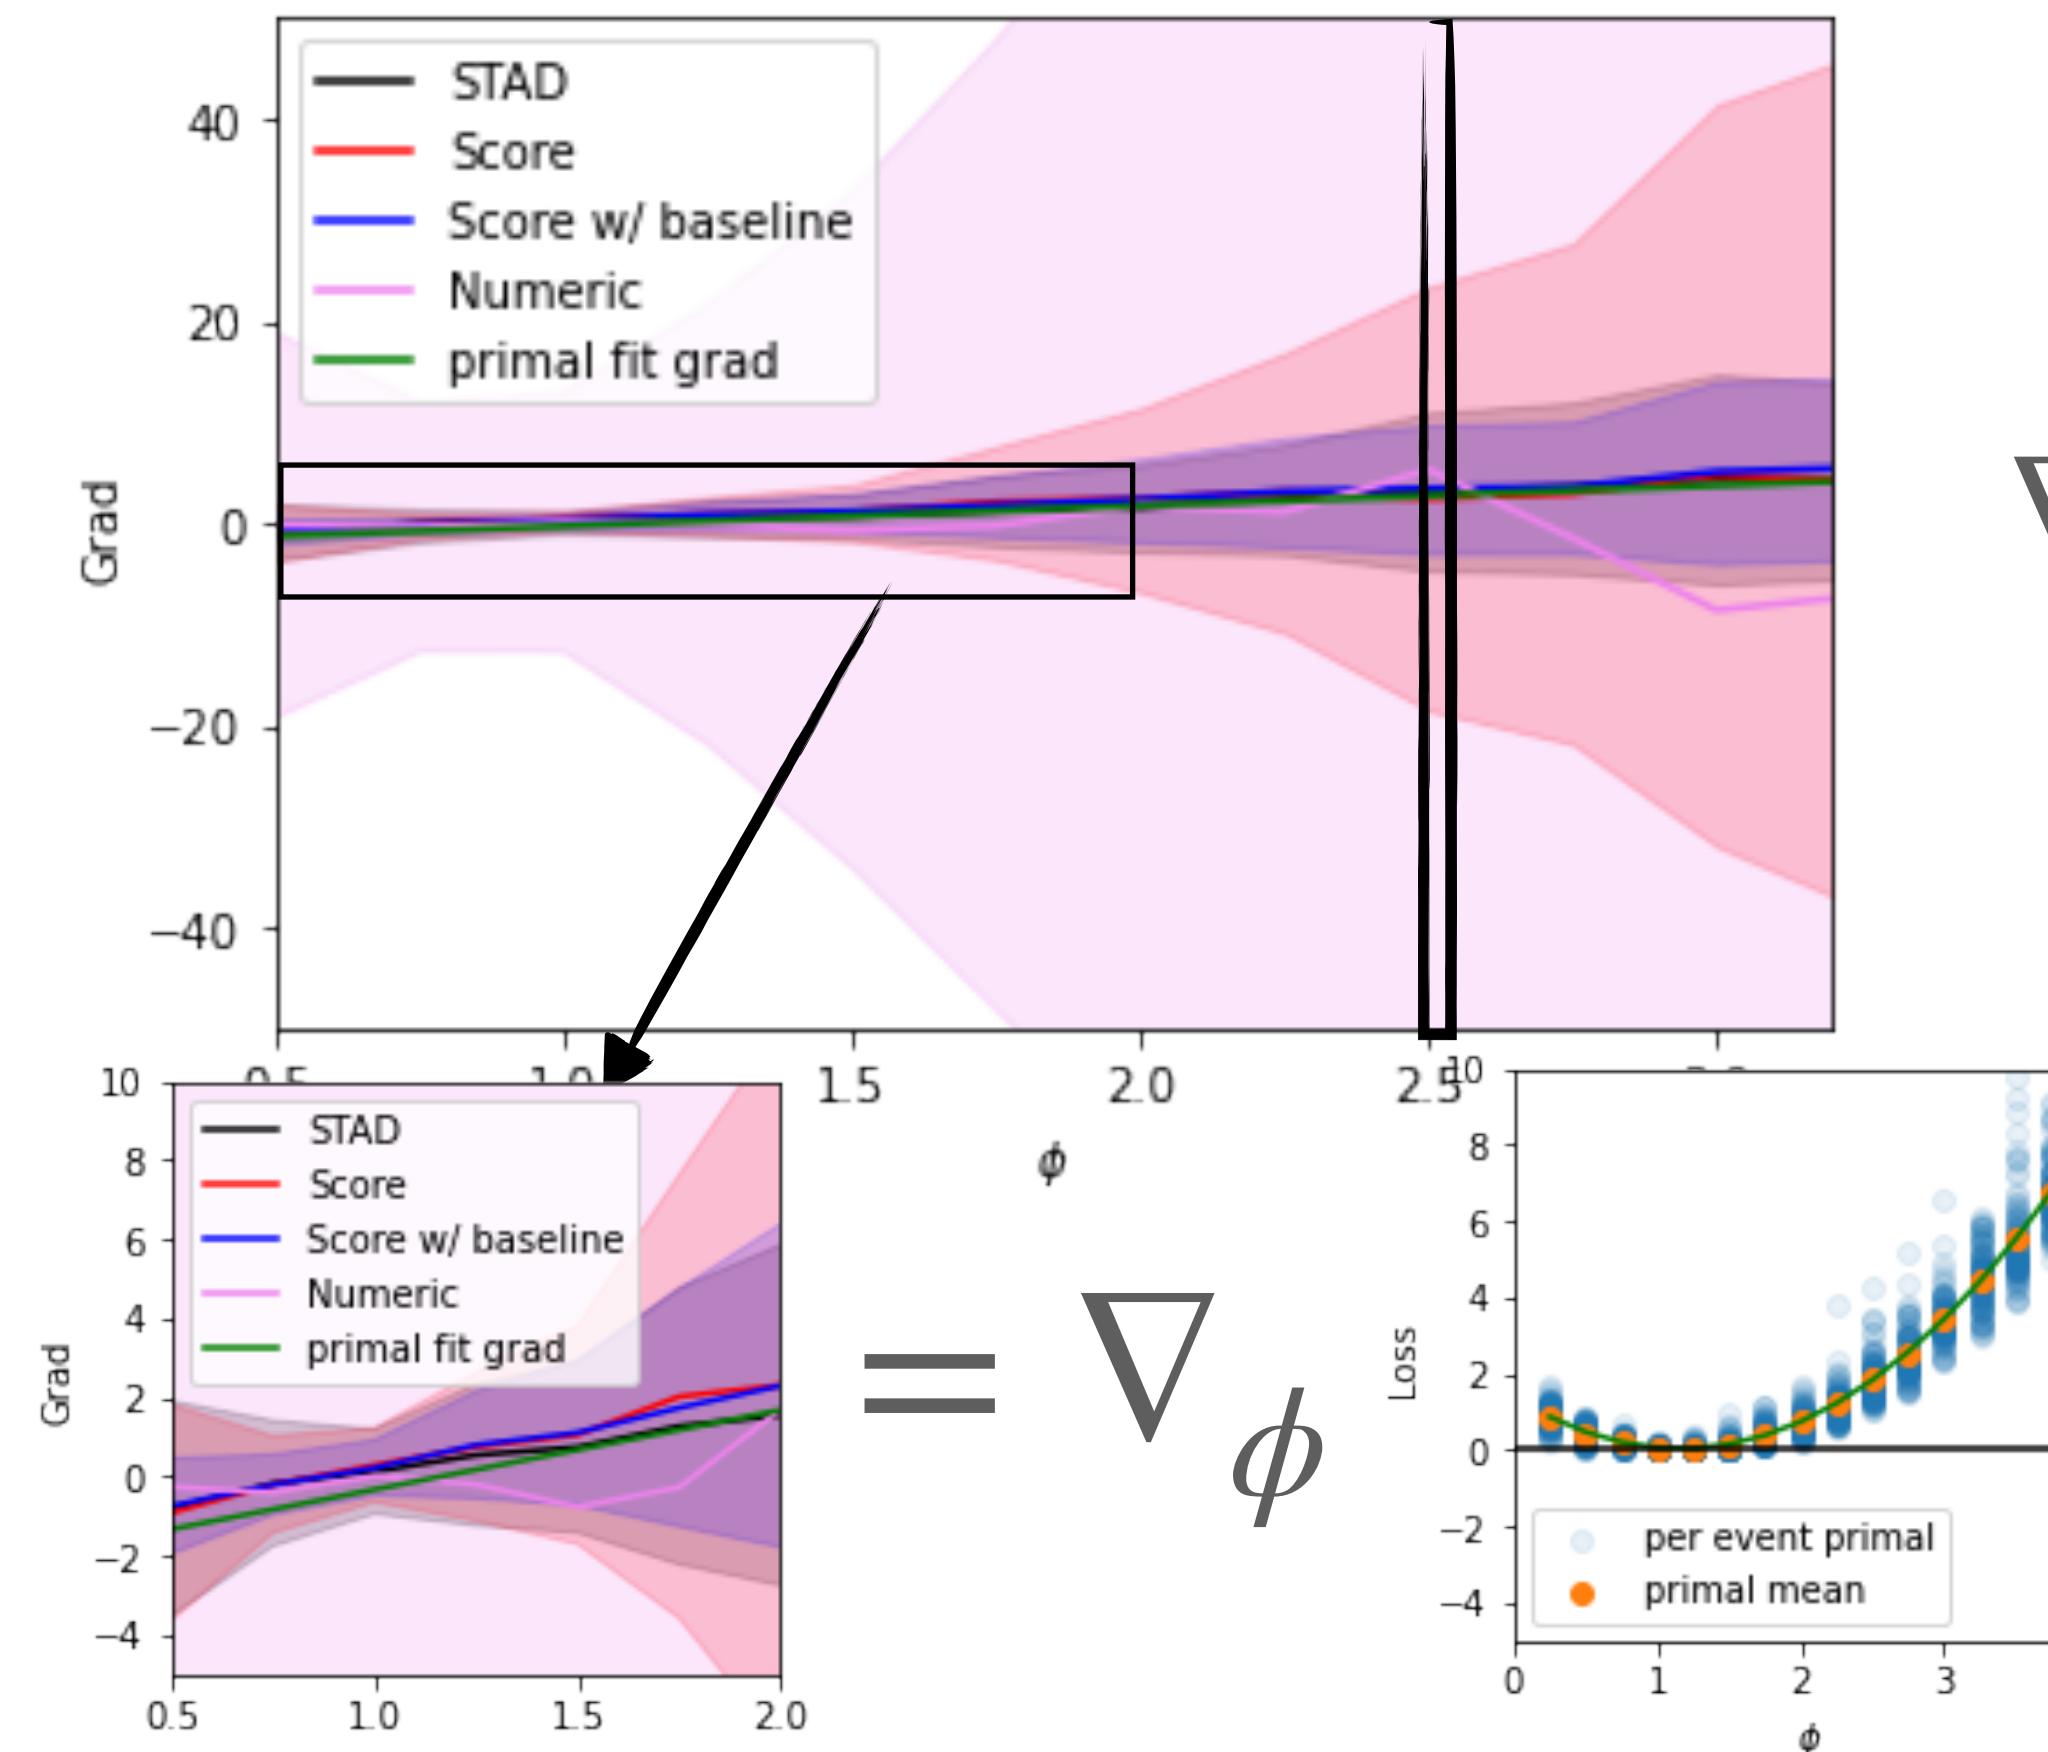
Objective: Mean Hit Radius  $\rightarrow 2.0$



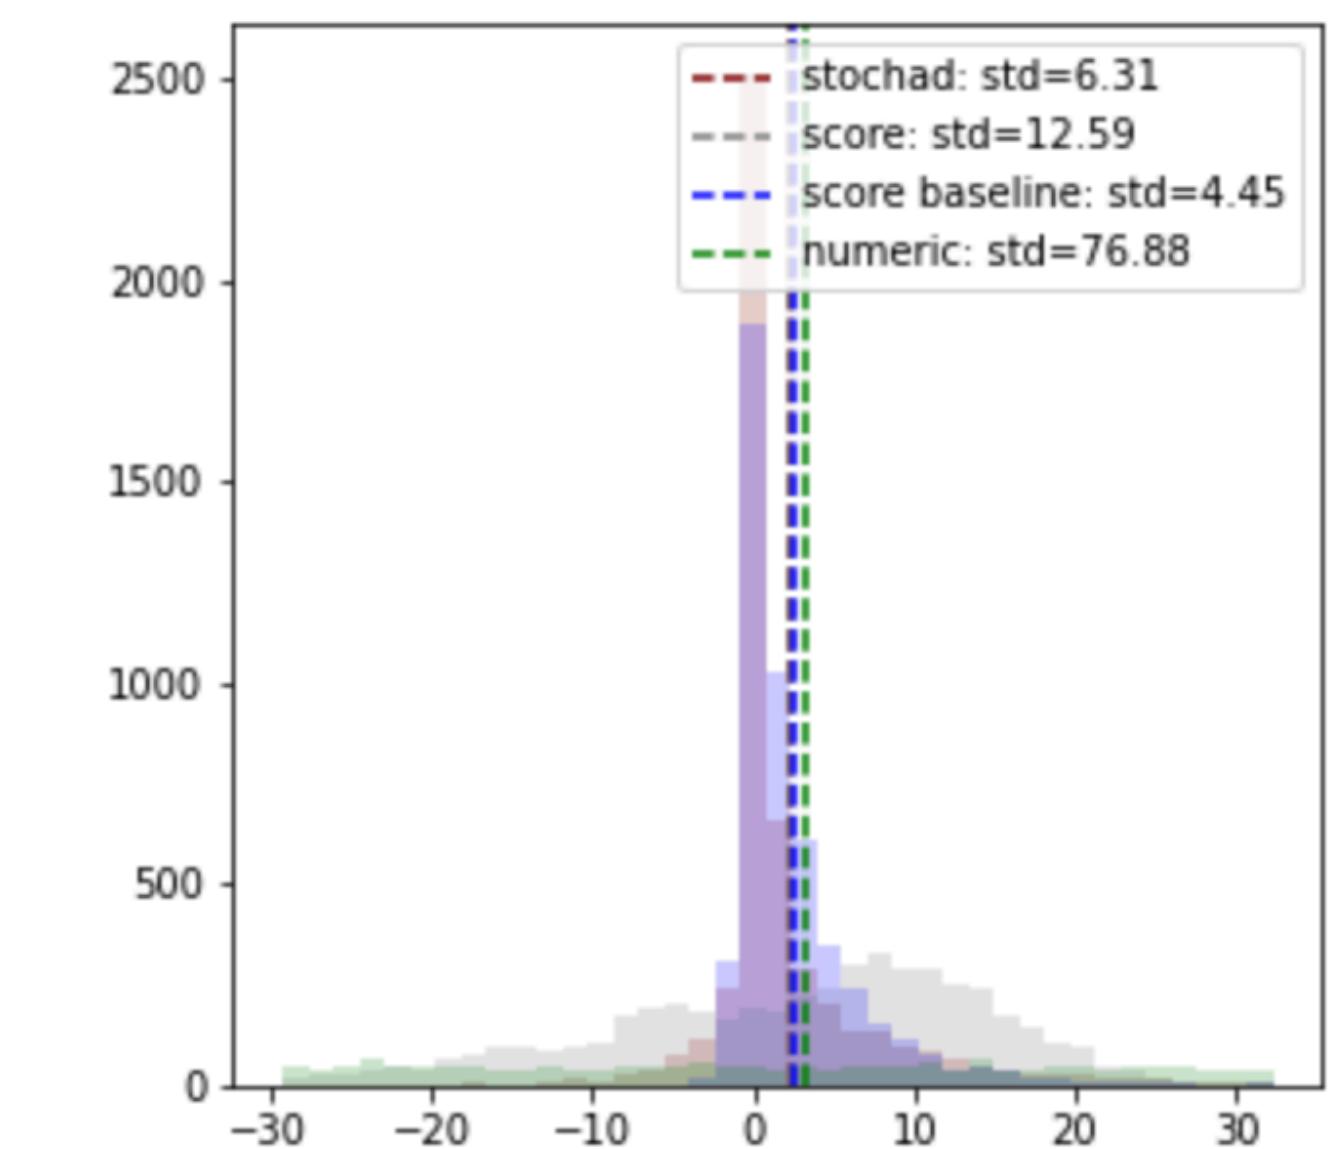
We do see **noisy gradients**, but:  
as long as direction is OK (in expectation), **minimization works**

# Comparison

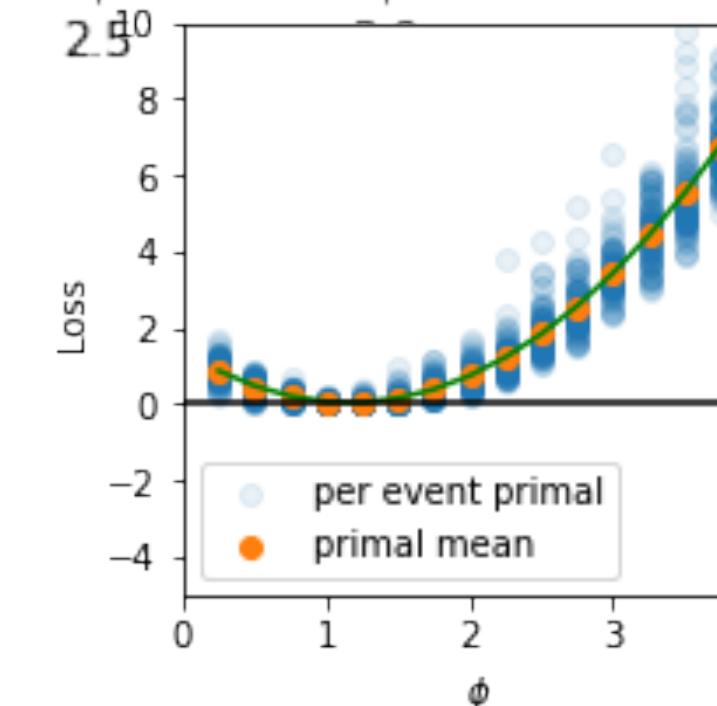
**Both score function and Stochastic AD gradient estimates massively reduce variance and optimization is possible**



$$\nabla_{\phi} L \Big|_{\phi=2.5}$$



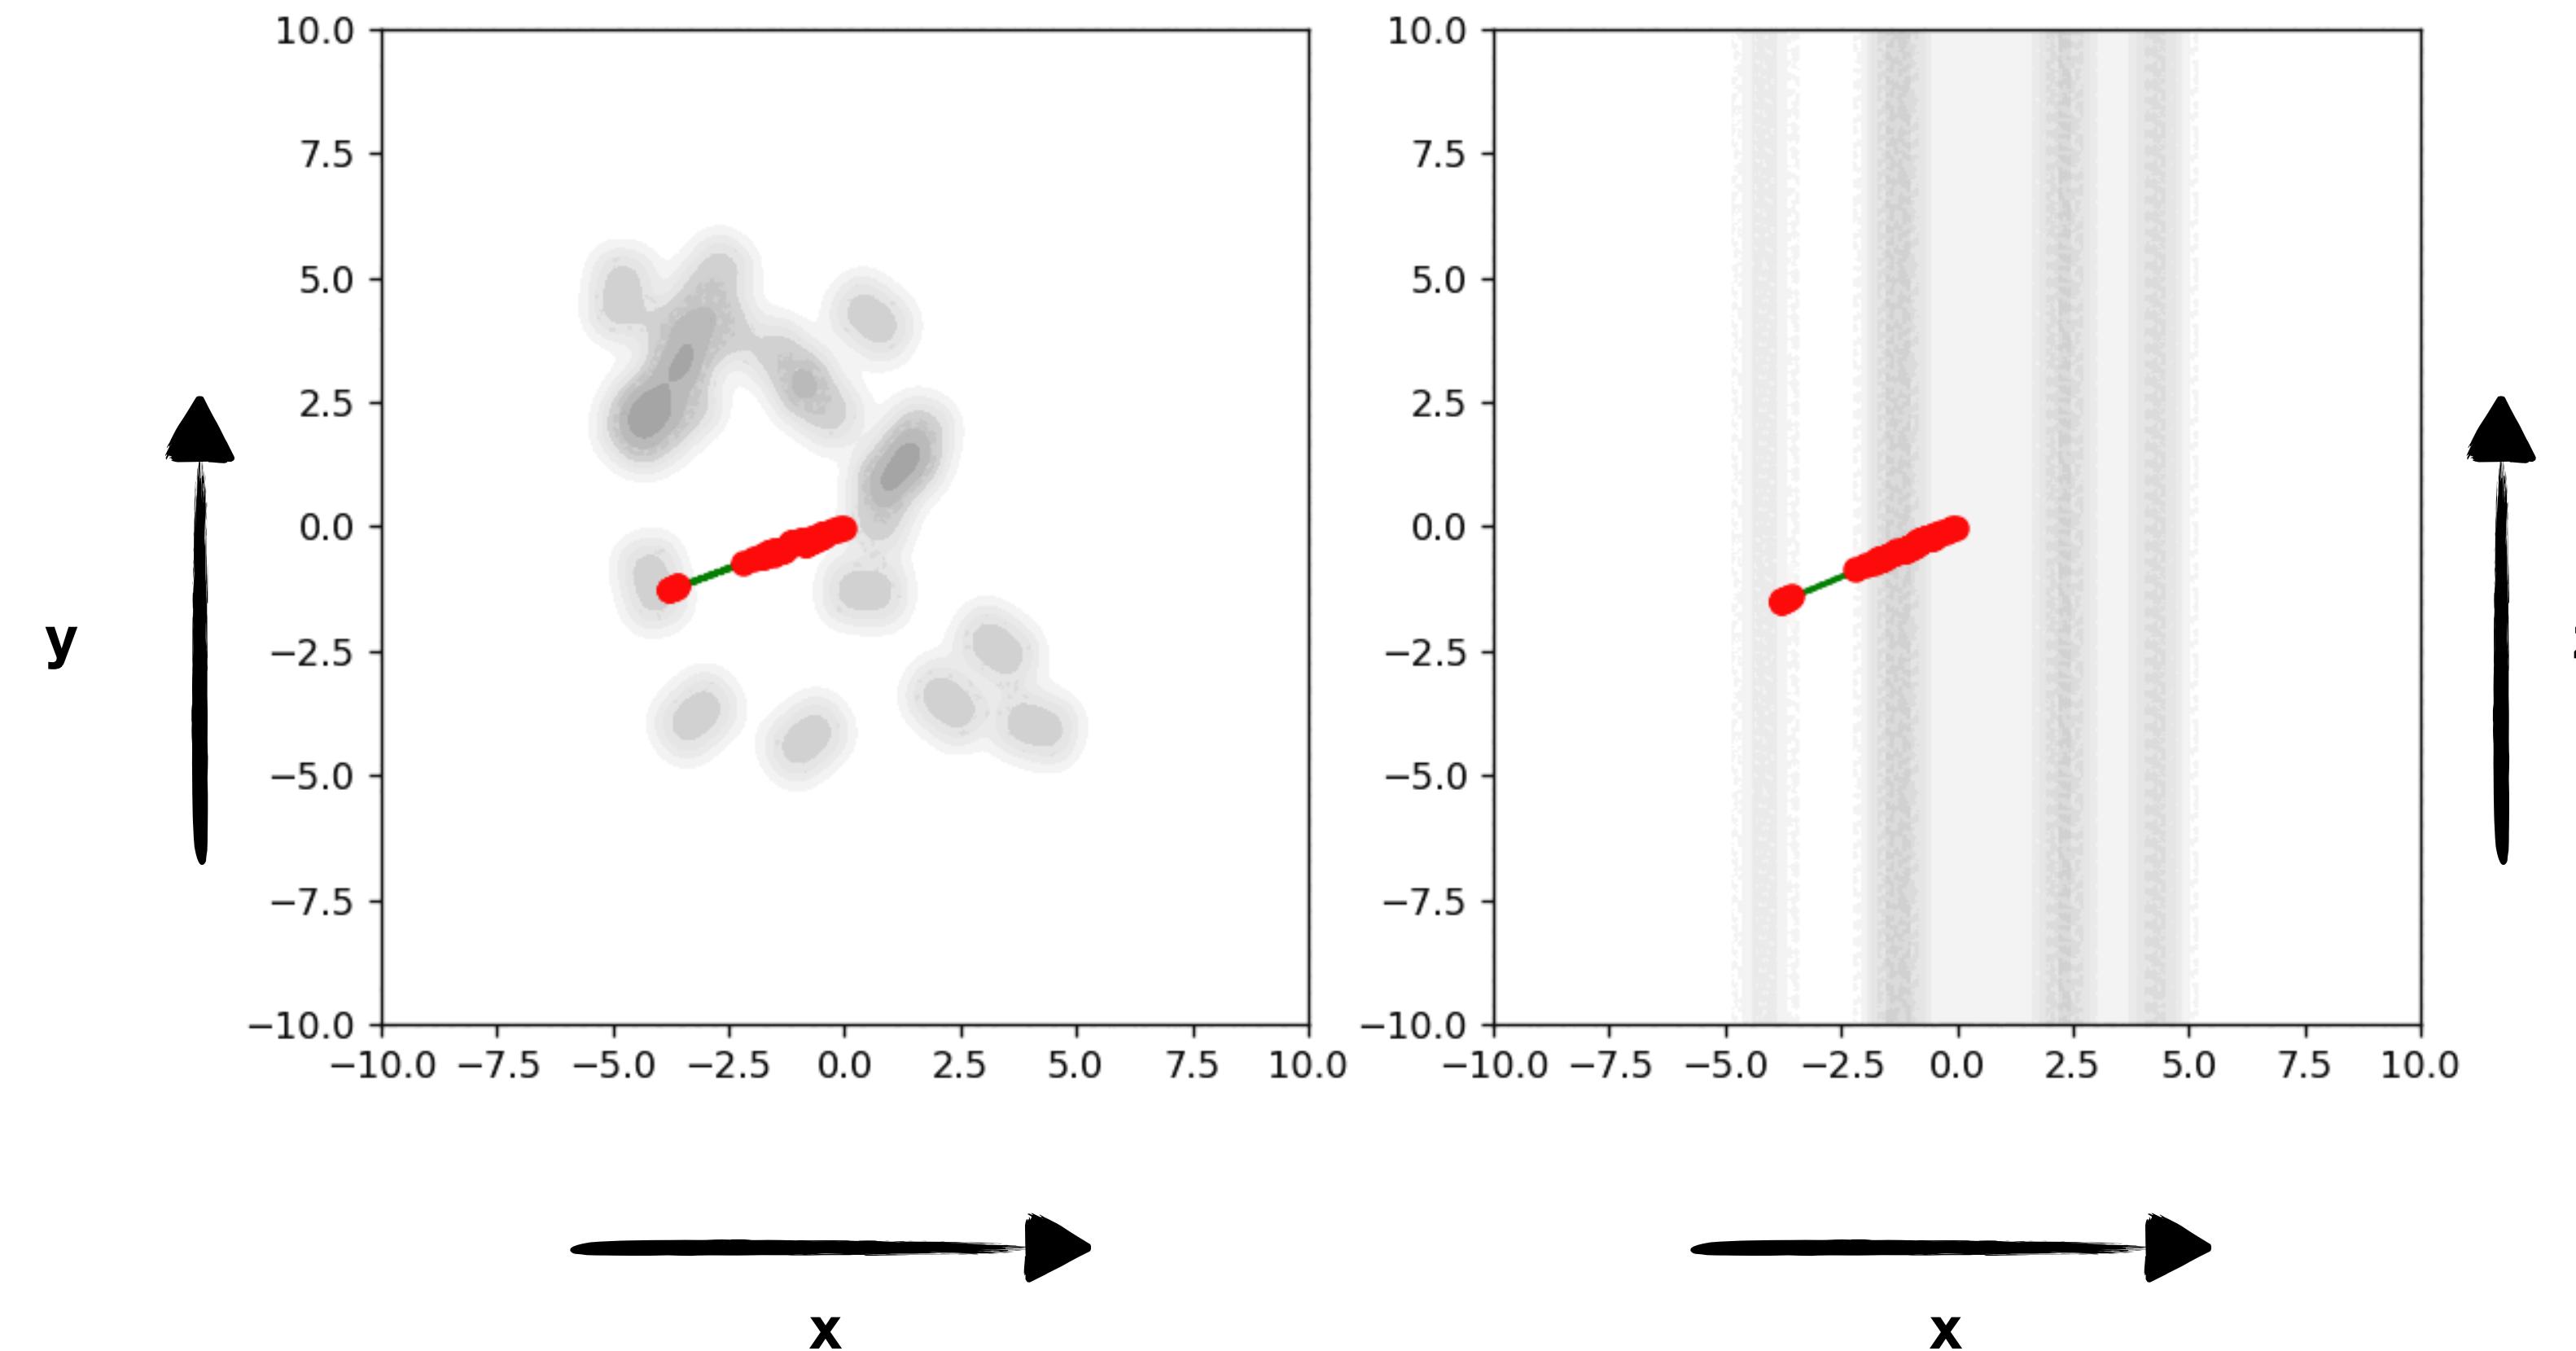
$$= \nabla_{\phi}$$



# A fun experiment:

Alternatively: 20 free-floating blocks of detectors.

**Self-assembly into a circular shell to contain showers**



# Next Steps & Outlook

**Takeaway message:**

The presence of *discrete stochastic randomness* makes *differentiable programming* difficult ...

**But not impossible.** With appropriate techniques, we can differentiate the seemingly undifferentiable

**But,** may require integration of more complex programming constructs than just AD.