Al-assisted design of experiments at the frontiers of computation

methods and new perspectives

ALPACA Workshop, ECT^{*}, Trento

Dr. Pietro Vischia pietro.vischia@cern.ch @pietrovischia



If you are reading this as a web page: have fun! If you are reading this as a PDF: please visit

https://www.hep.uniovi.es/vischia/persistent/2023-11-

22_AIAssistedDesignAtFrontiersOfComputationAtALPACAWorkshop_vischia.html

to get the version with working animations

Complex Experiments



Complex Data



Likelihood and information

• Data sample X_{obs}

$$\mathcal{L}(X; heta):=P(X| heta)|_{X_{obs}}$$

- The Likelihood Principle: The likelihood function $L(\vec{x}; \theta)$ contains all the information available in the data sample relevant for the estimation of θ
 - ✓ Bayesian statistics
 - X Frequentist statistics

$$I(heta) = -Eigg[\Big(rac{\partial^2}{\partial heta^2} ln L(X; heta) \Big)^2 | heta_{true} igg]$$



We like low-dim summaries



• Physical observable for inference



Brain activity...



...approximated...



...using computers



Image from zomato, elaborated with image from twinkl

Santiago Ramón y Cajal



Santiago Ramón y Cajal







Computationally heavy



Simplified Neurons

Bulletin of Mathematical Biology Vol. 52, No. 1/2, pp. 99-115, 1990. Printed in Great Britain. 0092-8240/90\$3.00 + 0.00 Pergamon Press plc Society for Mathematical Biology

A LOGICAL CALCULUS OF THE IDEAS IMMANENT IN NERVOUS ACTIVITY*

WARREN S. MCCULLOCH AND WALTER PITTS University of Illinois, College of Medicine, Department of Psychiatry at the Illinois Neuropsychiatric Institute, University of Chicago, Chicago, U.S.A.

Because of the "all-or-none" character of nervous activity, neural events and the relations among them can be treated by means of propositional logic. It is found that the behavior of every net can be described in these terms, with the addition of more complicated logical means for nets containing circles; and that for any logical expression satisfying certain conditions, one can find a net behaving in the fashion it describes. It is shown that many particular choices among possible neurophysiological assumptions are equivalent, in the sense that for every net behaving under one assumption, there exists another net which behaves under the other and gives the same results, although perhaps not in the same time. Various applications of the calculus are discussed.

Perceptrons

$$y=f\Big(b_i+\sum w_ix_i\Big)$$



Artificial Neural Networks



Learn in different ways



Gradient Descent

- Optimize/learn by finding the minimum of a function $\mathcal{L}:\mathbb{R}^n
 ightarrow\mathbb{R}$
- Nonconvex problems: saddle points, manifolds of minima
- Empirical risk minimization

 $\hat{L}(f) = rac{1}{n} \sum_{i=1}^n |f(x_i) - f^*(x_i)|^2$

• Generalization (for learning problems)





Backpropagation

$$\mathbf{J}(\mathbf{W}) = rac{1}{n}\sum_{i=1}^n \mathcal{L}(f(x^{(i)};\mathbf{W}),y^{*(i)}), \qquad \mathbf{W}^0 = argmin_{\mathcal{W}}\mathbf{J}(\mathbf{W})$$

 $\mathbf{W} \leftarrow \mathbf{W} + \eta rac{\partial \mathbf{J}(\mathbf{W})}{\partial \mathbf{W}}$



Derive



Automatic differentiation

has many names

- Automatic differentiation
- Algorithmic differentiation
- AD
- Autodiff
- Algodiff
- Autograd

Automatic differentiation

 $z(x,y) = 2x + x \sin(y) + y^3$



Forward mode

- To the extreme, $f:\mathbb{R}
 ightarrow\mathbb{R}^m$
- Evaluates $\left(\frac{\partial f_1}{\partial x}, \ldots, \frac{\partial f_m}{\partial x}\right)$

Reverse mode

- To the extreme, $f:\mathbb{R}^n
 ightarrow\mathbb{R}$
- Evaluate $abla f(\mathbf{x})(rac{\partial f}{\partial x_1},\ldots,rac{\partial f}{\partial x_n})$
- Computational cost of calculating $\mathbf{J}_f(\mathbf{x})$ for $f:\mathbb{R}^n o\mathbb{R}^m$ in $\mathbb{R}^n imes\mathbb{R}^m$

 $\mathcal{O}(n \operatorname{time}(f))$

 $\mathcal{O}(m \operatorname{time}(f))$

Forward and reverse (==backprop) modes

Primal: independent to dependent

Adjoint (derivatives): dependent to independent

$$y({f x})=2x_0+x_0\,sin(x_1)+x_1^3$$

<i>Fwd</i> <i>Primal</i> <i>Trace</i> Atomic operation	Value in $(1,2)$	Fwd Tangent Trace (set $\dot{x}_0 =$ 1 to compute $\frac{\partial y}{\partial x_0}$) Atomic operation	Value in $(1,2)$
$egin{array}{l} v_0 = x_0 \ v_1 = x_1 \end{array}$	$1 \\ 2$	$egin{array}{lll} \dot{v}_0 = \dot{x}_0 \ \dot{v}_1 = \dot{x}_1 \end{array}$	1 0
$egin{aligned} &v_2 = 2v_0 \ v_3 = \ sin(v_1) \ v_4 = \ v_0v_3 \ v_5 = v_1^3 \ v_6 = \ v_2 + \ v_4 + v_5 \end{aligned}$	$2 \\ 0.9093 \\ 0.9093 \\ 8 \\ 10.9093$	$egin{aligned} \dot{v}_2 &= 2\dot{v}_0 \ \dot{v}_3 &= \dot{v}_1 cos(v_1) \ \dot{v}_4 &= \dot{v}_0 v_3 + \ v_0 \dot{v}_3 \ \dot{v}_5 &= 3\dot{v}_1 v_1^2 \ \dot{v}_6 &= \dot{v}_2 + \dot{v}_4 + \ \dot{v}_5 \end{aligned}$	$\begin{array}{c} 2 \times 1 \\ 0 \times \\ -0.41 \\ 1 \times \\ 0.9093 + \\ 1 \times 0 \\ 3 \times 0 \times 4 \\ 2 + \\ 0.9093 + \\ 0 \end{array}$
$y = v_6$	10.9093	$\dot{y}=\dot{v}_{6}$	2.9093

<i>Fwd</i> <i>Primal</i> <i>Trace</i> Atomic operation	Value in $(1,2)$	Rev Adjoint Trace (set $\bar{y} =$ 1 to compute $\frac{\partial v}{\partial y}$) Atomic operation	Value in $(1,2)$
$egin{array}{l} v_0 = x_0 \ v_1 = x_1 \end{array}$	$1 \\ 2$	$ar{x}_0 = ar{v}_0 \ ar{x}_1 = ar{v}_1$	$2.9093 \\ 11.5839$
$egin{aligned} &v_2 = 2v_0 \ &v_3 = \ sin(v_1) \ &v_4 = \ &v_0v_3 \ &v_5 = v_1^3 \ &v_6 = \ &v_2 + \ &v_4 + v_5 \end{aligned}$	$2 \\ 0.9093 \\ 0.9093 \\ 8 \\ 10.9093$	$ar{v}_0 = ar{v}_0 + ar{v}_2 \partial v_2 / \partial v_0 \ ar{v}_0 = ar{v}_4 \partial v_4 / \partial v_0 \ ar{v}_1 = ar{v}_1 + ar{v}_3 \partial v_3 / \partial v_1 \ ar{v}_1 = ar{v}_5 \partial v_5 / \partial v_1 \ ar{v}_2 = ar{v}_6 \partial v_6 / \partial v_2 \ ar{v}_3 = ar{v}_4 \partial v_4 / \partial v_3 \ ar{v}_4 = ar{v}_6 \partial v_6 / \partial v_4 \ ar{v}_5 = ar{v}_6 \partial v_6 / \partial v_5$	$egin{array}{llllllllllllllllllllllllllllllllllll$

 $y = iew_{6}$ Vischi10.9993ed d \overline{w}_{6} grat \overline{y} he Frontiers of Computation - 2023.11.22 --- 23/108

Designed to be simple in software

```
import torch, math
x0 = torch.tensor(1., requires_grad=True)
x1 = torch.tensor(2., requires_grad=True)
p = 2*x0 + x0*torch.sin(x1) + x1**3
print(p)
p.backward()
print(x0.grad, x1.grad)
```

yielding

```
Primal: tensor(10.9093, grad_fn=<AddBackward0>)
Adjoint: tensor(2.9093) tensor(11.5839)
```

- Torch (and similar software) will correctly differentiate only when the atomic operations are supported within it
 - Common operations are overloaded (__mul__ rewritten by torch._mult_)
 - Operations from libraries (math.sin()) must be replaced by their differentiation-aware equivalents (torch.sin())

Differentiable Programming

Execute differentiable functions (programs) via automatic differentiation



Yann LeCun ♥ January 5, 2018 · ⑤

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

Yeah, Differentiable Programming is little more than a rebranding of the modern collection Deep Learning techniques, the same way Deep Learning was a rebranding of the modern incarnations of neural nets with more than two layers.

But the important point is that people are now building a new kind of software by assembling networks of parameterized functional blocks and by training them from examples using some form of gradient-based optimization.

An increasingly large number of people are defining the networks procedurally in a data-dependent way (with loops and conditionals), allowing them to change dynamically as a function of the input data fed to them. It's really very much like a regular progam, except it's parameterized, automatically differentiated, and trainable/optimizable. Dynamic networks have become increasingly popular (particularly for NLP), thanks to deep learning frameworks that can handle them such as PyTorch and Chainer (note: our old deep learning framework Lush could handle a particular kind of dynamic nets called Graph Transformer Networks, back in 1994. It was needed for text recognition).

People are now actively working on compilers for imperative differentiable programming languages. This is a very exciting avenue for the development of learning-based AI.

Important note: this won't be sufficient to take us to "true" AI. Other concepts will be needed for that, such as what I used to call predictive learning and now decided to call Imputative Learning. More on this later....



186 Comments 464 Shares

Maps as a Tool of Understanding



"Representation" simplifies tasks





Animation and picture from FastForward Labs

Impressive results



Impressive results



• Busco colaboraciones para aplicaciones médicas de inteligencia artificial

Go to INFERNO: syst-aware inference opt.



(b) profile-likelihood comparison Pietro Vischia - Al-assisted design at the Frontiers of Computation - 2023.11.22 --- 30 / 108

Measurement-aware analysis opt.



neos

Measurement-aware detector opt.!

- Joint optimization of design parameters w.r.t. inference made with data
- MODE White Paper, 10.1016/j.revip.2023.100085 (2203.13818), 117-pages document, physicists + computer scientists (published last month!!!)



Guarantee feasibility within constraints

- Monetary cost
- Case-specific technical constraints

 $\mathcal{L}_{ ext{cost}} = c(heta, \phi)$

- θ : local, specific to the technology used (e.g. active components material)
- ϕ : global, describing overall detector conception (e.g. number, size, position of detector modules)
- Fixed costs can be added separately to the loss function

In general



Thrive in asymmetries


Large gains to be had

- MUonE: proposed 150 GeV muon beam experiment to be built at CERN
 - \circ Measure precisely the q^2 differential cross section in electron-muon scattering
 - 40 tracking stations and a calorimeter
- Dramatic improvement in the resolution on q^2 even from a simple grid search



Assist the physicist with a landscape of solutions

- Cannot parameterize everything
- The optimal solution: unrealistic
- Provide feasible solutions near optimality
- The physicist will fine tune



How far from optimality?

- Can we define in a general way an acceptable increase in loss?
 - Tradeoff performance/cost
- For sure we can regularize the loss landscape to select our scale of interest



Method of choice depends on scale



- 1. Grid/random search
- 2. Bayesian opt, simulated annealing, genetic algos, ...
- 3. Gradient-based optimization (Newton, BFGS, gradient descent, ...)

Cosmic rays from supernovae

• High-energy primary cosmic rays produced by supernovae



Primary cosmic rays

- 89% hydrogen nuclei (protons)
- the rest is Helium, Carbon, Oxigen, and other less abundant elements



Muons from cosmic rays

- Cosmic ray muons produced when primary cosmic rays impact with earth's atmosphere
- 1990, Gaisser formula for flux at sea level

$$\frac{dI_{\mu}}{dE_{\mu}} = 0.14 \left(\frac{E_{\mu}}{GeV}\right)^{-2.7} \left[\frac{1}{1 + \frac{1.1E_{\mu}\cos\theta}{115\ GeV}} + \frac{0.054}{1 + \frac{1.1E_{\mu}\cos\theta}{850\ GeV}}\right]$$

- Valid only if:
 - \circ Earth curvature negligible ($heta < 70 ext{ deg}$)
 - $\circ~$ Muon decay negligible ($E_{\mu}> 100/cos heta$ GeV)



Fig. 1. The relation of the observed zenith angle of muons, θ^* , to the zenith angle at the muon production point in the atmosphere, θ . R is the radius of the Earth. Adopted from [3][4]

Muons from cosmic rays

- Improved formula by Guan et al. (2015)
- Account for Earth's curvature
- θ at ground and θ^* at production differ

$$\cos\theta^* = \sqrt{\frac{(\cos\theta)^2 + P_1^2 + P_2(\cos\theta)^{P_3} + P_4(\cos\theta)^{P_5}}{1 + P_1^2 + P_2 + P_4}}$$

$$\frac{1}{\frac{P_1 + P_2 + P_3 + P_4 + P_5}{0.102573 - 0.068287 + 0.958633 + 0.0407253 + 0.817285}}$$

• Correction at low energies

$$\begin{split} \frac{dI_{\mu}}{dE_{\mu}} &= 0.14 \left[\frac{E_{\mu}}{GeV} \left(1 + \frac{3.64GeV}{E_{\mu}(\cos\theta^{*})^{1.29}} \right) \right]^{-2.7} \\ &\times \left[\frac{1}{1 + \frac{1.1E_{\mu}\cos\theta^{*}}{115GeV}} + \frac{0.054}{1 + \frac{1.1E_{\mu}\cos\theta^{*}}{850GeV}} \right] \end{split}$$



Fig. 1. The relation of the observed zenith angle of muons, θ^* , to the zenith angle at the muon production point in the atmosphere, θ . R is the radius of the Earth. Adopted from [3][4]

Absorption or scattering

- Absorption: measure missing flux
 - Pyramids, volcanoes...
- Scattering: measure deflection of muon trajectories
 - Containers, furnaces, statues...



Heart of steel

- Map electrode position in electric furnaces
- Wear of inner walls of pipes



Portable muon detectors

- (A) Maximize angular resolution
- (B) Maximize acceptance
- (C) 3D tomography
- (D) 2D projections in orthogonal planes



Border control

$$L = \left(1 + e^{k(c_ heta - c_0)}
ight) \sum_z \left[w_{imp}(Z) m^{ ext{concealed}}_{50,lpha}[s(Z)]
ight]$$



Scanning a passive volume

- Want to infer properties (e.g. 3D map of elemental composition) of unknown volume
 - Shipping container, archeological site, nuclear waste dump, industrial machinery, etc.
- Muons from cosmic rays traverse us all the time
 - On average, 1 muon per cm^2 per minute
 - \circ Change in kinematics provides handle for inference on X_0



Domain knowledge is not enough

- Domain knowledge typically provides heuristics based on proxy objectives
- Will likely have a budget
 - Money, heat, power, positioning of detectors, imaging time...
- Will likely have varying purposes
 - Today want to spot uranium, tomorrow e.g. drugs



Example 1: Muons measured precisely but less efficiently Example 2: Muons measured less precisely but more efficiently

TomOpt

- Differential optimization of muon-tomography detectors (ongoing project)
 - Giles C. Strong, Maxime Lagrange, Aitor Orio, Anna Bordignon, Florian Bury, Tommaso Dorigo, Andrea Giammanco, Mariam Heikal, Jan Kieseler, Max Lamparth, Pablo Martínez Ruíz del Árbol, Federico Nardi, Pietro Vischia, Haitham Zaraket
 - 2309.14027 submitted to journal, shorter version accepted by NeurIPS MLPS Workshop!
 - Modular design in python, autodiff via PyTorch



Figure from the TomOpt proje Gradient descent

Pietro Vischia - Al-assisted design at the Frontiers of Computation - 2023.11.22 --- 50 / 108

Muon Generation

- Formulas by 2015 and 2016 models
- Tensor of muons $(x,y,p, heta,\phi)$
 - $\circ \ heta, p$ from flux model
 - $\circ x, y$ from ranges
 - $\circ \phi$ uniformly in $[0,2\pi]$
- Code handles many muons at once (MuonBatch)
 - Propagate the muon position (can snapshot to track)
 - $\circ~$ Scatter ($dxy, d heta d\phi$) at each step



Volume Specification

- Volume made up of stacked layers in *z*
- Passive layers scatter muon
 - PDG and GEANT models both available
 - Voxelized passive layers (x, y)
- Active layers record muon hits
 - Parameterized efficiency and resolution (cost per m^2 , physics constraints)
 - Budget is a volume attribute, and is assigned to detector layers



Panel specification



Figure 2: Example of detector panel modeling with sigmoid function used during optimisation (blue) and with rectangle function used during validation (orange).

Monte Carlo Truth

- Per each scenario, can build voxelized random volumes
 - Each voxel can be a different material
 - Next: material mixture per voxel



Make muon hits differentiable

- Associate a distribution to resolution and efficiency
 - e.g. Gaussian centered on panel and width equal to panel span
 - p.d.f. of the muon position is now differentiable
- Further generalization: Gaussian Mixture models



From hits to tracks

- Analytic maximum likelihood fit
 - considering uncertainty and efficiency of hits
 - fully differentiable w.r.t. detector parameters
- Provides track parameters and their uncertainties



- POCA (POint of Closest Approach)
 - assume one scattering in one point
 - \circ invert model to compute X_0



$\circ~$ average X_0 per voxel



• Assume one scattering \rightarrow bias!



Extended POCA

• Assume one scattering \rightarrow bias!



Detector Optimization in TomOpt





ume.

Heavy Metal

- Transport liquid steel to fill moulds
 - Lack of enough metal: moulds not filled
 - Too much metal: remnants, scraps
 - Slags hides metal from optical inspection





Heavy Metal



Encouraging results

- But POCA provides a biased estimator, with bias increasing with fill level (more metal, more bias)
- Debiasing via parametric correction



Optimize crappy detector



(b) Detector configuration after stage one optimisation process.

Optimize crappy detector



Refine a good detector



(a) Detector configuration after stage one optimi- (b) Detector configuration after stage two optimisation process. sation process.

Refine a good detector



Recover human baselines



Method of choice depends on scale



- 1. Grid/random search
- 2. Bayesian opt, simulated annealing, genetic algos, ...
- 3. Gradient-based optimization (Newton, BFGS, gradient descent, ...)

Need for new paradigma

• If you are interested in Neuromorphic computing or Quantum computing, drop me a line!



Technology readiness?

Thank you!



Backup
It's all about summaries

- Promising performance
- Muon track quantities differentiable
 - can compute uncertainties due to spatial resolution
 - useful for aggregating
- X_0 predictions biased and inaccurate
 - Uranium is denser than other materials
 - \circ Under U hypothesis, large difference between mean of the n lowest X_0 voxels and mean X_0 of remaining voxels
 - One-dimensional summary statistic!



• Ultimately, the summary can be learned with graph neural networks)

• Use Quasim et al: distance-weighted representations, information-sharing along learnable connections



relevant muons.

Represent voxels

voxels.

• Use Quasim et al: distance-weighted representations, information-sharing along learnable connections



Predict volume (binary classification: there or not there)

• Use Quasim et al: distance-weighted representations, information-sharing along learnable connections



Imaging behaviour

Infer the volume

- Expectation maximization: iteratively
 - estimate scatter density based on current estimate of the image
 - update estimate of the image based on the estimated scatter density



Summary of the EM algorithm for muon tomography. Schultz et al. 2007



Figure 2: Illustration of commonly used algorithms for reconstructing images in muon scattering tomography. Barnes et al. 2023

EM preliminary performance



Volume inference

- The main boundary is that inference algorithm must be differentiable
- Basic approach of inverting scatter model to compute X_0 is highly biased
- Maybe a task-specific summary statistic would work better



Pipeline



Encouraging results

- Inference via low-dim summaries
- E.g. identify uranium in container





Cosmic rays from supernovae

• High-energy primary cosmic rays produced by supernovae



Primary cosmic rays

- 89% hydrogen nuclei (protons)
- the rest is Helium, Carbon, Oxigen, and other less abundant elements



Muons from cosmic rays

- Cosmic ray muons produced when primary cosmic rays impact with earth's atmosphere
- 1990, Gaisser formula for flux at sea level

$$\frac{dI_{\mu}}{dE_{\mu}} = 0.14 \left(\frac{E_{\mu}}{GeV}\right)^{-2.7} \left[\frac{1}{1 + \frac{1.1E_{\mu}\cos\theta}{115\ GeV}} + \frac{0.054}{1 + \frac{1.1E_{\mu}\cos\theta}{850\ GeV}}\right]$$

- Valid only if:
 - Earth curvature negligible ($heta < 70 ext{ deg}$)
 - $\circ~$ Muon decay negligible ($E_{\mu} > 100/cos heta$ GeV)



Fig. 1. The relation of the observed zenith angle of muons, θ^* , to the zenith angle at the muon production point in the atmosphere, θ . R is the radius of the Earth. Adopted from [3][4]

Muons from cosmic rays

- Improved formula by Guan et al. (2015)
- Account for Earth's curvature
- θ at ground and θ^* at production differ

$$\cos\theta^* = \sqrt{\frac{(\cos\theta)^2 + P_1^2 + P_2(\cos\theta)^{P_3} + P_4(\cos\theta)^{P_5}}{1 + P_1^2 + P_2 + P_4}}$$

$$\frac{1}{\frac{P_1 + P_2 + P_3 + P_4 + P_5}{0.102573 - 0.068287 + 0.958633 + 0.0407253 + 0.817285}}$$

• Correction at low energies

$$\begin{split} \frac{dI_{\mu}}{dE_{\mu}} &= 0.14 \left[\frac{E_{\mu}}{GeV} \left(1 + \frac{3.64GeV}{E_{\mu}(\cos\theta^{*})^{1.29}} \right) \right]^{-2.7} \\ &\times \left[\frac{1}{1 + \frac{1.1E_{\mu}\cos\theta^{*}}{115GeV}} + \frac{0.054}{1 + \frac{1.1E_{\mu}\cos\theta^{*}}{850GeV}} \right] \end{split}$$



Fig. 1. The relation of the observed zenith angle of muons, θ^* , to the zenith angle at the muon production point in the atmosphere, θ . R is the radius of the Earth. Adopted from [3][4]

Absorption or scattering

- Absorption: measure missing flux
 - Pyramids, volcanoes...
- Scattering: measure deflection of muon trajectories
 - Containers, furnaces, statues...



Heart of steel

- Map electrode position in electric furnaces
- Wear of inner walls of pipes



Portable muon detectors

- (A) Maximize angular resolution
- (B) Maximize acceptance
- (C) 3D tomography
- (D) 2D projections in orthogonal planes



Border control

$$L = \left(1 + e^{k(c_ heta - c_0)}
ight) \sum_z \left[w_{imp}(Z) m^{ ext{concealed}}_{50,lpha}[s(Z)]
ight]$$



Scanning a passive volume

- Want to infer properties (e.g. 3D map of elemental composition) of unknown volume
 - Shipping container, archeological site, nuclear waste dump, industrial machinery, etc.
- Muons from cosmic rays traverse us all the time
 - On average, 1 muon per cm^2 per minute
 - \circ Change in kinematics provides handle for inference on X_0



Domain knowledge is not enough

- Domain knowledge typically provides heuristics based on proxy objectives
- Will likely have a budget
 - Money, heat, power, positioning of detectors, imaging time...
- Will likely have varying purposes
 - Today want to spot uranium, tomorrow e.g. drugs



Example 1: Muons measured precisely but less efficiently Example 2: Muons measured less precisely but more efficiently

TomOpt

- Differential optimization of muon-tomography detectors (ongoing project)
 - Giles C. Strong, Tommaso Dorigo, Andrea Giammanco, Pietro Vischia, Jan Kieseler, Maxime Lagrange, Mariam Safieldin, Federico Nardi, Anna Bordignon, Haitham Zaraket, Max Lamparth, Federica Fanzago, Oleg Savchenko, Nitesh Sharma
 - Modular design in python, autodiff via PyTorch



Muon Generation

- Formulas by 2015 and 2016 models
- Tensor of muons $(x,y,p, heta,\phi)$
 - $\circ \ heta, p$ from flux model
 - $\circ x, y$ from ranges
 - $\circ \phi$ uniformly in $[0,2\pi]$
- Code handles many muons at once (MuonBatch)
 - Propagate the muon position (can snapshot to track)
 - $\circ~$ Scatter ($dxy, d heta d\phi$) at each step



Volume Specification

- Volume made up of stacked layers in *z*
- Passive layers scatter muon
 - PDG and GEANT models both available
 - Voxelized passive layers (x, y)
- Active layers record muon hits
 - Parameterized efficiency and resolution (cost per m^2 , physics constraints)
 - Budget is a volume attribute, and is assigned to detector layers



Monte Carlo Truth

- Per each scenario, can build voxelized random volumes
 - Each voxel can be a different material
 - Next: material mixture per voxel



Make muon hits differentiable

- Associate a distribution to resolution and efficiency
 - e.g. Gaussian centered on panel and width equal to panel span
 - p.d.f. of the muon position is now differentiable
- Further generalization: Gaussian Mixture models



From hits to tracks

- Analytic maximum likelihood fit
 - considering uncertainty and efficiency of hits
 - fully differentiable w.r.t. detector parameters
- Provides track parameters and their uncertainties



- POCA (POint of Closest Approach)
 - assume one scattering in one point
 - \circ invert model to compute X_0



 $\circ~$ average X_0 per voxel

Infer the volume

- Expectation maximization: iteratively
 - estimate scatter density based on current estimate of the image
 - update estimate of the image based on the estimated scatter density



Summary of the EM algorithm for muon tomography. Schultz et al. 2007



Figure 2: Illustration of commonly used algorithms for reconstructing images in muon scattering tomography. Barnes et al. 2023

EM preliminary performance



Volume inference

- The main boundary is that inference algorithm must be differentiable
- Basic approach of inverting scatter model to compute X_0 is highly biased
- Maybe a task-specific summary statistic would work better



It's all about summaries

- Promising performance
- Muon track quantities differentiable
 - can compute uncertainties due to spatial resolution
 - useful for aggregating
- X_0 predictions biased and inaccurate
 - Uranium is denser than other materials
 - \circ Under U hypothesis, large difference between mean of the n lowest X_0 voxels and mean X_0 of remaining voxels
 - One-dimensional summary statistic!



• Ultimately, the summary can be learned with graph neural networks)

• Use Quasim et al: distance-weighted representations, information-sharing along learnable connections



relevant muons.

Represent voxels

voxels.

• Use Quasim et al: distance-weighted representations, information-sharing along learnable connections



Predict volume (binary classification: there or not there)

• Use Quasim et al: distance-weighted representations, information-sharing along learnable connections



Imaging behaviour

Optimization

- Regular gradient descent of a loss function
 - account for cost of the detector and other constraints
 - standard optimisers to update detector parameters



Pipeline



Encouraging results

- Inference via low-dim summaries
- E.g. identify uranium in container





The MODE Collaboration

https://mode-collaboration.github.io/

• Joint effort (created 11.2020) of particle physicists, nuclear physicists, astrophysicists, and computer scientists

At INFN and Università of Padova Dr. Tommaso Dorigo, Dr. Pablo De Castro Manzano, Dr. Federica Fanzago, Dr. Lukas Layer, Dr. Giles Strong, Dr. Mia Tosi, and Dr. Hevjin Yarar At Université catholique de Louvain Dr. Andrea Giammanco, Prof. Christophe Delaere, Mr. Maxime Lagrange, and Dr. Pietro Vischia At Université Clermont Auvergne, Prof. Julien Donini, and Mr. Federico Nardi (joint with Universitá di Padova) At the Higher School of Economics of Moscow, Prof. Andrey Ustyuzhanin, Dr. Alexey Boldvrey, Dr. Denis Derkach, and Dr. Fedor Ratnikov At the Instituto de Física de Cantabria. Dr. Pablo Martínez Ruíz del Árbol At CERN, Dr. Jan Kieseler, Dr. Sofia Vallecrosia At University of Oxford Dr. Atilim Gunes Baydin At New York University Prof. Kyle Cranmer At Université de Liège Prof. Gilles Louppe At GSI/FAIR Dr. Anastasios Belias At Rutgers University Dr. Claudius Krause At Uppsala Universitet Prof. Christian Glaser At TU-München, Prof. Lukas Heinrich and Mr. Max Lamparth At Durham University Dr. Patrick Stowell At Lebanese University Prof. Haitham Zaraket At Technische Universitäet Kaiserslautern Mr. Max Aehle, Prof. Nicolas Gauger, Dr. Lisa Kusch At Technische Universitäet Worms Prof. Ralf Keidel At Princeton University Prof. Peter Elmer At University of Washington Prof. Gordon Watts At SLAC Dr. Ryan Roussel The Scientific Coordinator of the MODE Collaboration is Dr. Tommaso Dorigo, INFN-Sezione di Padova The Steering Board of the MODE Collaboration includes: Prof. Julien Donini, UCA Dr. Tommaso Dorigo, INFN-PD

- Dr. Andrea Giammanco, UCLouvain
- Dr. Fedor Ratnikov, HSE
- Dr. Pietro Vischia, UCLouvain
Series of yearly workshop

- First installment in Louvain-la-Neuve (Belgium)
- Second installmentnt in Kolymbari (Greece)
 - 37 talks, 9 posters, one data challenge with prizes, recordings will be online soon
- You are all invited to the Third MODE Workshop, to be held in Princeton (USA)



If you are interested, join us!

https://mode-collaboration.github.io/