

Differentiable programming and detector design optimization

Seminaire LPNHE - 14/02/2022

Julien Donini – Université Clermont Auvergne / LPC









Differentiable programming for HEP applications

- (a bit formal) introduction to **automatic differentiation**
- Optimization use-cases: analysis optimization, detector design
- MODE collaboration
- Ongoing projects

Acknowledgements

Many of the material presented in this talk was shown at the 1st Workshop on Differentiable Programming for experimental design organised by the MODE collaboration last September

. . .

Thanks to Atilim Gunes Baydin Tommaso Dorigo Giles Strong Nathan Simpson



https://indico.cern.ch/event/1022938/

Disclaimer



"You know nothing, Jon Snow"

Warm up: ML basics

5	3			7				
6			1	9	5			
	9	8					6	
8				6				3
4			8		3			1
7				2				6
	6					2	8	
			4	1	9			5
				8			7	9

Machine Learning Basics



Training Neural Networks



Training Neural Networks



Backpropagation in NN

Example: MLP network with 2 layers (1 hidden, 1 output)

Backward pass

Input data \mathbf{X} Hidden laye $\mathbf{s^{(1)}} = \mathbf{W^{(1)}}\mathbf{x} + \mathbf{b^{(1)}}$ $\downarrow \\ \mathbf{x}^{(\mathbf{1})} = f(\mathbf{s}^{(\mathbf{1})})$ Forward pass $\mathbf{s^{(2)}} = \mathbf{W^{(2)}}\mathbf{x^{(1)}} + \mathbf{b^{(2)}}$ \downarrow $\mathbf{x^{(2)}} = f(\mathbf{s^{(2)}})$ Output laye $\mathbf{y}(\mathbf{x}) = \mathbf{x}^{(2)}$ **NN** output

Use chain rule to compute derivatives of the loss $\ell(\mathbf{y}, \mathbf{t})$

$$\begin{aligned} \frac{\partial \ell}{\partial \mathbf{W}^{(2)}} &= \frac{\partial \ell}{\partial \mathbf{y}} \frac{\partial \mathbf{y}}{\partial \mathbf{s}^{(2)}} \frac{\partial \mathbf{s}^{(2)}}{\partial \mathbf{W}^{(2)}} \\ &= \frac{\partial \ell}{\partial \mathbf{y}} \frac{\partial f(\mathbf{s}^{(2)})}{\partial \mathbf{s}^{(2)}} \mathbf{x}^{(1)} \\ \frac{\partial \ell}{\partial \mathbf{W}^{(1)}} &= \frac{\partial \ell}{\partial \mathbf{y}} \frac{\partial \mathbf{y}}{\partial \mathbf{s}^{(2)}} \frac{\partial \mathbf{s}^{(2)}}{\partial \mathbf{x}^{(1)}} \frac{\partial \mathbf{x}^{(1)}}{\partial \mathbf{s}^{(1)}} \frac{\partial \mathbf{s}^{(1)}}{\partial \mathbf{W}^{(1)}} \\ &= \frac{\partial \ell}{\partial \mathbf{y}} \frac{\partial f(\mathbf{s}^{(2)})}{\partial \mathbf{s}^{(2)}} \frac{\partial \mathbf{s}^{(2)}}{\partial \mathbf{x}^{(1)}} \frac{\partial \mathbf{s}^{(2)}}{\partial \mathbf{s}^{(1)}} \frac{\partial f(\mathbf{s}^{(1)})}{\partial \mathbf{s}^{(1)}} \mathbf{x} \end{aligned}$$

Automatic differentiation

$$\frac{\partial}{\partial a} \ln f_{a,\sigma^{2}}(\xi_{1}) = \frac{(\xi_{1} - a)}{\sigma^{2}} f_{a,\sigma^{2}}(\xi_{1}) = \frac{1}{\sqrt{2\pi\sigma}} \int_{\mathbb{R}^{d}} T(x)f(x,\theta)dx = \int_{\mathbb{R}^{d}} \frac{\partial}{\partial \theta} T(x)f(x,\theta)dx = \int_{\mathbb{R}^{d}} T(x) \cdot \frac{\partial}{\partial \theta} \int_{\mathbb{R}^{d}} f(x,\theta)dx = M\left(T(\xi) \cdot \frac{\partial}{\partial \theta} \ln L(\xi,\theta)\right) \int_{\mathbb{R}^{d}} \frac{\partial}{\partial \theta} \int_{\mathbb{R}^{d}} \frac{\partial}{\partial \theta} \int_{\mathbb{R}^{d}} f(x,\theta)dx = \int_{\mathbb{R}^{d}} \frac{\partial}{\partial \theta} \int_{\mathbb{R}^{d}} f(x,\theta)$$

How to code derivatives ?



(Note: numerical differentiation is also an option but it gives approximate results)

Baydin et al. [1502.05767]

How to code derivatives ?



Baydin et al. [1502.05767]

How to code derivatives ?



Julien Donini – Differentiable programming and detector design optimization

Automatic (algorithmic) differentiation (AD)

- Numerical derivative evaluations rather than derivative expressions
- Composition of operations for which derivatives are known
- No need to rearrange the code in a closed-form expression
- Accurate at machine precision

For each function a computational graph is constructed
→ evaluation of the function (forward pass)
→ calculation of gradient (backward pass)

Computational graph (example)



(see: https://pytorch.org/blog/overview-of-pytorch-autograd-engine/)

Julien Donini – Differentiable programming and detector design optimization

backward

<u> </u>*θ*ν2

SinBackward

Automatic differentiation

Two **main modes**, both based on chain rule



 $\bar{v}_i = \frac{\partial f}{\partial v_i}$

Forward mode

Example

$$y = f(x_1, x_2) = \ln(x_1) + x_1 x_2 - \sin(x_2)$$

Each **intermediate** $\dot{v}_i = \dot{v}_i$

$$v_i = \frac{\partial v_i}{\partial x}$$



F	orward Primal Tra	ice	F	orwa	rd Tangent (Derivative	e) Trace		
1	$v_{-1} = x_1$	=2	T	\dot{v}_{-1}	$x_1 = \dot{x}_1$	= 1	$\dot{v}_{\dot{\cdot}} =$	$\frac{\partial v_i}{\partial v_i}$
	$v_0 = x_2$	= 5		\dot{v}_0	$=\dot{x}_2$	= 0		∂x_1
	$v_1 = \ln v_{-1}$	$=\ln 2$		\dot{v}_1	$=\dot{v}_{-1}/v_{-1}$	= 1/2		
	$v_2 = v_{-1} \times v_0$	$= 2 \times 5$		\dot{v}_2	$= \dot{v}_{-1} \times v_0 + \dot{v}_0 \times v_{-1}$	$= 1 \times 5 + 0$	$\times 2$	
	$v_3 = \sin v_0$	$=\sin 5$		\dot{v}_3	$=\dot{v}_0 \times \cos v_0$	$= 0 \times \cos 5$		
	$v_4 = v_1 + v_2$	= 0.693 + 10		\dot{v}_4	$=\dot{v}_1+\dot{v}_2$	= 0.5 + 5		
	$v_5 = v_4 - v_3$	= 10.693 + 0.959		\dot{v}_5	$=\dot{v}_4-\dot{v}_3$	= 5.5 - 0		
♦	$y = v_5$	= 11.652	▼	\dot{y}	$=\dot{v}_{5}$	= 5.5		

Forward mode example, evaluated at $(x_1, x_2) = (2, 5)$ and setting $\dot{x}_1 = 1$ to compute $\dot{y} = \frac{\partial y}{\partial x_1}$

[1502.05767]

Digression: Dual numbers

Forward mode can be viewed as evaluating a function using dual numbers

Numbers defined as $v + \dot{v}\epsilon$ where $\epsilon \neq 0$ and $\epsilon^2 = 0$

Properties (using Taylor expansion):

 $f(v + \dot{v}\epsilon) = f(v) + f'(v)\dot{v}\epsilon$

$$f(g(v + \dot{v}\epsilon)) = f(g(v) + g'(v)\dot{v}\epsilon))$$

= $f(g(v)) + f'(g(v))g'(v)\dot{v}\epsilon$

Composite function derivative !

In practice, implement **specific code** to handle the **dual operations** so that **function f and its derivative are simultaneously computed** (operator overloading)

(For a concrete example with code see R. Lange post)

Reverse mode (backpropagation)

Example $y = f(x_1, x_2) = \ln(x_1) + x_1 x_2 - \sin(x_2)$	$x_1 \xrightarrow{x_1} \underset{v_{-1}}{\overset{v_1}{\underset{v_1}{\underset{v_1}{\underbrace{v_1}{v_1}{\underbrace{v_1}{\underbrace{v_1}{\underbrace{v_1}{\underbrace{v_1}{\underbrace{v_1}{\underbrace{v_1}{\underbrace{v_1}{\underbrace{v_1}{\underbrace{v_1}{\underbrace{v_1}{\underbrace{v_1}{\underbrace{v_1}{v_1}{\underbrace{v_1}{v_1}{\underbrace{v_1}{\underbrace{v_1}{v_1}{\underbrace{v_1}{v_1}{v_1}{\underbrace{v_1}{v_1}{\underbrace{v_1}{v_1}{\underbrace{v_1}{v_1}{v_1}{v_1}{v_1}{v_1}{v_1}{v_1}$		- y
Propagates derivatives $\bar{v}_i = \frac{\partial y}{\partial v_i}$	$x_2 \xrightarrow{v_0} x_2$	-v ₃ sin	$v_5 \longrightarrow f(x_1, x_2)$
Forward Primal Trace	Reverse Adjoint (Derivative) Trace		
$v_{-1} = x_1 \qquad = 2$	$\blacklozenge \bar{x}_1 = \bar{v}_{-1}$	= 5.5	
$v_0 = x_2 = 5$	$ar{x}_2 = ar{v}_0$	= 1.716	
$v_1 = \ln v_{-1} = \ln 2$	$\bar{v}_{-1} = \bar{v}_{-1} + \bar{v}_1 \frac{\partial v_1}{\partial v_{-1}} = \bar{v}_{-1} + \bar{v}_1 / v$	$_{1} = 5.5$	
$v_2 = v_{-1} \times v_0 = 2 \times 5$	$\bar{v}_0 = \bar{v}_0 + \bar{v}_2 \frac{\partial v_2}{\partial v_0} = \bar{v}_0 + \bar{v}_2 \times v$	$_{1} = 1.716$	
	$\bar{v}_{-1} = \bar{v}_2 \frac{\partial v_2}{\partial v_{-1}} \qquad = \bar{v}_2 \times v_0$	= 5	
$v_3 = \sin v_0 \qquad = \sin 5$	$\bar{v}_0 = \bar{v}_3 \frac{\partial v_3}{\partial v_0} = \bar{v}_3 \times \cos v_0$	= -0.284	
$v_4 = v_1 + v_2 = 0.693 + 10$	$\bar{v}_2 = \bar{v}_4 \frac{\partial v_4}{\partial v_2} \qquad = \bar{v}_4 \times 1$	= 1	
	$\bar{v}_1 = \bar{v}_4 \frac{\partial v_4}{\partial v_1} \qquad = \bar{v}_4 \times 1$	= 1	
$v_5 = v_4 - v_3 = 10.693 + 0.959$	$\bar{v}_3 = \bar{v}_5 \frac{\partial v_5}{\partial v_3} = \bar{v}_5 \times (-1)$	= -1	
	$\bar{v}_4 = \bar{v}_5 \frac{\partial v_5}{\partial v_4} \qquad = \bar{v}_5 \times 1$	= 1	
$\checkmark y = v_5 = 11.652$	$\overline{v}_5 = \overline{y} = 1$		

Reverse mode example, evaluated at $(x_1, x_2) = (2, 5)$. Both $\frac{\partial y}{\partial x_1}$ and $\frac{\partial y}{\partial x_2}$ are computed on the same reverse pass starting from the output

$$\bar{v}_5 = \bar{y} = \frac{\partial y}{\partial y} = 1$$

[1502.05767]

Current state of differentiable programming

Evolution of frameworks

From: coarse-grained (module level) backprop Towards: fine-grained, general-purpose automatic differentiation



Auto-diff tools: http://www.autodiff.org/

[slide G. Baydin]

theano

TensorFlow

torch

Julien Donini – Differentiable programming and detector design optimization

Pytorch example (simple)

PyTorch enables automatic differentiation of computational graphs

Example: differentiation of $f(x) = x^4$

f = 1

Pure python



for i in range(4):
 f = f * x

def fun(x):
 return x.pow(4)

x = torch.tensor(2.,requires_grad=True)
f = fun(x)



```
\rightarrow f(2) = 16, f'(2) = 32, f''(2) = 48
```

Note: exact numerical value of derivative is computed (no analytical form)

Pytorch example (differential geometry)

Calculate partial derivatives of metric tensor w.r.t to coordinates $x^{\alpha} = (ct, r, \theta, \phi)$

$$g_{\mu\nu} = \begin{pmatrix} 1 - \frac{R_s}{R} & 0 & 0 & 0\\ 0 & -\left(1 - \frac{R_s}{R}\right)^{-1} & 0 & 0\\ 0 & 0 & -R^2 & 0\\ 0 & 0 & 0 & -R^2 \sin^2\theta \end{pmatrix}$$

Jacobian matrix:

grad = functional.jacobian(metric_tensor,coord)
print(grad.T)



2nd order derivatives:

```
def fgrad(inputs):
    return torch.autograd.functional.jacobian(metric_tensor,inputs, create_graph=True)
hessian = torch.autograd.functional.jacobian(fgrad,coord, create_graph=True)
```

print(hessian)

$$\longrightarrow \frac{\partial^2 g_{\mu\nu}}{\partial x^{\alpha} \partial x^{\beta}}$$
 (16 matrices

Julien Donini – Differentiable programming and detector design optimization

Differentiable programming

Gradient-based **optimization** methods:

→Code composed of differentiable and parameterized building blocks

→ Software optimized via automatic differentiation



Gradient descent can write code better than you. I'm sorry.

3:56 PM - 4 Aug 2017



OK, Deep Learning has outlived its usefulness as a buzz-phrase. Deep Learning est mort. Vive Differentiable Programming!

Differentiable programming in HEP

Incorporating automatic differentiation in HEP software

In Analysis Code

- Optimize free parameters with respect to the desired physics objective
- End-to-end differentiable analysis workflow

In Simulation Code

- Compute gradient for with respect to parameters of simulation
- EFT, Cosmology, MadGraph (evaluation of matrix element), ...

Differentiable programming in High Energy Physics, SnowMass 2021

Differentiable analysis: NEOS

End-to-end optimized analysis pipelines that use the analysis sensitivity including systematic uncertainties as the objective function

github.com/gradhep/neos



[figure N. Simpson]

Differentiable analysis: NEOS

End-to-end optimized analysis pipelines that use the analysis sensitivity including systematic uncertainties as the objective function

github.com/gradhep/neos



[Slide Nathan Simpson]

INFERNO

Inference Aware Neural Optimization

[1806.04743, de Castro, Dorigo]

 Include nuisance parameters in the loss function and directly minimize precision of parameters of interest (e.g. signal strenght measurement)



Profiled likelihood around the expectation value for the parameter of interest for **inference-aware** models and **cross-entropy** loss based models.

INFERNO algorithm



INFERNO algorithm



NN loss function: uncertainty on parameter of interest (U)

Obtained by computing full **hessian** of the **Likelihood** with respect to all **nuisance** parameters

INFERNO algorithm



Algorithm 1 Inference-Aware Neural Optimisation.

Input 1: differentiable simulator or variational approximation $g(\theta)$.

Input 2: initial parameter values θ_s .

Input 3: parameter of interest $\omega_0 = \theta_k$.

Output: learned summary statistic $s(D; \phi)$.

- 1: **for** i = 1 to *N* **do**
- 2: Sample a representative mini-batch G_s from $g(\boldsymbol{\theta}_s)$.
- 3: Compute differentiable summary statistic $\hat{s}(G_s; \phi)$.
- 4: Construct Asimov likelihood $\mathcal{L}_A(\boldsymbol{\theta}, \boldsymbol{\phi})$.
- 5: Get information matrix inverse $I(\boldsymbol{\theta})^{-1} = \boldsymbol{H}_{\boldsymbol{\theta}}^{-1}(\log \mathcal{L}_A(\boldsymbol{\theta}, \boldsymbol{\phi})).$

6: Obtain loss
$$U = I_{kk}^{-1}(\boldsymbol{\theta}_s)$$
.

- 7: Update network parameters $\phi \to \text{SGD}(\nabla_{\phi} U)$.
- 8: end for

For a detailed explanation of the algorithm see G. Strong blog post

Julien Donini – Differentiable programming and detector design optimization

Optimization of detector design

Can automatic differentiation be applied to detector optimization ?



Optimization of detector design

Design of detectors traditionally relies on individual optimization of subdetector

Track first, destroy later

• First detect ionization tracks in tracker, then measure energy deposits from destructive interaction with thick calorimeters

Per-subdetector optimization

- subdetector-specific figures of merit (e.g. momentum resolution)
- Impact on physics goals typically considered in a second step

Optimization of a **joint problem** ≠ different from **individual optimization**

 $argmax_{x,y}(\mathcal{L}(x,y)) \neq \left[argmax_x(\int \mathcal{L}(x,y)dy), argmax_y(\int \mathcal{L}(x,y)dx)\right]$

Proof of concept: MUonE experiment

YWW

hadrons

Example of geometry optimization: **MUonE** experiment

- MUonE: high precision muon-electron differential cross section
 - \rightarrow hadronic contributions to g-2 muon anomaly





Generic optimization pipeline

LAIES1: 10.17 UPDAIE) CHANGES IN VERSION 10.17: THE CPU NO LONGER OVERHEATS WHEN YOU HOLD DOWN SPACEBAR. COMMENTS: LONGTIME USERY WRITES: THIS UPDATE BROKE MY WORKFLOW! MY CONTROL KEY IS HARD TO REACH, 50 I HOLD SPACEBAR INSTEAD, AND I CONFIGURED EMACS TO INTERPRET A RAPID TEMPERATURE RISE AS "CONTROL". ADMIN WRITES: THAT'S HORRIFYING. LONGTIMEUSER4 WRITES: LOOK, MY SETUP WORKS FOR ME. JUST ADD AN OPTION TO REENABLE SPACEBAR HEATING.

EVERY CHANGE BREAKS SOMEONE'S WORKFLOW.

Generic optimization pipeline



Minimization of objective function through automatic differentiation

Generic optimization pipeline

What if simulator is not differentiable ? Try generative surrogate



Black-Box Optimization with Local Generative Surrogates, S. Shirobokov, V. Belavin, M. Kagan, A. Ustyuzhanin, A. G. Baydin, https://arxiv.org/abs/2002.04632

Julien Donini – Differentiable programming and detector design optimization

Muon shielding in SHIP

Minimize muon background fluxes in the SHIP steel magnet by varying its geometry



Local generative surrogate solution is shorter and has lower mass than other proposal, hence improving efficacity of the experiment and reducing its cost



Julien Donini – Differentiable programming and detector design optimization

Machine-Learning Optimized Design of Experiments MODE Collaboration

https://mode-collaboration.github.io

A. G. Baydin⁵, A. Boldyrev⁴, K. Cranmer⁸, P. de Castro Manzano¹, T. Dorigo¹, C. Delaere², D. Derkach⁴, J. Donini³, A. Giammanco², J. Kieseler⁷, G. Louppe⁶, L. Layer¹, P. Martinez Ruiz del Arbol⁹, F. Ratnikov⁴, G. Strong¹, M. Tosi¹, A. Ustyuzhanin⁴, P. Vischia², H. Yarar¹ + more members that joined recently

1 INFN, Sezione di Padova (and associates from Padova and Naples Universities), Italy 2 Université Catholique de Louvain, Belgium

- 3 Université Clermont Auvergne, France
- 4 Laboratory for big data analysis of the Higher School of Economics, Russia
- 5 University of Oxford
- 6 Université de Liege
- 7 CERN
- 8 New York University
- 9 IFCA



MODE ultimate goals

The target of **MODE** is to design and offer to the community a scalable, versatile **architecture** that can provide **end-to-end optimization of particle detectors**, proving it on a number of **different applications** across different **domains**



Toward Machine Learning Optimization of Experimental Design, Nuclear Physics News, 2021

Differentiable programming for muography

Tomography: exploit **atmospheric muon** flux to **map** the interior of **objects**

Muon absorption

Muon scattering



[images : A. Giammanco]

Muon tomography

Volume with unknown composition sandwiched between detectors



$High X_0 = Iow$	Low $X_0 = high$
scattering	scattering

Infer X_0 (radiation length) of volume by measuring **muon scattering**

How should detectors be positionned for best performances ?

- i.e Muon detection accuracy, resolution on X₀
- But also: cost, size, ...

[see G. Strong talk]

TomOpt: Tomography Optimization

Muon scan of volume of unknown density

- **Detector layers** are regions of space containing a fixed number of panels
 - Each panel exists and has a fixed resolution and efficiency
 - The (x,y,z) position and xy-span are parameters to be optimised
 - **Cost** of each panel scales with its area

Inference is fully differentiable w.r.t. detector parameters



Python package for differential optimisation of muon-tomography detectors **G.Strong**, T.Dorigo, F.Fanzago, A.Giammanco, M.Lagrange, M.Lamparth, F.Nardi, and P.Vischia

TomOpt: Tomography Optimization



Example volume

- Block of lead $(X_0 = 0.005612m)$
- Surrounded by beryllium ($X_0 = 0.3528m$)

Prediction based on 10k muons

- Lead block clearly visible
- The X0 predictions are **biased** due the assigning of the entirety of the muon scattering to a single point
- high z uncertainty in scatter location causes 'ghosting' above and below

[G. Strong]

TomOpt: Tomography Optimization

Loss contains:

- Predictive performance
- And detector cost

Live **feedback** of optimisation

Both loss and detector states



Important milestone for this use case

[G. Strong]

Calorimetry for Muon Collider

ECAL geometry optimization studies

- Focus on role of **Beam Induced Background** (mostly photons)
- Mitigate impact of BIB: granularity, segmentation, reconstruction, timing
- Aim: use Diff. Programming to **optimize geometry configuration**





Detector scheme

ECAL barrel and shielding "nozzle"

https://muoncollider.web.cern.ch/design/muon-collider-detector

Julien Donini – Differentiable programming and detector design optimization

Differentiable programming paradigm opens to many different applications

For HEP: end-to-end optimization of analysis, simulators, detectors, ...

MODE collaboration: ML optimization of **detector design**

- Several projects: muon tomography (advanced), muon collider detector shielding (ongoing), Hybrid calorimeter (staring) + others considered
- We know this is a **challenging** and **ambitious** task !
- Objective is **not to substitute experts** in detector design
- **Domain knowledge crucial** in setting up analysis workflow
- Consider joining and bring you use case

Backup material



Surrogates for differentiability



- Run simulator many times
- Generate a (large) dataset of input output pairs capturing simulator's behavior



• Use the dataset to learn a differentiable approximation of the simulator (e.g., a deep generative model)



[slides G. Baydin]

Surrogates for differentiability

Algorithm 1 Local Generative Surrogate Optimization (L-GSO) procedure

- **Require:** number N of ψ , number M of x for surrogate training, number K of x for ψ optimization step, trust region U_{ϵ} , size of the neighborhood ϵ , Euclidean distance d
- 1: Choose initial parameter ψ
- 2: while ψ has not converged do
- 3: Sample ψ_i in the region U^{ψ}_{ϵ} , i = 1, ..., N
- 4: For each ψ_i , sample inputs $\{x_j^i\}_{j=1}^M \sim q(x)$
- 5: Sample $M \times N$ training examples from simulator $y_{ij} = F(x_j^i; \psi_i)$
- 6: Store $\boldsymbol{y}_{ij}, \boldsymbol{x}_j^i, \boldsymbol{\psi}_i$ in history H $i = 1, \dots, N; j = 1, \dots, M$
- 7: Extract all y_l, x_l, ψ_l from history H, iff $d(\psi, \psi_l) < \epsilon$
- 8: Train generative surrogate model $S_{\theta}(\boldsymbol{z}_l, \boldsymbol{x}_l; \boldsymbol{\psi}_l)$, where $\boldsymbol{z}_l \sim \mathcal{N}(0, 1)$
- 9: Fix weights of the surrogate model θ
- 10: Sample $\bar{\boldsymbol{y}}_k = S_{\theta}(\boldsymbol{z}_k, \boldsymbol{x}_k; \boldsymbol{\psi}), \boldsymbol{z}_k \sim \mathcal{N}(0, 1),$ $\boldsymbol{x}_k \sim q(\boldsymbol{x}), \ k = 1, \dots, K$

11:
$$\nabla_{\boldsymbol{\psi}} \mathbb{E}[\mathcal{R}(\bar{\boldsymbol{y}})] \leftarrow \frac{1}{K} \sum_{k=1}^{K} \frac{\partial \mathcal{R}}{\partial \bar{\boldsymbol{y}}_k} \frac{\partial S_{\theta}(\boldsymbol{z}_k, \boldsymbol{x}_k; \boldsymbol{\psi})}{\partial \boldsymbol{\psi}}$$

12:
$$\boldsymbol{\psi} \leftarrow \text{SGD}(\boldsymbol{\psi}, \nabla_{\boldsymbol{\psi}} \mathbb{E}[\mathcal{R}(\bar{\boldsymbol{y}})])$$

13: end while

$$egin{aligned} oldsymbol{\psi}^* &= rgmin_{oldsymbol{\psi}} \mathbb{E}[\mathcal{R}(oldsymbol{y})] = rgmin_{oldsymbol{\psi}} \int \mathcal{R}(oldsymbol{y}) p(oldsymbol{y} | oldsymbol{x}; oldsymbol{\psi}) q(oldsymbol{x}) doldsymbol{x} doldsymbol{y} \ &pprox rgmin_{oldsymbol{\psi}} \ rac{1}{N} \sum_{i=1}^N \mathcal{R}(F(oldsymbol{x}_i; oldsymbol{\psi})) \end{aligned}$$

$$abla_{oldsymbol{\psi}} \mathbb{E}[\mathcal{R}(oldsymbol{y})] pprox rac{1}{N} \sum_{i=1}^{N}
abla_{oldsymbol{\psi}} \mathcal{R}(S_{ heta}(oldsymbol{z}_i,oldsymbol{x}_i;oldsymbol{\psi}))$$

TomOpt: end-to-end optimization



- We can now compute the analytical effects of detector parameters (position, size, resolution, etc.) on system outputs Kne
- Now express the desired task as a loss function
 - E.g. error on X₀ predictions, detector costs, time to achieve desired resolution
- We can now backpropagate the loss gradient to detector parameters and optimise via gradient descent
 - Just like a neural network



[G. Strong talk]

Joining MODE ?

If you are doing experimental research in HEP, astro-HEP, neutrino physics, or high-energy nuclear physics, or if you are working at spin-offs involving, *e.g.*, muon tomography, hadron therapy, or other endeavours which operate with instruments that extract information from the interaction of energetic radiation with matter, you are very likely to have a use case – a system liable to benefit from a study with differentiable programming.

The idea of MODE is to bring together ML experts who are developing the interfaces for these applications, with the researchers who have problems to solve in their area of interest

We cannot offer a solution to any given problem (we lack the manpower to work on-demand), but together we may work toward it

□ Consider joining MODE, and bring your use case!

Do I have a use case checklist:

Are you involved in the design, assembly, or upgrade of an instrument?

Can you specify one or a set of desirable scientific goals from its use?

Are those goals achieved through information processing?

If your answers to all are «yes», you have something to optimize and chances are this can't be done without a deep learning model of the full information extraction chain.

[slide T. Dorigo]

Muon collider detector



https://muoncollider.web.cern.ch/design/muon-collider-detector