Toward the End-to-End Optimization of Detector Design with Differentiable Programming

Tommaso Dorigo

INFN, Sezione di Padova

Foreword/1

The recent publication of the 2020 update of the European Strategy for Particle Physics (EUSUPP) [1] encourages feasibility studies for new large, long-term projects which will once again push our technological skills to their limits.

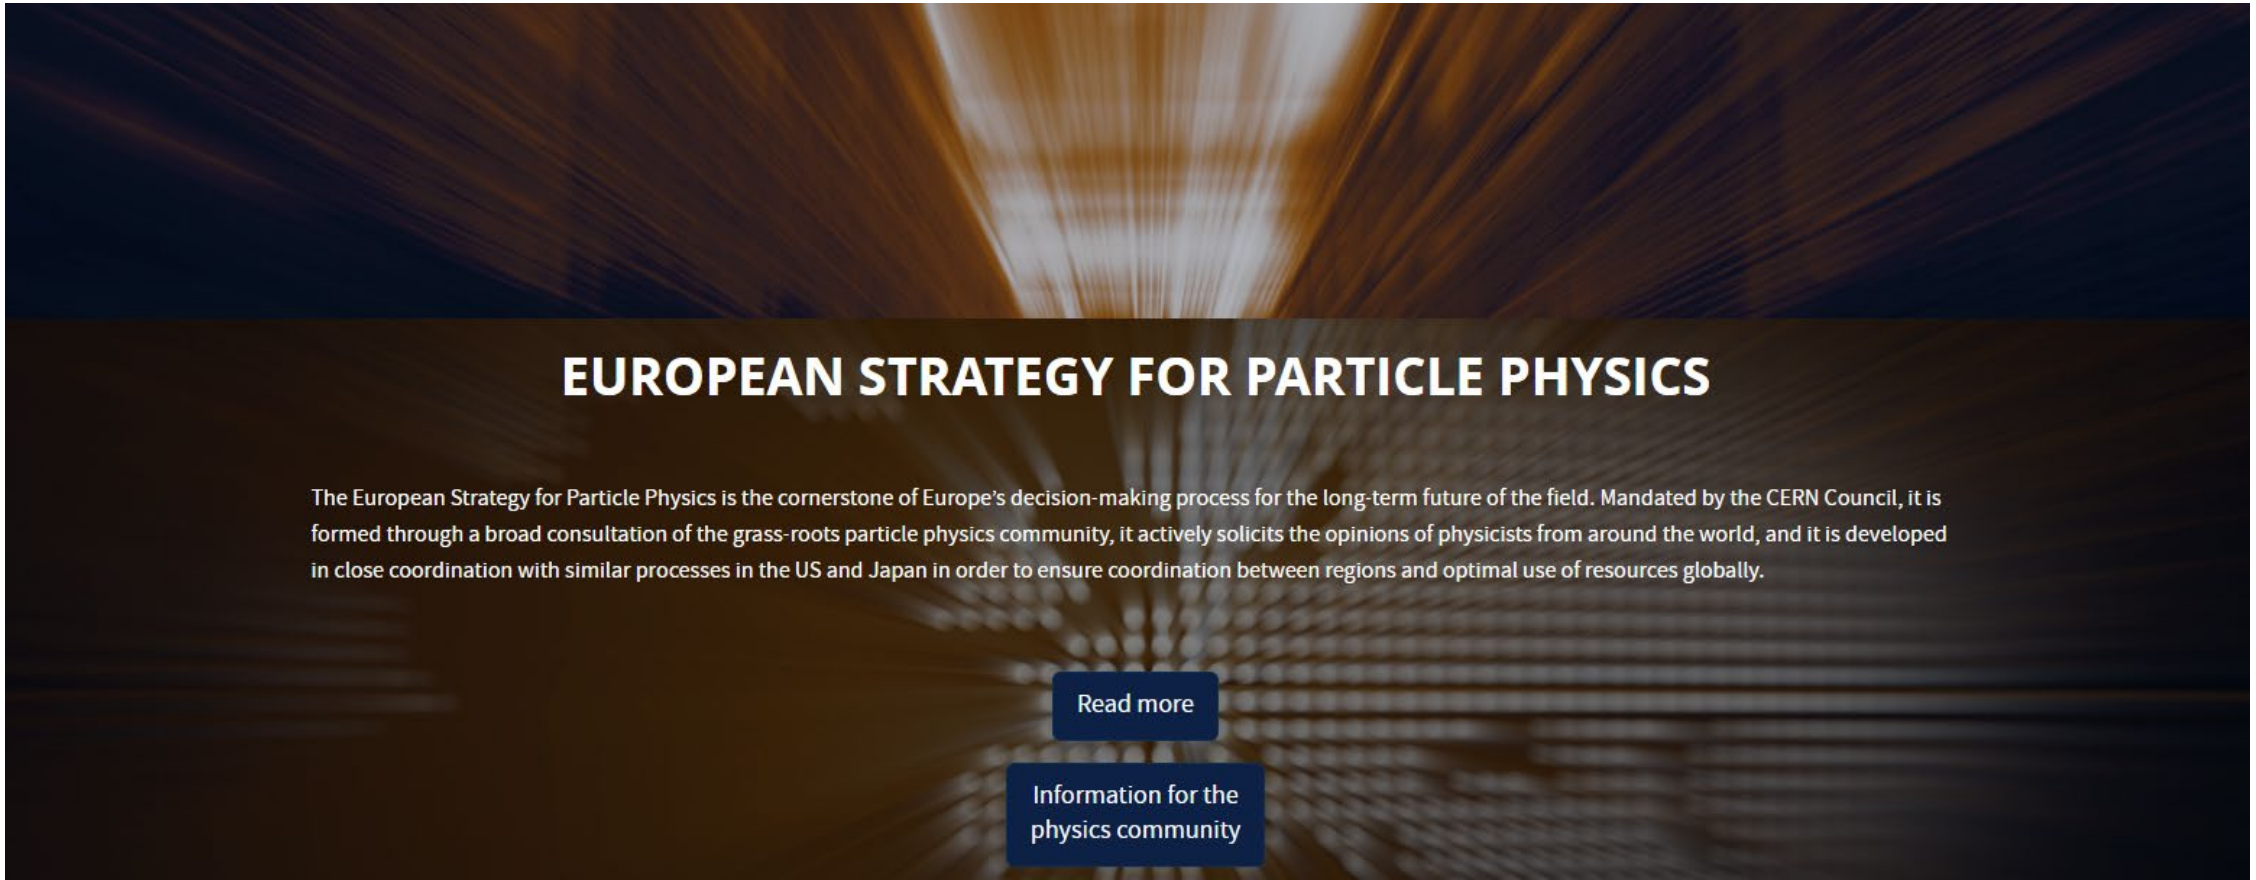

Foreword/2

The recent publication of the 2020 update of the European Strategy for Particle Physics (EUSUPP) [1] encourages feasibility studies for new large, long-term projects which will once again push our technological skills to their limits.

At the same time, humanity faces unprecedented global challenges (climate change, pandemics, overpopulation) which demand the use of our resources to seek solutions through applied science innovations, rather than investing in fundamental research.

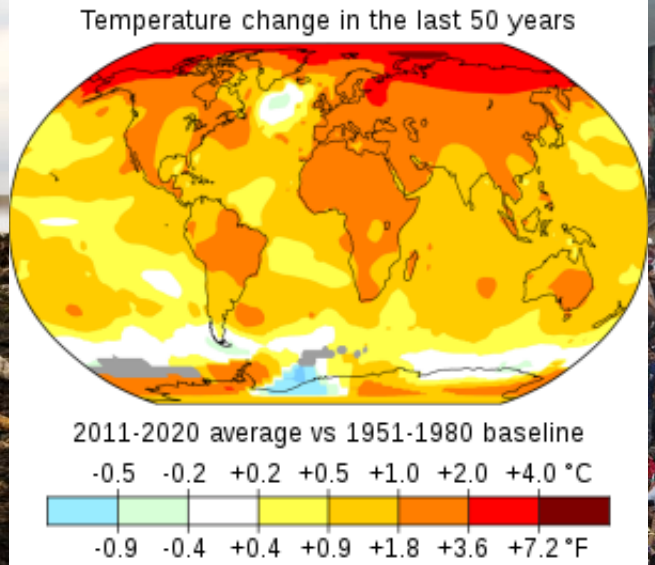

Foreword/3

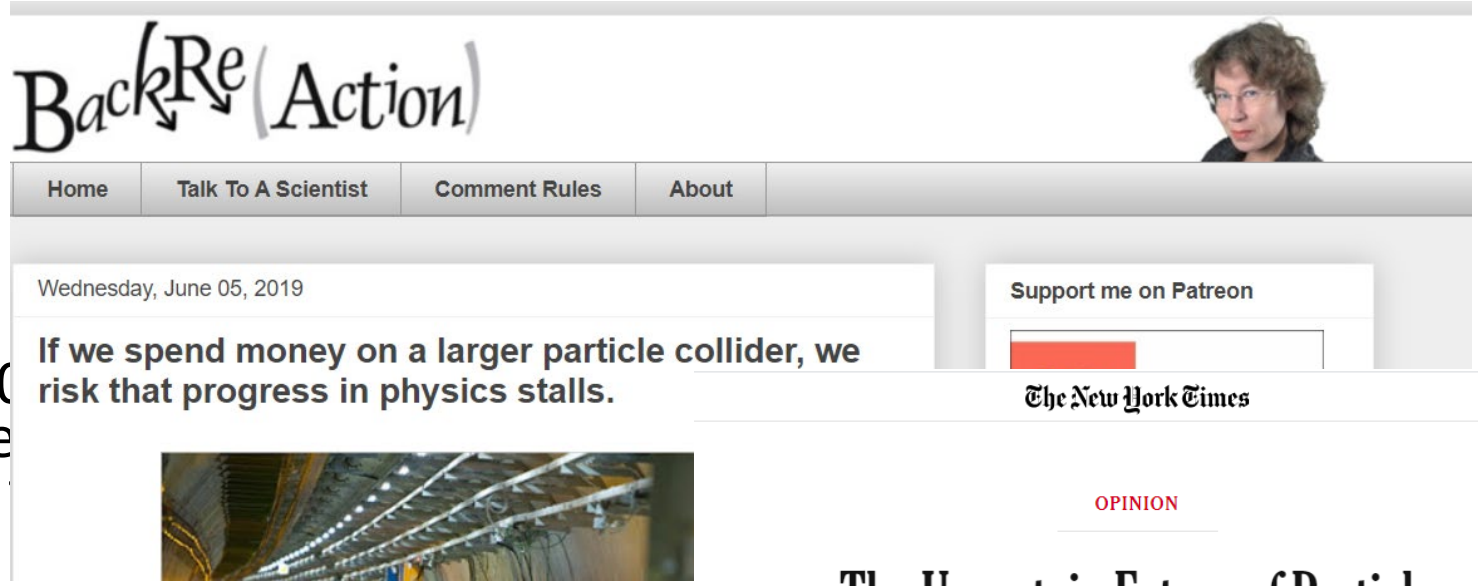

The recent publication of the 20
Physics (EUSUPP) [1] encourage
which will once again push our

At the same time, humanity faces unprecedented global c
pandemics, overpopulation) which demand the use of our
through applied science innovations, rather than investing

Furthermore, there are indications that wide sectors of society no longer consider a top priority the furthering of our understanding of matter at the smallest distance scales through projects that require large effort and significant funding [2].

In this situation, **ensuring the maximum exploitation of any resources spent on fundamental research is a moral imperative**, and it may be a key to ensure that the long-term projects envisioned by the EUSUPP may be undertaken and sustained.

The Status Quo

The design of new detectors in the past 50 years systematically leveraged the most performant available technologies for particle detection, often fostering significant further advancements and spin-offs [3], yet the crucial underlying paradigms of experimental design remained mostly unchallenged across decades, e.g.

- "**Track first, destroy later**": charged particles can be traced by their ionization in low-material-budget elements, while neutral ones must undergo destructive interactions in a calorimeter to be detected → standard setup of experiments in HEP and astro-HEP typically involves a low-material tracker followed by a thick calorimeter.
- Focus on significant **redundancy of detection systems**, ensuring robustness and enabling cross-calibration of resulting measurements
- **Symmetrical layouts** leveraging conservation laws, also simplifying reconstruction

While these choices have served us very well for a long time, **they are not meant to be "optimal":** *i.e.*, they **do not directly maximize a high-level utility function**, such as the highest discovery reach for a physical process, or measurement precision for a given physics parameter.

Optimal for what?

The reason why detectors are so damnedly complex is not only that the studied physics is hard: a lot has to do with Science being a competitive job.
We want to study *everything* and do it *better* than our predecessors

So, what does it mean for a detector to be *optimal*?

What loss function do we aim to minimize?

Does it make sense of speaking of a single utility function?

Concerning the last question: I am convinced that it does, and I will try to convince you, too, in the next few slides.

Recipe for a perfect dinner


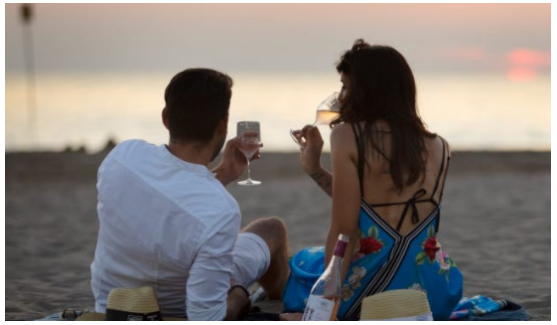







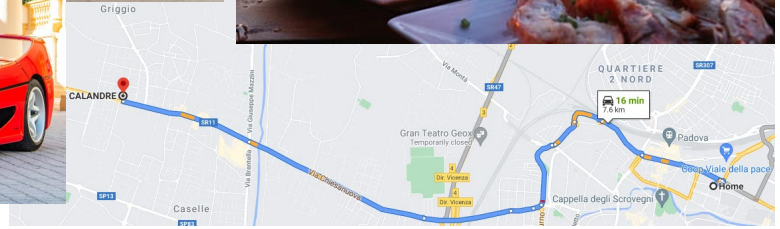

We are not alien to confidently taking complex decisions in a multi-objective space. We actually do it routinely...

Of course, we are not deterred by knowing that our optimization target is not universal!

Recipe for a perfect trigger

Similarly, we have actually grown *used* to create multi-target optimization strategies, e.g., to allocate resources for the trigger menu of a collider detector!

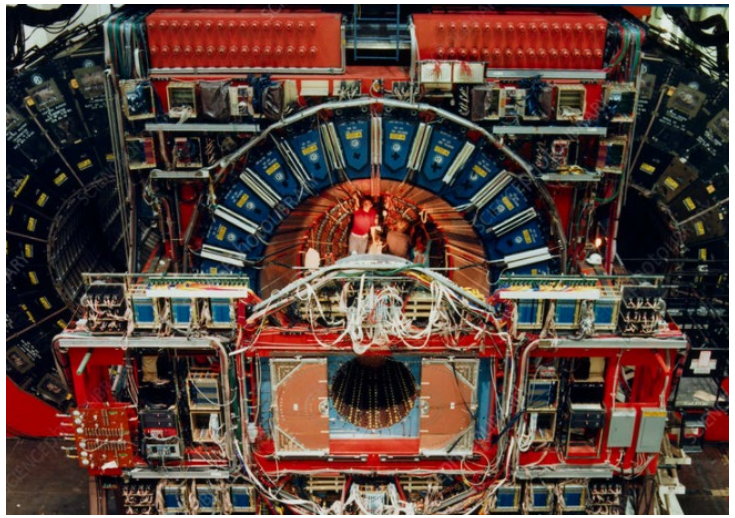

Consider CDF, Run 1 (1992-96): taking in a rate of 300 kHz of proton-antiproton collisions and having to select 50 Hz of writable data created some of the most heated scientifically-driven, rationally motivated, painfully well argumented debates I ever listened to. The top quark had to be discovered, but it was not the only goal of the experiment...

Recipe for a perfect detector

So - what does an end-to-end detector optimization looks like in blueprint?

1. Model as steep functions the cost of overriding budget and commissioning time
2. Assess the scientific impact of each achievable scientific results, optionally as a continuous function of their precision
3. Create a differentiable model of geometry, components, information-extraction procedures, and utility function
4. Construct a pipeline with those modules, enabling backpropagation and gradient descent functionality
5. **Let the chain rule of differential calculus do the hard work for you**!

Experiment proposals undergo a review, wherein the **utility function is well explicitated to the funding agents**... So, we invariably know what we are shooting at!

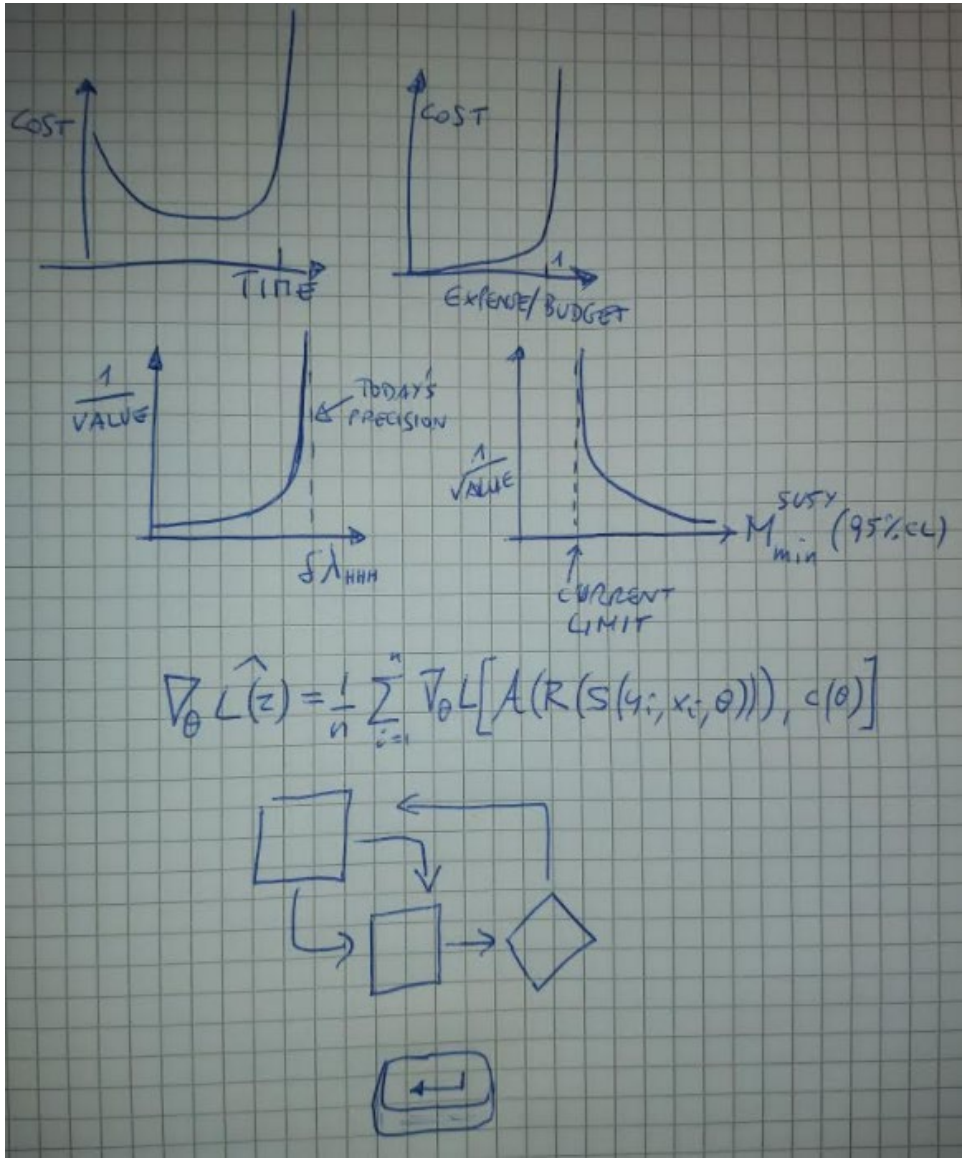

Makeshift surrogates of objectives

When we design the sensors for a tracking device, operate choices on budget allocations, define requirements for the various resolutions of detection elements, or choose composition and layout of active and passive material of calorimeter cells, **we are implicitly trying to find an optimal working point in a loosely-constrained feature space of hundreds of dimensions**. Such a task is clearly super-human.

Because of that, we set our aim on makeshift surrogates of our real objectives.

- E.g., we might desire our objective to be "*the highest precision on the Higgs boson self-couplings our budget can ensure*", but all we can do is stick to useful proxies suggested by past experience, and rather focus on the "*highest achievable energy resolution for isolated photons*", ignoring the rest of the parameter space
- In a neutrino detector this would sound as "the highest precision on theta_13 we can get", when the focus becomes instead maximizing the number of reconstructed interactions and reducing the background level.
- Our simulations only allow us to probe the result of specific choices, not to map interdependencies.

Makeshift surrogates of objectives

When we design the sensors for a tracking device, operate choices on budget allocations, define requirements for the various resolutions of detection elements, or choose composition and layout of active and passive material of calorimeter cells, **we are implicitly trying to find an optimal working point in a loosely-constrained feature space of hundreds of dimensions**.

THIS IS A SUPER-HUMAN TASK

Because of that, we set our aim on makeshift surrogates of our real objectives.

- E.g., we might desire our objective to be "*the highest precision on the Higgs boson self-couplings our budget can ensure*", but all we can do is stick to useful proxies suggested by past experience, and rather focus on the "*highest achievable energy resolution for isolated photons*", ignoring the rest of the parameter space
- In a neutrino detector, this would instead sound as "the highest precision on theta_13 we can get", when the focus becomes instead maximizing the number of interactions and reducing the background level.
- Our simulations only allow us to probe the result of specific choices, not to map interdependencies.

Evolving from this modus operandi to directly goal-informed decisions enables potentially enormous performance gains.

The design space is large – no, larger

New technological advancements are crucially enabling a better optimization of our instruments by reducing the cost of complex layouts.

- 3D printing of scintillation detectors are being explored for neutrino physics [4]
- Very thin layouts of resistive AC-coupled silicon detector elements may provide large gains in spatial and temporal resolution [5].

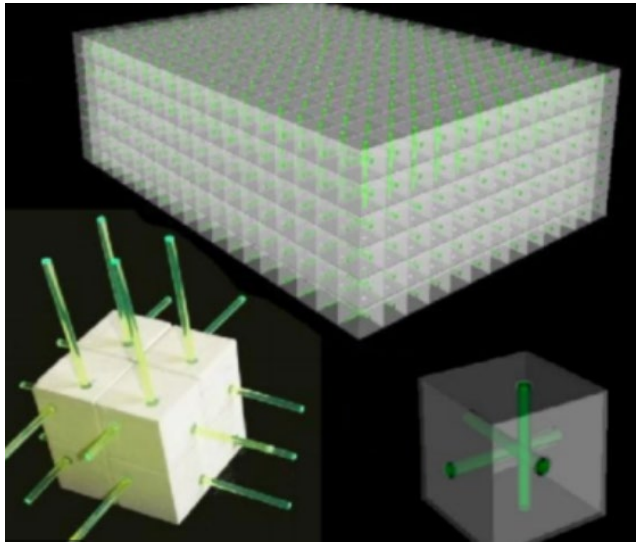

The geometry space has become larger and more complex to explore.

New performance demands are also arising:

- HL-LHC studies made clear that tracking in dense environments requires AI solutions
- In space, payload and power consumption remain major constraints, making applications especially sensitive to hard-to-make design decisions
- Boosted jet tagging at high p_T –all the rage for NP searches at the LHC –demands us to invest in better calorimeters

(Background: boosted decays and fat jets)

Energetic LHC collisions may produce heavy objects with large momentum (top quarks, or W, Z, H bosons). When these decay, they yield a collimated stream of particles – a single (fat) hadron jet.

A number of techniques allow the extraction of features sensitive to the heavy object decay

High granularity and effective identification of individual constituents within dense environments has become unavoidable

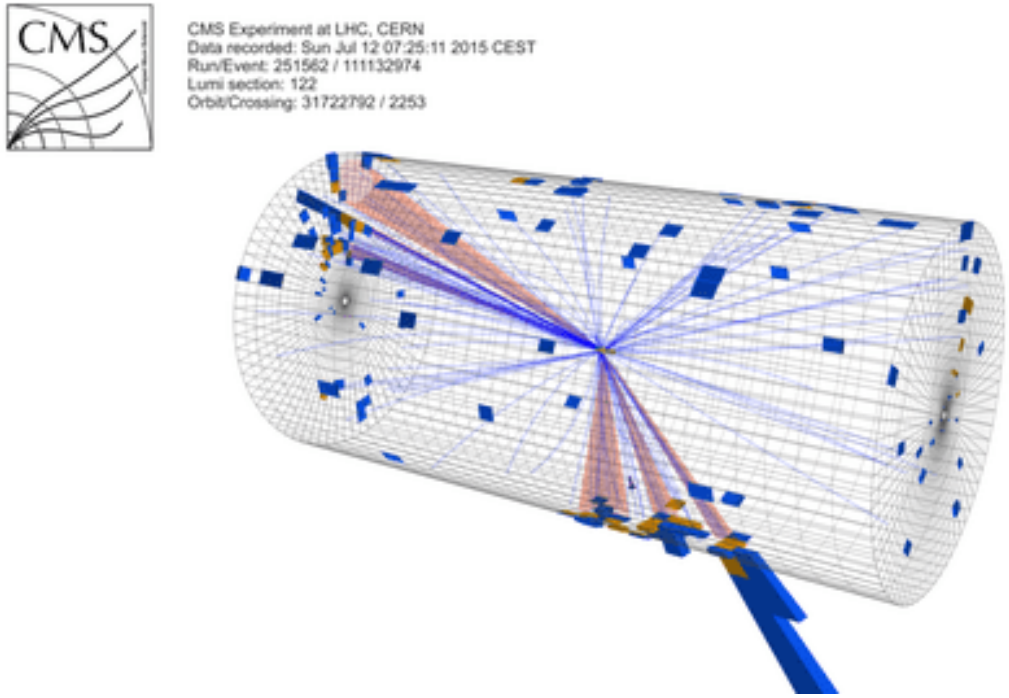

Above: a top-pair decay produces two fat jets, where the individual subjects are visible

Speaking of calorimeters…

High granularity has thus become a compelling requirement.

- The CMS HGCAL detector [6] is a step in that direction; its design will improve by a large margin usable information about the showers, development, pointing, and composition.
- For different reasons, similar developments and improvements are planned for other projects (*e.g.,* CALICE [7] or CaloCUBE [8]).

However, an end-to-end optimization of the design of such instruments has not been attempted yet; **nor have models of the future potential of machine learning in pattern recognition been considered so far in the design phase**

As a telling example, the HGCAL detection elements are arranged in a hexagonal symmetry which offers construction benefits but significantly complicates the most common imaging techniques employing convolutional neural networks (CNN) to shower reconstruction. While solutions to this issue do exist (*e.g.*, see [9]), this is an example of misalignment between design and potential exploitation.

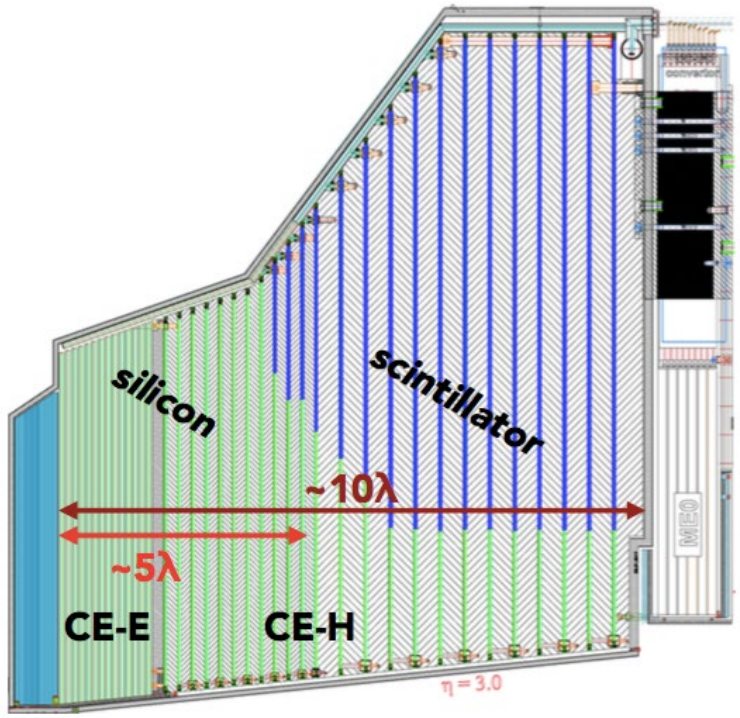

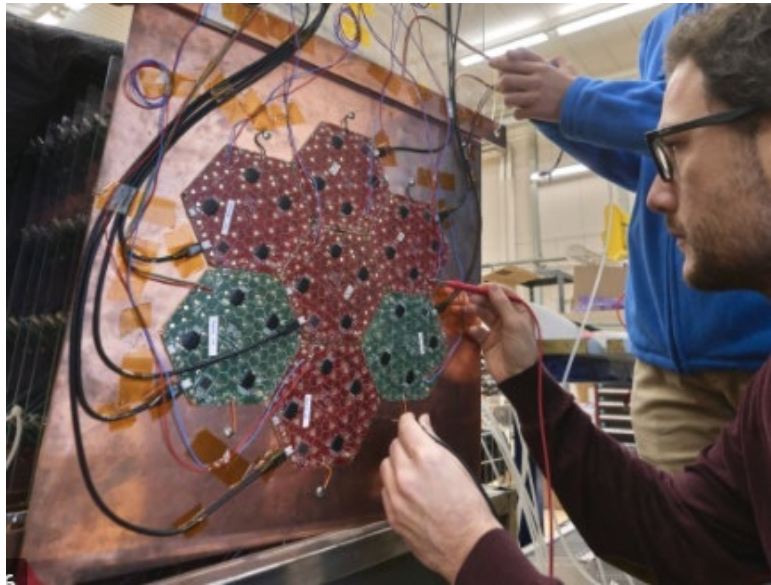

One further note on calorimetry

Take the LHC experiments for a telling example. CMS was originally endowed with a less performant hadron calorimeter than ATLAS. Only later did hadron calorimetry get recognized as crucial for new physics searches involving boosted jets.

CMS regained the lost ground through the high performance of its "particle flow" reconstruction algorithm [10].

This was only possible thanks to the high magnetic field integral of CMS, which spreads out charged particles of different momenta within jets, easing their matching to calorimeter deposits.

This *post-hoc* exploitation of the solenoid characteristics, whose original specifications were rather driven by compactness and transverse momentum resolution of charged particles, is a **striking example of how the combined search of hardware and software solutions may be proficuous** to inform the optimization of a modern particle detector.

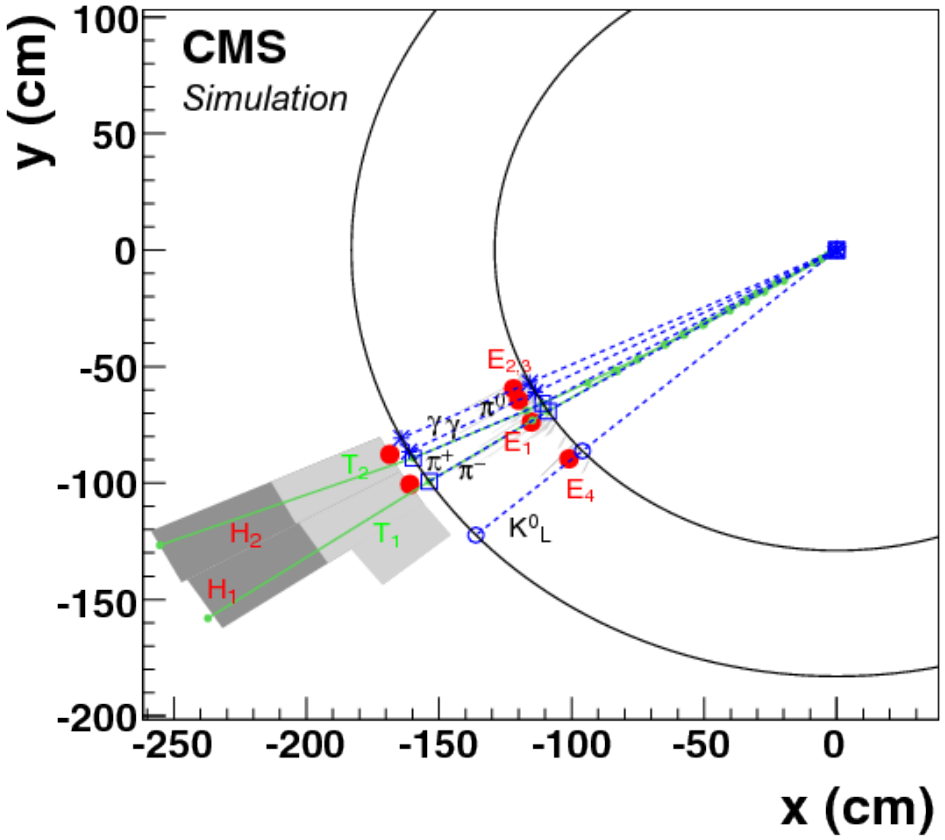

A hybrid calorimeter?

As particle flow techniques allow the tracing of individual particles and the complete reconstruction of dense, collimated jets, we must have more of that:

- Optimizing the design of a detector for a long-timescale project based on reconstruction capabilities which will be available in the future [11] must be pursued
- Integrating tracking and calorimetry layers may improve the «image» reconstruction of energetic hadronic jets, shown to be crucial for high-mass new particles
- Measuring muons from their radiative loss in a dense environment using convolutional neural networks was first shown to be viable (J. Kieseler, CERN) by «accident», as a tangential observation of NN outputs from shower reconstruction in the phase-2 CMS endcap calorimeter.
  → This will become imperative in future colliders at the energy frontier (see *infra*)
- Nuclear interactions have always been dreaded in a tracker, but in combination with calorimetry they may strengthen particle-ID (using probabilistic information coming from nuclear cross sections of different species)
  → of special interest to a number of applications

Muon energy measurement in a calorimeter?

Muons interact with matter by ionization, pair production, bremsstrahlung, and photonuclear reactions. The E loss is dominated by the high-end of the Landau distribution (knock-on electrons).

The total release is very modest and stochastic, so we have to rely on magnetic bending for inference on muon momentum

Bending measurements break down for TeV energies: in 2T, a 1 TeV muon traversing 2m of field is deflected by less than a mm
→ resolution scales: e.g. in ATLAS, **σ(p)/p = 0.2 p**

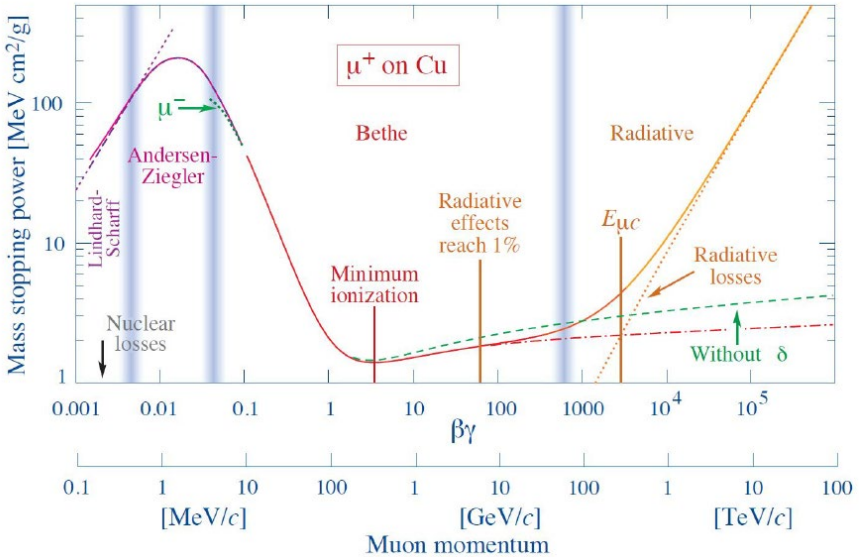

Left: mass stopping power for positive muons in Cu,showing the radiative energy loss onset above 1 TeV

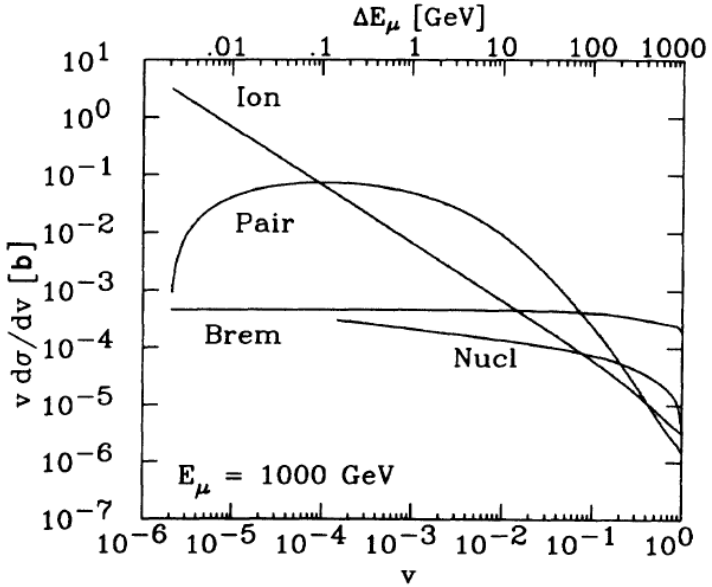

FIG. 3. The ionization, bremsstrahlung, pair production and photnuclear cross sections of a 1000-GeV muon incident on an iron atom. The top scale gives the muon energy loss ΔE_μ, which corresponds to the fractional muon energy loss v. Note that the ordinate is the logarithmic derivative $d\sigma/d(\ln v) = d\sigma/d(\ln \Delta E_\mu)$.

(From a CCFR study [xx])

Energy regression with CNNs

We studied how a granular calorimeter may determine the energy of multi-TeV muons using a customized deep learning architecture, which combines convolutional blocks and dense layers using both high-level features and raw «image-like» energy deposits in 3D space

Of relevance is the point that the pattern of radiation deposits
contains information useful to regress to true muon energy

Results sh[illegible]w that [illegible] ecov[illegible]r 20-25% resolution for muons of u[illegible] [illegible]ing the radiation loss infor[illegible] [illegible]ation (we assumed a relative [illegible] 20% from magnetic bend[illegible] [illegible]y ATLAS in mid-rapidity region[illegible]

See Lukas Layer's talk Aug 25!

arXiv:2107.02119v1 [physics.ins-det] 5 Jul 2021

Calorimetric Measurement of Multi-TeV Muons via Deep Regression

Jan Kieseler[1], Giles C. Strong[2,3], Filippo Chiandotto[2], Tommaso Dorigo[3], and Lukas Layer[4,3]

[1]CERN
[2]Università di Padova
[3]INFN, Sezione di Padova
[4]Università di Napoli "Federico II"

June 30, 2021

Abstract

The performance demands of future particle-physics experiments investigating the high-energy frontier pose a number of new challenges, forcing us to find improved solutions for the detection, identification, and measurement of final-state particles in subnuclear collisions. One such challenge is the precise measurement of muon momentum at very high energy, where an estimate of the curvature provided by conceivable magnetic fields in realistic detectors proves insufficient for achieving good momentum resolution when detecting, *e.g.*, a narrow, high mass resonance decaying to a muon pair.

In this work we show the feasibility of an entirely new avenue for the measurement of the energy of muons based on their radiative losses in a dense, finely segmented calorimeter. This is made possible by exploiting spatial information of the clusters of energy from radiated photons in a regression task. The use of a task-specific deep learning architecture based on convolutional layers allows us to treat the problem as one akin to image reconstruction, where images are constituted by the pattern of energy released in successive layers of the calorimeter. A measurement of muon energy with better than 20 % relative resolution is shown to be achievable for ultra-TeV muons.

How large are the gains of a full optimization?

I recently provided a clear example [13] of how experimental design as is carried out today leaves ample room for improvement from the systematic study of even seemingly irrelevant choices for, *e.g.*, the placement of active and passive material in a simple detector.

The chance of doing so was offered by my refereeing work of the detector proposed by the MUonE collaboration [14], which aims at determining with high precision the cross section of elastic muon-electron scattering.

In the cited study I demonstrated, through the direct exploration of the parameter space of detector geometry, how **large gains** in suitable utility functions (related to the resolution in the event q^2) can be obtained by moving away from choices dictated by past experience

One example of geometry optimization: MUonE

MUonE [14] aims to determine with high precision the muon-electron elastic scattering differential cross section, to extract hadronic contributions and reduce the systematics of the g-2 muon anomaly

The experiment must be sensitive to hadronic loop effects particularly at high q^2, where a 10^{-4} measurement may substantially improve the theoretical understanding of the g-2 value

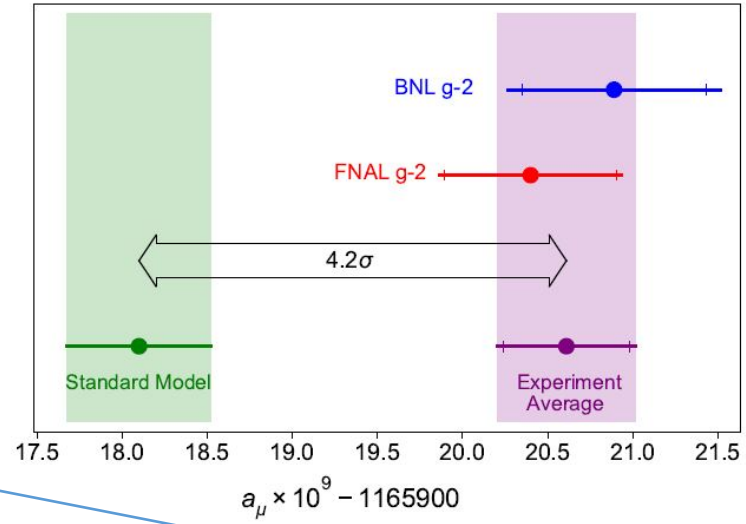

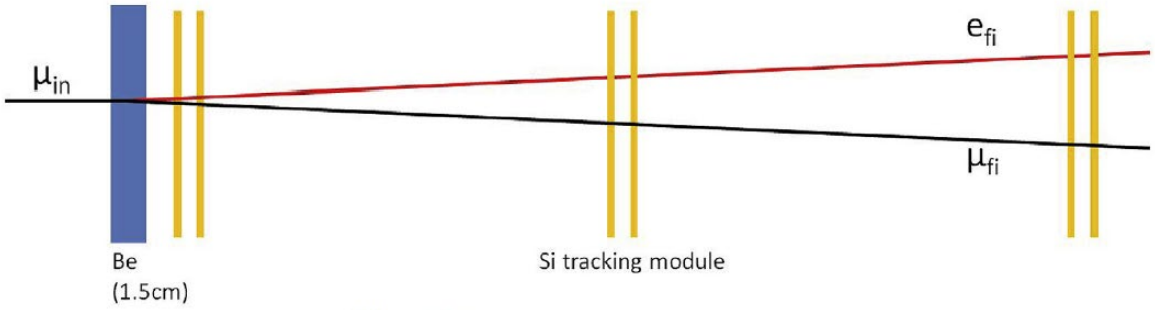

Above: layout of one of 40 1m-long stations

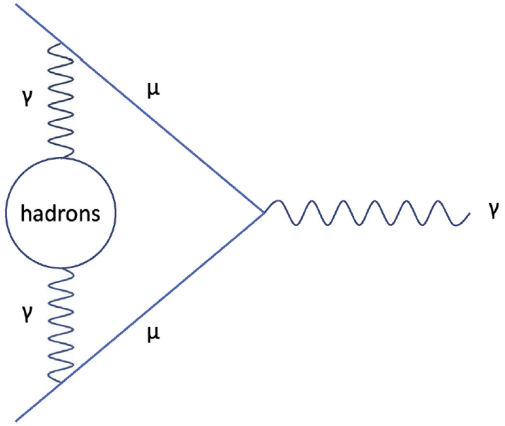

Virtual hadronic loop

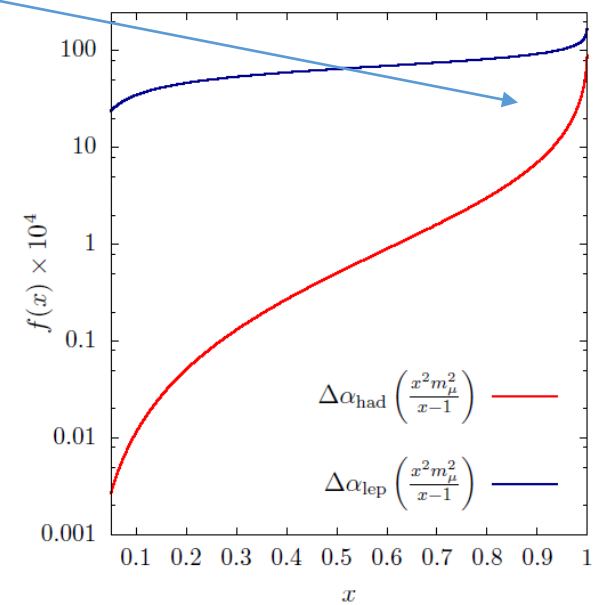

MUonE optimization

By optimizing layout with a discrete sampling, I proved how an improvement of **a factor of 2** in the relevant metric could be achieved without increase in detector cost

The study also proved how dreaded systematic effects from positioning uncertainties could be nullified by software means

ELSEVIER

Physics Open
Volume 4, September 2020, 100022

Geometry optimization of a muon-electron scattering detector

Tommaso Dorigo [1]

Show more

https://doi.org/10.1016/j.physo.2020.100022

Abstract

A high-statistics determination of the differential cross section of elastic muon-electron scattering as a function of the transferred four-momentum squared, $d\sigma_{el}(\mu e \to \mu e)/dq^2$, has been argued to provide an effective constraint to the hadronic contribution to the running of the fine-structure constant, $\Delta\alpha_{had}$, a crucial input for precise theoretical predictions of the anomalous magnetic moment of the muon. An experiment called "MUonE" is being planned at the north area of CERN for that purpose. We consider the geometry of the detector proposed by the MUonE collaboration and offer a few suggestions on the layout of the passive target material and on the placement of silicon strip sensors, based on a fast simulation of elastic muon-electron scattering events and the investigation of a number of possible solutions for the detector geometry. The employed methodology for detector

Sample results

The study was *not* performed with deep learning technologies, as that was not strictly necessary given the reduced space of design choices I wished to investigate.

By optimizing layout with a discrete sampling, I proved how an improvement of **a factor of 2** in the relevant metric could be achieved without increase in detector cost

The results prove that design optimization is not something alien to our reach, but rather, something we should pay more attention to!

We can only guess how large are the gains in the final experimental objectives possible if a fully differentiable model is created for detectors of significantly higher complexity than MUonE.

My guess: huge.

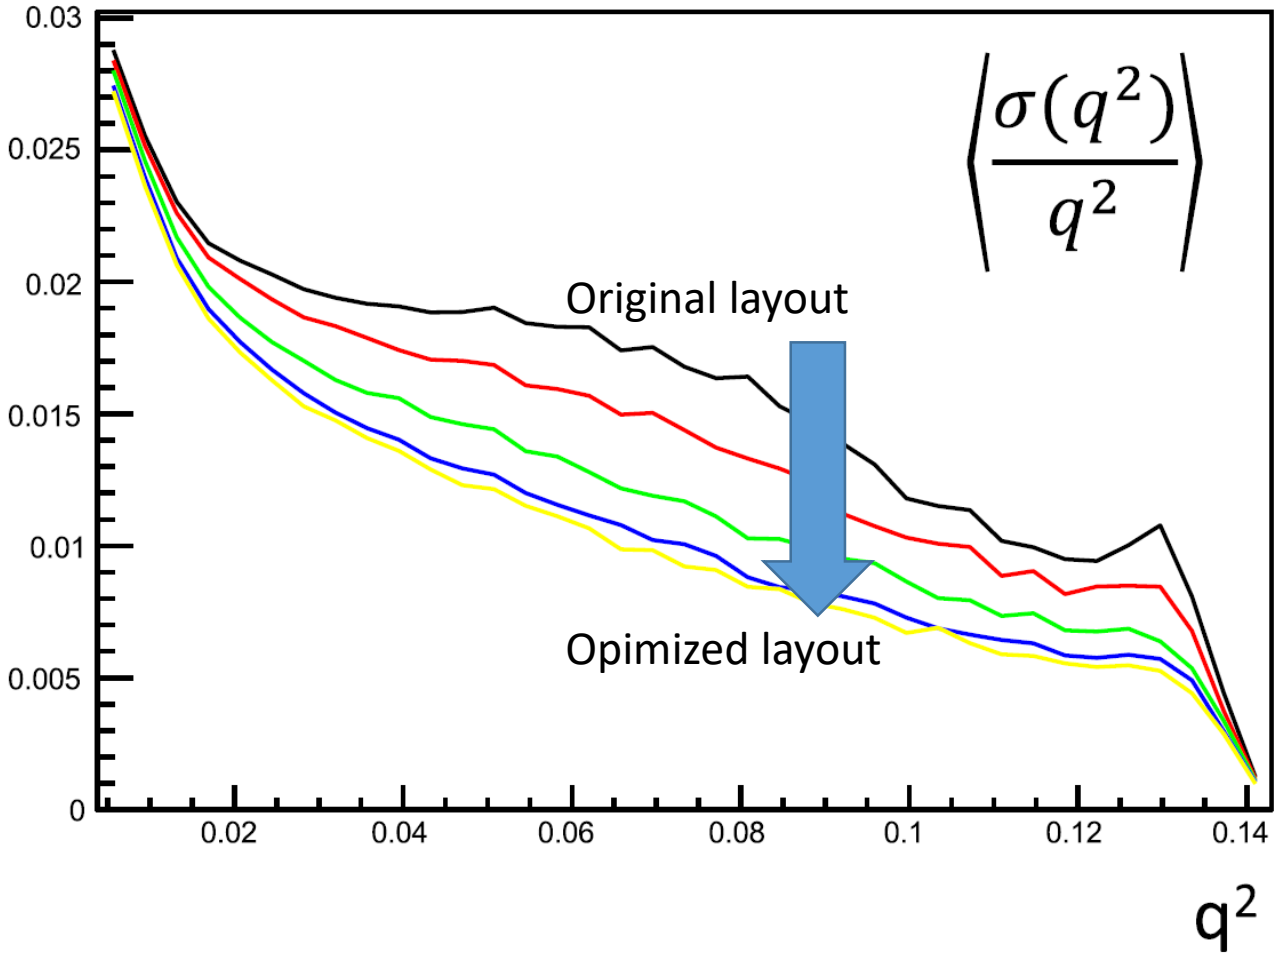

Above: relative resolution in event q^2 for different configurations (the higher, black line is the original proposal by the MUonE coll.)

Speaking of systematic uncertainties,

MUonE correctly identified the need for locating the scattering vertex to within 10μm along the beam axis (it has a strong impact on the q^2 resolution), and proceeded to design a very fancy holographic laser system, to be mounted on each station (=40 systems) to monitor the sensors locations

Cost: several hundred kEuro

My optimization study showed that with 5' of muon beam data, the location, tilt and bow of all detector and target elements can be determined with O(1μ) accuracy by a global fit to the vertex

This is an example of the dividends that the study of a full model of (physics)+(detector)+(reconstruction method)+(inference extraction) can provide

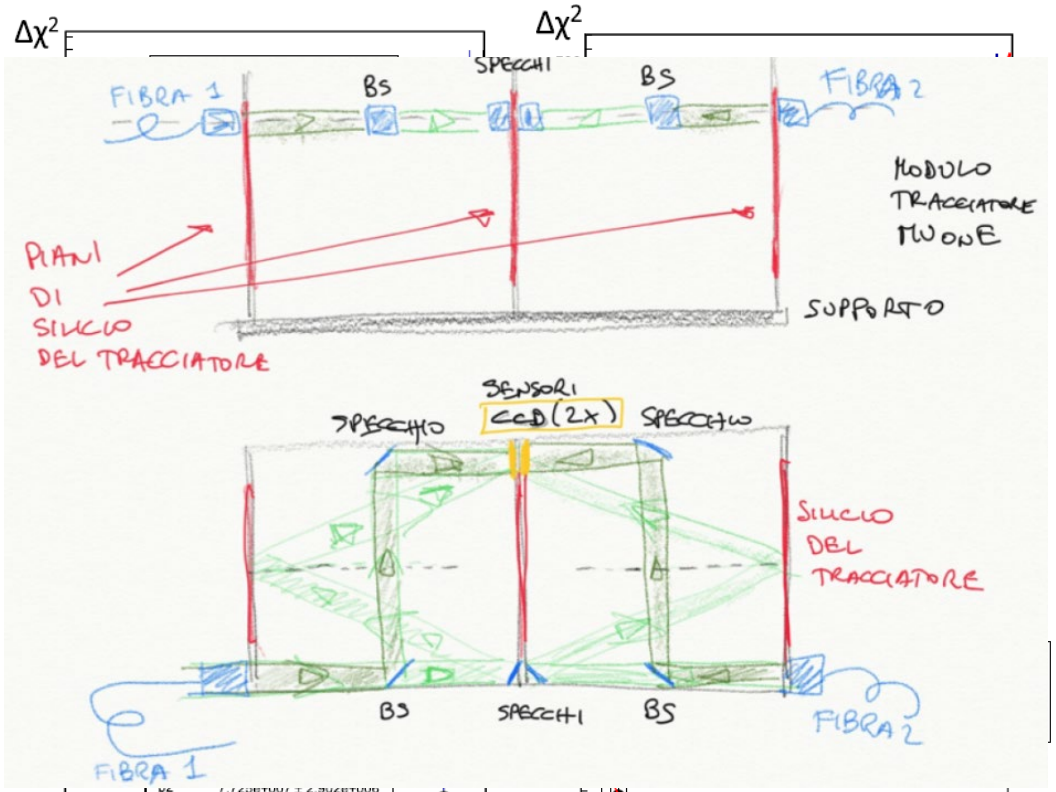

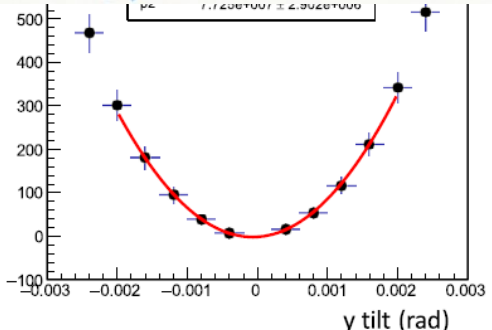

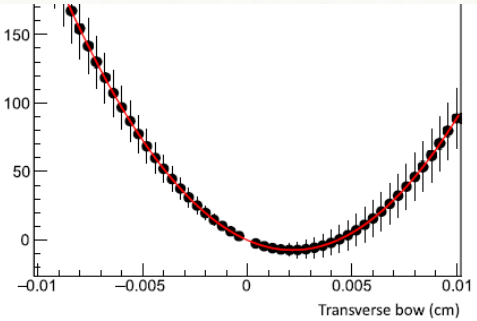

Computer science to the rescue

Progress in CS redefined performance standards of our technologies, and reshaped the way we think about optimization, by providing us with deep learning algorithms that revolutionize common tasks and surpass human performance. We can today identify AI ingredients in, *e.g.*, language translation, speech recognition, self-driving vehicles.

→ Of course, that AI is not general but application-specific: its potential of providing new solutions to old tasks **depends on our ability to create the right interfaces**.

In HEP, ML applications caught up rather slowly, but NNs and gradient boosting techniques eventually operated a paradigm shift, improving the performance of our measurements by large amounts.

A new paradigm shift is now offered by differentiable programming [15], which eases the systematic search of minima of arbitrarily complex multi-dimensional functions; by casting the whole problem in a differentiable framework a full end-to-end optimization becomes possible.

INFERNO

To show what differentiable programming can do for our data analyses, I designed with P. de Castro an innovative algorithm [16] using automatic differentiation to construct a loss function that directly targets the information content of the statistical summary produced by the neural network.

If the loss function incorporates the effects of nuisance parameters on the measurement objective, virtual optimality of the classification task and large improvements in precision can be achieved over procedures that only account for nuisance parameters downstream of the NN training.

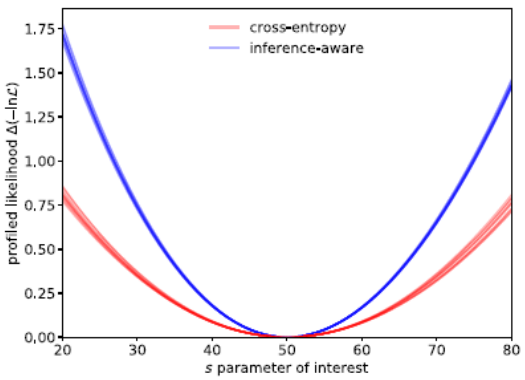

(b) profile-likelihood comparison

Left: profile likelihood on the parameter of interest for a neural network with (blue) and without (red) the feedback on effect of nuisances provided by INFERNO

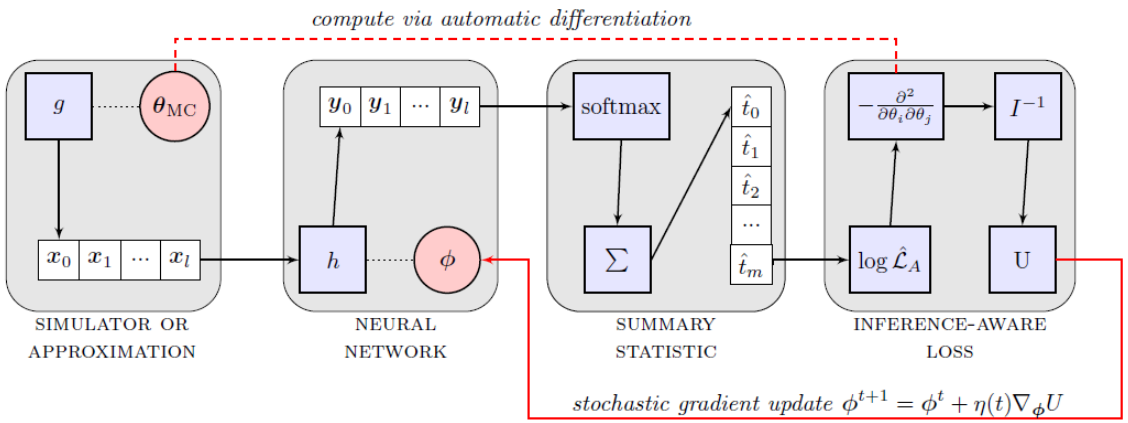

Top: Control flow of INFERNO algorithm, which extracts optimal NN parameters given a final analysis objective, such that the resulting measurement becomes maximally robust to nuisance parameters. Data $\boldsymbol{x}$ *generated by a simulator* $\boldsymbol{g}$ *(left block) which depend on nuisances* $\boldsymbol{\theta_{MC}}$ *are used to train a NN* $\boldsymbol{h}$ *(second block), producing output* $\boldsymbol{y}$*. The output is made differentiable by a softmax function (third block) and used to construct a summary statistic* $\boldsymbol{t}$*, which is histogrammed to compute a saturated likelihood* $\boldsymbol{L_A}$ *used to perform inference (right block). The expected variance* $\boldsymbol{U}$ *on the parameter of interest may be derived from the information matrix* $\boldsymbol{I}$*, and used as NN loss, whose parameters* $\boldsymbol{\varphi}$ *are optimized by backpropagation.*

A study of muon shielding in SHIP

In another seminal work [17], local generative surrogates of the gradient of the objective function were proven to allow for the minimization by SGD and a strong reduction in muon background fluxes in the SHIP experiment

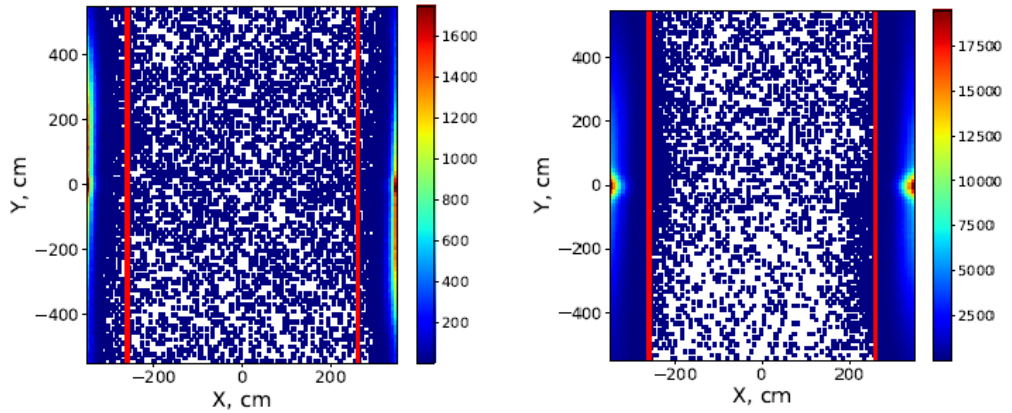

Figure 7. Muon hits distribution in the detection apparatus (depicted as red contour) obtained by Bayesian optimization (Left) and by L-GSO (Right), showing better distribution. Color represents number of the hits in a bin.

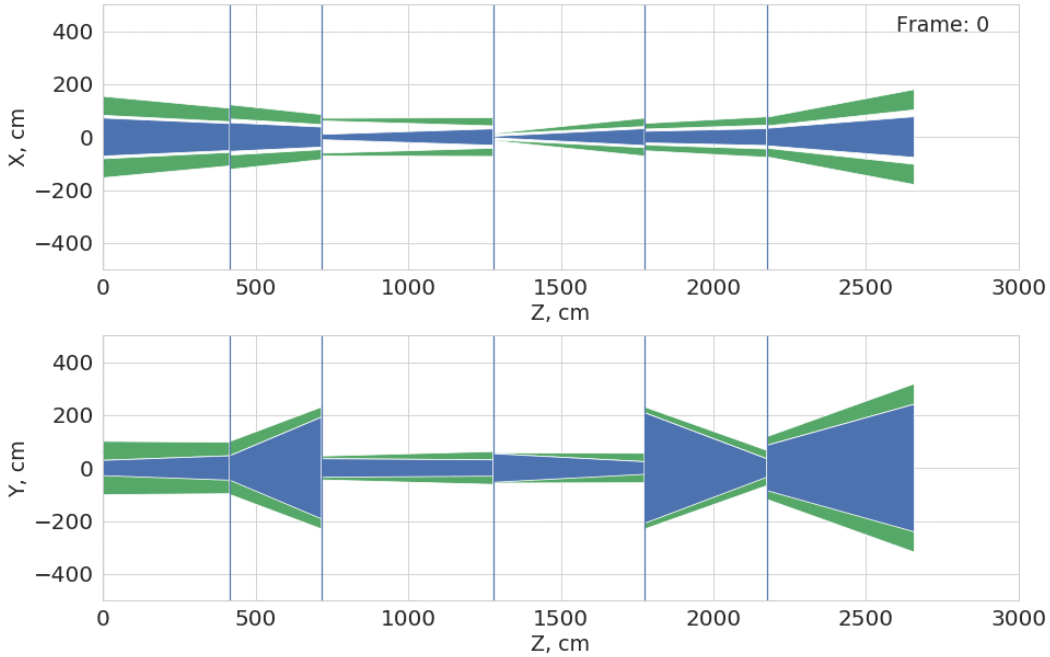

Geometry optimization at work in real time!

Realigning design choices and ultimate goal: MODE

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

Study cases:

- Demonstration of muon energy measurement in optimized calorimeter → article in preparation
- Muon tomography detector optimization [18] → in progress
- Hybrid calorimeter design integrating tracking layers → activity starting

Other use cases being considered include:

- Optimization of Cherenkovs for SWGO
- Hadron therapy (iMPACT project [19]);
- Muon collider detector shielding [20];
- Optimization of MUonE calorimeter;
- Optimized search for long-lived signatures at FCC-ee

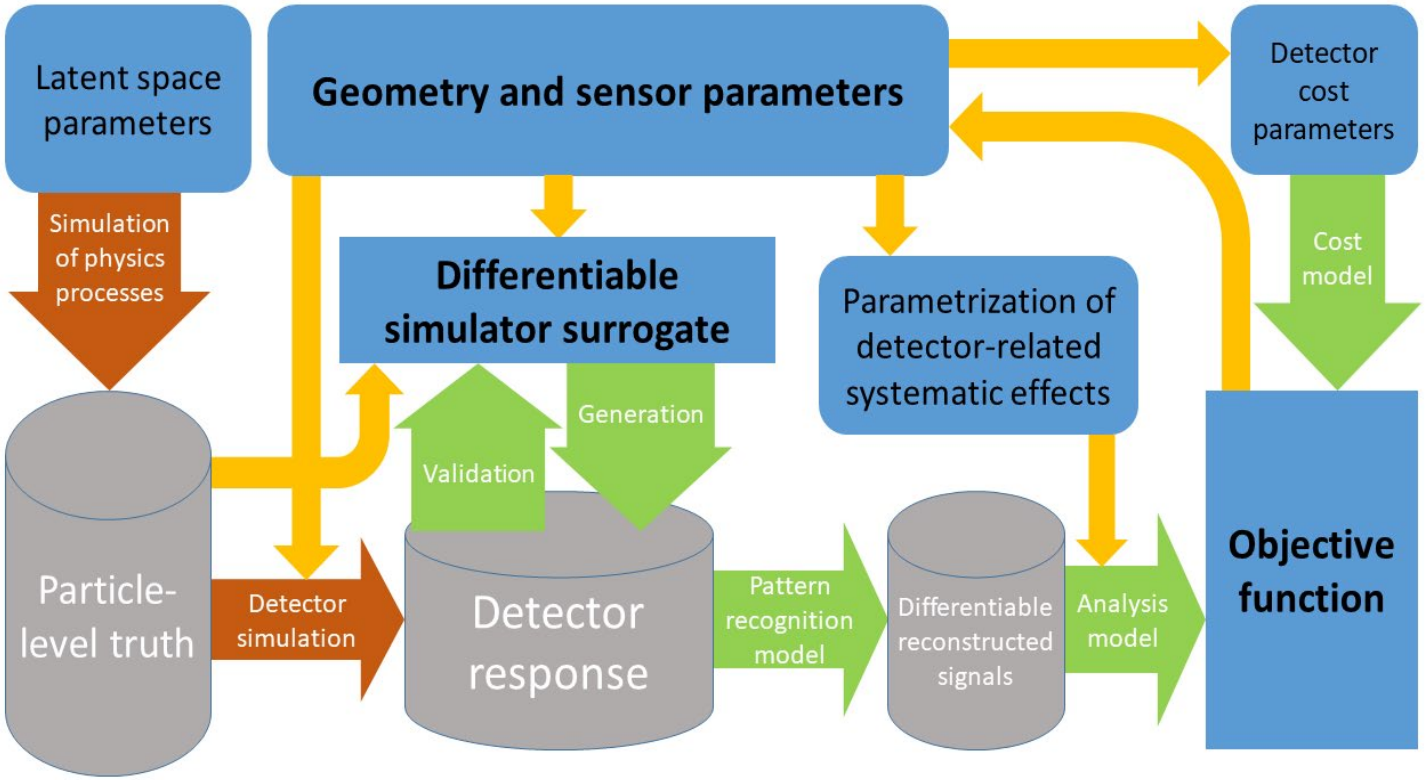

And for the time being…

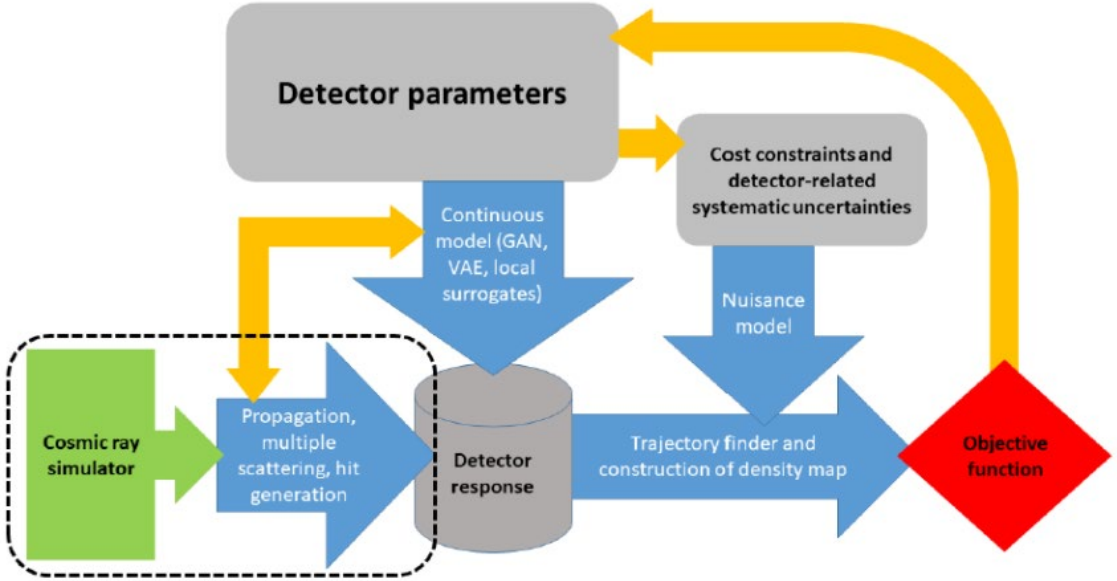

«Simple» use case: muon tomography. We need no surrogate of a simulator, yet all other pieces of the puzzle still need to be carved and set in.

For a simple test, we model a scanned volume including a Pb block of 0.5x0.1x0.1 m^3 inside a 0.6x1x1m^3 of lower-Z material

The system «learns» how to compromise cost and precision, and where detector elements are less useful

A number of shortcuts have been taken to develop this purposedly crude model – but once we have something that «breathes», we may start building into it functionality and detail

Right: scheme of the modeled apparatus

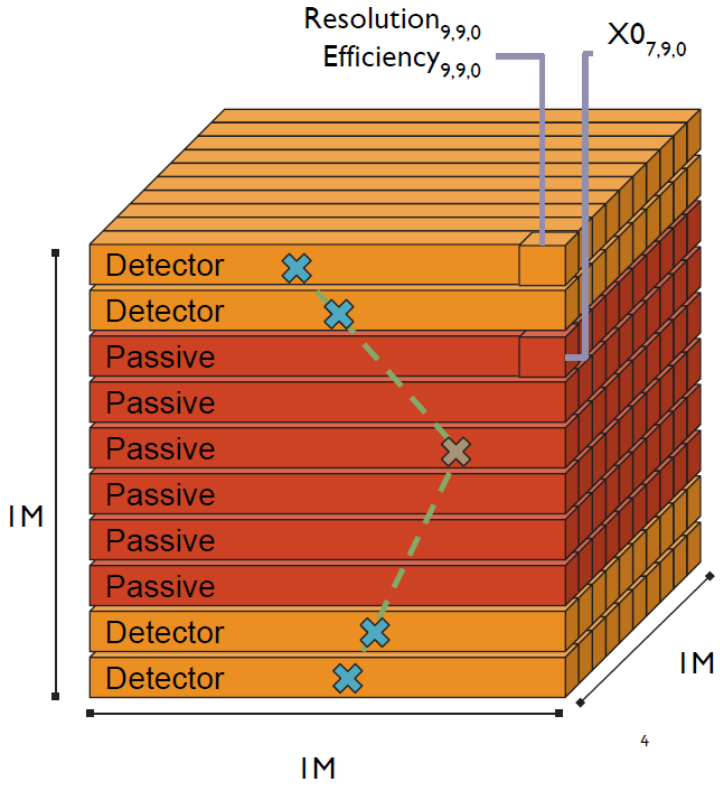

VERY preliminary results

The code and results shown have been produced by **Giles Strong**

The graphs show the result of a run of 100 epochs training, followed by a prediction with 100k muons

First proof of principle (very low statistics) of correct training of a differentiable model of a schematic muon tomography apparatus.

The **loss** is a combination of detector cost (itself a function of sensors efficiency and resolution) and RMSE on rad length estimate

Left graph, top to bottom: loss, loss composition, resolution map, and efficiency map of detection elements after minimization

Still a looong way to go, but an important milestone for this use case

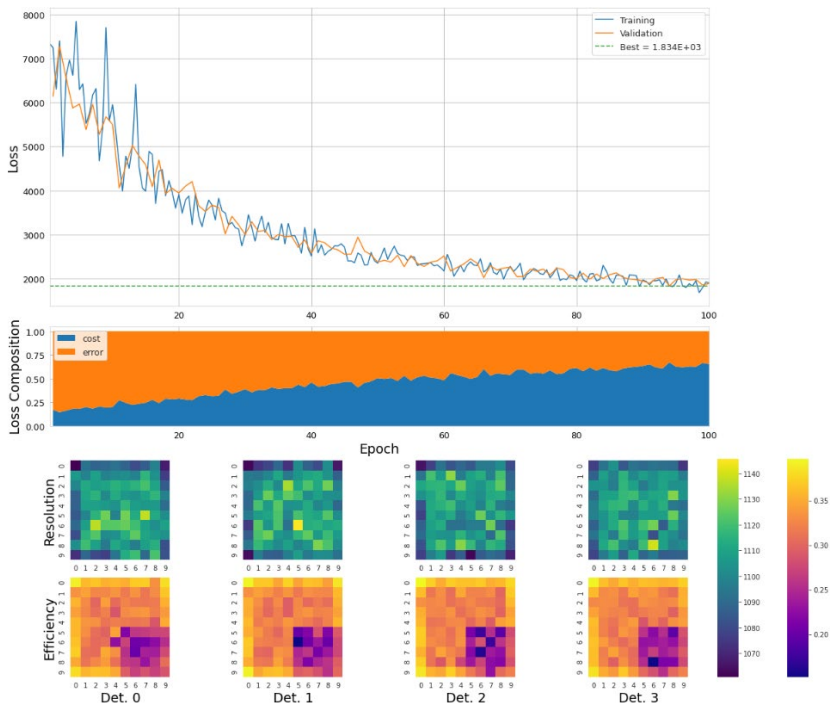

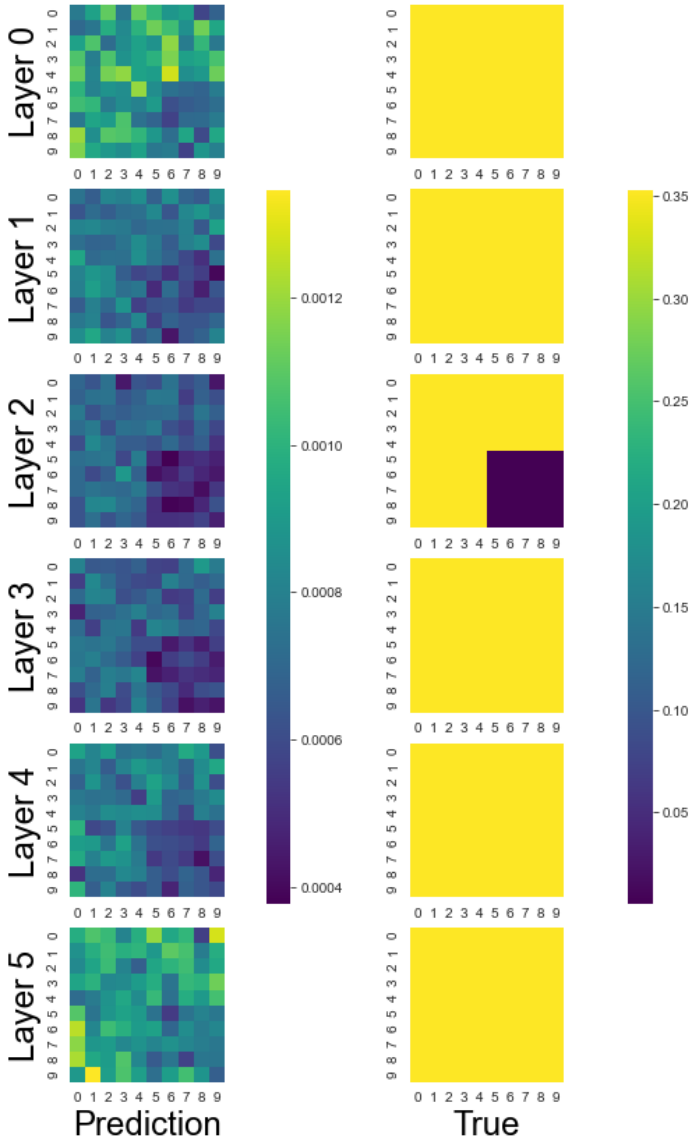

Right: predicted and true X_0 of passive volume

Realigning our design choices to future AI

A point which cannot be stressed enough is that if we design today something that will operate 10 or 20 years in the future, we need to account for the pattern recognition capabilities of future automated systems

In 20 years, will we use a Kalman filter to reconstruct trajectories in our trackers, or photon energy and direction in our calorimeters?

No, we won't. We will employ AI technology, streamlined by a decade of consolidation in similar tasks.

Shouldn't we then build those devices by considering how AI technology could best exploit them? If we do not, we will suffer a misalignment of our design choices and the future capabilities of the software we will end up using.

How to get around this problem?

We can and should try to model increasingly performant pattern recognition in our optimization loops, and verify whether there are discontinuities in the solutions space.

It is not going to be easy, but it is IMHO absolutely necessary to start getting equipped.

Four-slide description of a differentiable model for design optimization - 1

An end-to-end detector design optimization task can be briefly formalized in the following way.

We start with a simulation of the physics processes of relevance for the considered application, which generates a multi-dimensional, stochastic input variable **x**, distributed with a PDF **f(x)**.

The input is turned by the simulation of the detection apparatus into sensor readouts **z** distributed with a PDF

$$\mathbf{p(z|x,\theta)},$$

which constitute the observed low-level features of the physical process; readouts **z** depend through **p( )** on parameters **θ** that describe the physical properties of the detector and its geometry.

Four-slide description of a differentiable model for design optimization - 2

The observations **z** are used by a reconstruction model **R( )** that produces high-level features

$$\zeta(\theta) = R[z,\theta,\nu(\theta)]$$

(*e.g.* particle four-momenta), by employing knowledge of the detector parameters as well as a model of the detector-driven nuisance parameters **ν(θ)** which affect the pattern recognition task.

In turn, high-level features **ζ(θ)** constitute the input of a further, less dramatic, dimensionality reduction, the data analysis step: this is typically performed by a classifier or regressor **NN( )** powered by a neural network.

Once properly trained for the task at hand, the network produces a low-dimensional summary statistic

$$s = NN[\zeta(\theta)]$$

with which inference can finally be carried out to produce the desired goal of the experiment.

Four-slide description of a differentiable model for design optimization - 3

In general, one may formally specify the problem of identifying optimal detector parameters as that of finding estimators $\widehat{\theta}$ that satisfy

$$\widehat{\theta} = arg \min_{\theta} \int L[NN(\zeta), c(\theta)] p(z|x, \theta) f(x) dx dz$$

c(θ) is a function modeling the cost of the considered detector layout of parameters **θ**, and the loss function **L[NN,c]** is constructed to appropriately weight the result of the measurement in terms of its desirable goals, as well as to obey cost constraints and other use-case-specific limitations.

Since in the cases of interest the PDF **p(z|x,θ)** is not available in closed form –the considered models are implicit–, we must rely on forward simulation: we approximate $\widehat{\theta}$ with a sample of n events:

$$\widehat{\theta}_a = arg \min_{\theta} \frac{1}{n} \sum_{i=1}^{n} L\left[NN\left(R(z_i)\right), c(\theta)\right]$$

where z_i is distributed as $F(x_i,\theta)$ to emulate **p(z|x,θ)** as x_i is sampled from its PDF **f( )** by the simulator. One may thus obtain an estimate of the loss function and the detector parameters which minimize it.

Four-slide description of a differentiable model for design optimization - 4

It has been shown how in applications such as those of our interest it is viable to approximate the non-differentiable stochastic simulator **F( )** with a local surrogate model,

$$z = S(y,x,\theta),$$

that depends on a parameter **y** describing the stochastic variation of the approximated distribution[49]. This allows to descend to the minimum of the approximated loss $\widehat{L(z)}$ by following its surrogate gradient

$$\nabla_\theta \widehat{L(z)} = \frac{1}{n}\sum_{i=1}^{n} \nabla_\theta L[NN\left(R\left(S(y_i, x_i, \theta)\right)\right), c(\theta)].$$

The above recipe requires to learn the differentiable surrogate **S( )**: this is a task liable to be carried out independently from the optimization procedure.

The modular structure of a differentiable pipeline modeling the optimization cycle allows the user to turn on and off specific parts of the chain, helping the system in its exploration of the feature space.

Machine-Learning Optimized Design of Experiments

MODE Collaboration

https://mode-collaboration.github.io

A. G. Baydin[5], A. Boldyrev[4], K. Cranmer[8], P. de Castro Manzano[1], T. Dorigo[1], C. Delaere[2], D. Derkach[4], J. Donini[3], A. Giammanco[2], J. Kieseler[7], G. Louppe[6], L. Layer[1], P. Martinez Ruiz del Arbol[9], F. Ratnikov[4], G. Strong[1], M. Tosi[1], A. Ustyuzhanin[4], P. Vischia[2], H. Yarar[1]

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

The strategy of MODE

We are fully aware that the dream of informing the design of a complex detector for a fundamental physics endeavour (be it HE, astro-, Neutrino, or Nuclear physics) entails a walk in the desert

Yet we must start it, as in 20 years the shortcomings of having designed experiments that are misaligned with goals and information-extraction procedures will otherwise be paid dearly.

The strategy is thus to start with easy use cases, where further proof may be brought of the gains of using DL architectures to parametrize the essential ingredients of the design problems

Hopefully, we will be able to convince the community, or else, we'll have to wait for a generation change.

The important observation is that the developed architectures for optimization are modular, hence we will be able to recycle part of the work for one application when we move to the next one.

Status and next steps

- An article describing the MODE program has been published last month in Nuclear Physics News International [21]
- A white paper on differentiable programming for detector design optimization is being drafted
- We are organizing a workshop on "*Differentiable Programming for Design Optimization*" on September 6-8 2021 in Louvain-la-Neuve, to allow interested scientists to join and discuss together the means and the possible applications
  - Extra support for this activity is provided by IRIS-HEP and JENAA

Every other solved application = a publication AND added knowledge base on solving these hard problems!

You are most welcome to participate to MODE activities, or propose to become a member

THANK YOU FOR YOUR ATTENTION!

References

1. The European Strategy Group, "*2020 Update of the European Strategy for Particle Physics*," CERN, Geneva, https://europeanstrategy.cern/, doi: http://dx.doi.org/10.17181/ESU2020Deliberation, 2020.

2. The topic has been discussed widely in blogs and other social media. *E.g.*, see Sabine Hossenfelder's discussion here: http://backreaction.blogspot.com/2019/06/if-we-spend-money-on-larger-particle.html, or see Alessandro Strumia's guest post at the same site: http://backreaction.blogspot.com/2020/06/guest-post-who-needs-giant-new-collider.html. Also see the New York Times op-ed by S. Hossenfelder, https://www.nytimes.com/2019/01/23/opinion/particle-physics-large-hadron-collider.html

References / cont'd

3. F. Hartmann, *Evolution of Silicon Sensor Technology in Particle Physics*, Springer Tracts in Modern Physics, 2017, ISBN 978-3-319-64436-3; P. Delpierre, *A history of hybrid pixel detectors, from high energy physics to medical imaging*, J. INST. 9 (2014) C05059, doi: 10.1088/1748-0221/9/05/C05059; P. Allport, *Applications of silicon strip and pixel-based particle tracking detectors*, Nature Reviews Physics 1 (2019) 567–576, doi: https://doi.org/10.1038/s42254-019-0081-z; G. Case, *New trends in silicon detector technology*, JINST 15 (2020), https://iopscience.iop.org/article/10.1088/1748-0221/15/05/C05057; M. Garcia-Sciveres and N. Wermes, *A review of advances in pixel detectors for experiments with high rate and radiation*, Rep. Prog. Phys. 81 (2018) 066101, doi: http://doi.org/10.1088/1361-6633/aab064.

References / cont'd

4. Y. Mishnayot *et al*., "*3D Printing of Scintillating Materials*", arXiv:1406.4817[cond-mat.mtrl-sci] (2014); Y. Abreu *et al*. (SoLID Collaboration), "*A novel segmented-scintillator antineutrino detector*", JINST 12 (2017) P04024, doi: 10.1088/1748-0221/12/04/P04024.

5. M. Mandurrino *et al*., "*Demonstration of 200 --, 100 --, and 50 micron Pitch Resistive AC Coupled Silicon Detectors (RSD) With 100% Fill Factor for 4D Particle Tracking*," IEEE Electron Device Letters 40, 11 (2019) 1780, arXiv:1907.03314[physics.ins-det]; M. Mandurrino *et al*., "*Analysis and numerical design of Resistive AC Coupled Silicon Detectors (RSD) for 4D particle tracking*", Nucl. Instr. Meth. A959 (2020) 163479, doi: 10.1016/j.nima.2020.163479; M. Tornago et al., "*Resistive AC-Coupled Silicon Detectors: Principles of operation and first results from a combined analysis of beam test and laster data*", NIM A 1003 (2021) 165319, https://www.researchgate.net/publication/343095980 .

6. CMS Collaboration, "*The Phase-2 Upgrade of the CMS Endcap Calorimeter*", CERN-LHCC-2017-023 (2018), https://cds.cern.ch/record/2293646?ln=en .

References / cont'd

7. J. Repond *et al.*, "*Hadronic Energy Resolution of a Combined High Granularity Scintillator Calorimeter System*", JINST 13 (2018) P12022, doi: 10.1088/1748-0221/13/12/P12022; C. Adloff *et al.*, "*Tests of a particle flow algorithm with CALICE test beam data*", JINST 6 (2011) P07005, doi: 10.1088/1748-0221/6/07/P07005.

8. P.W. Cattaneo *et al.*, "*CaloCube: a novel calorimeter for high-energy cosmic rays in space*", JINST 12 (2017) 06, C06004, doi: 10.1088/1748-0221/12/06/C06004 .

9. E. Hoogeboom *et al.*, "*HexaConv*", arXiv:1803.02108[cs.LG] (2018)).

10. M. Sirunyan *et al*. (CMS Collaboration), "*Particle-flow reconstruction and global event description with the CMS detector*", JINST 12 (2017) P10003, doi: 10.1088/1748-0221/12/10/P10003 .

References / cont'd

11. See *e.g.* S.R. Qasim, **J. Kieseler**, Y. Iiyama, and M. Pierini, "*Learning representations of irregular particle-detector geometry with distance-weighted graph networks*", Eur. Phys. J. C79 (2019) 7,608, doi: 10.1140/epjc/s10052-019-7113-9; **J. Kieseler**, "*Object condensation: one-stage grid-free multi-object reconstruction in physics detectors, graph and image data*", arXiv:2002.03605[physics.data-an] (2020); J. Alimena, Y. Iiymana, **J. Kieseler**, "*Fast convolutional neural networks for identifying long-lived particles in a high-granularity calorimeter*", arXiv:2004.10744 [hep-ex] (2020).

12. **T. Dorigo, J. Kieseler,** L. Layer and G. Strong, "*Muon Energy Measurement from Radiative Losses in a Calorimeter for a Collider Detector*", http://arxiv.org/abs/2008.10958 (2020).

13. **T. Dorigo**, "*Geometry Optimization of a Muon-Electron Scattering Detector*," Physics Open 4 (2020) 100022, arXiv:200200973[physics.ins-det], doi: 10.1016/j.physo.2020.100022.

14. G. Abbiendi *et al.*, "*Letter of Intent: The MUonE Project*", CERN-SPSC-2019-026/SPSC-I-252 (2019).

References / cont'd

15. **A. Güneş Baydin**, B.A. Pearlmutter, A.A. Radul, and J.M. Siskind, "*Automatic Differentiation in Machine Learning: a Survey*", Journal of Machine Learning Research (JMLR) 18 (153) (2018) 1, http://jmlr.org/papers/v18/17-468.html; Y. LeCun, Y. Bengio, and G. Hinton, "*Deep learning*", Nature, 521(7553) (2015) 436, doi: 10.1038/nature14539.

16. **P. de Castro Manzano** and **T. Dorigo**, "*INFERNO: Inference-Aware Neural Optimization*", Comp. Phys. Commun. 244 (2019) 170; Arxiv:1806.04743v2 [stat.ml] (2018), doi: 10.1016/j.cpc.2019.06.007.

17. S. Shirobokov, **A. Ustyuzhanin, A. Güneş Badyin** *et al*., "*Differentiating the Black-Box: Optimization with Local Generative Surrogates*", arXiv:2002.04632v1[cs.LG] (2020).

18. S. Wuyckens, **A. Giammanco**, P. Demin, and **E. Cortina Gil**, "*A Portable muon telescope based on small and gas-tight Resistive Plate Chambers*", Phil. Trans. Royal Soc. A377 (2019) 2137, arXiv:1806.06602v2[physics.ins-det] (2018), doi: 10.1098/rsta.2018.0139.

References / cont'd

19. P. Giubilato *et al*., "*iMPACT: innovative pCT scanner*", IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC) IEEE (2015), https://ieeexplore.ieee.org/abstract/document/7581240; S. Mattiazzo, *et al*., "*Advanced proton imaging in computed tomography*", Radiation protection dosimetry 166.1 (2015) 388, https://academic.oup.com/rpd/article/166/1-4/388/1612689; S. Mattiazzo *et al*., "*The iMPACT project tracker and calorimeter*", Nucl. Instr. Meth. A845 (2017) 664, https://www.sciencedirect.com/science/article/pii/S0168900216303412; N. Pozzobon *et al*. "*Calorimeter prototyping for the iMPACT project pCT scanner*", Nucl. Instr. Meth. A936 (2019) 1, https://ui.adsabs.harvard.edu/abs/2019NIMPA.936....1P/abstract.

20. N. Bartosik *et al*., "*Preliminary Report on the Study of Beam-Induced Background Effects at a Muon Collider*," arXiv:1905.03725[hep-ex] (2019); N. Bartosik *et al.*, "*Detector and Physics Performance at a Muon Collider*," J. Inst. 15 (2020) P05001, doi: 10.1088/1748-0221/15/05/P05001.

21. **A.G. Baydin** et al., *"Toward Machine Learning Optimization of Experimental Design"*, submitted to Nuclear Physics News International, 2021.

Muon energy measurement in a calorimeter?

In a preliminary study, we showed that resolutions of 30-35% are achievable for 2 TeV muons in a highly granular, homogeneous calorimeter [12]

- Genetically breeded kNN learners used for this study
- Spatial information proven to be crucial for the task (blue vs red double arrow for 68.3% CI, see below, left)

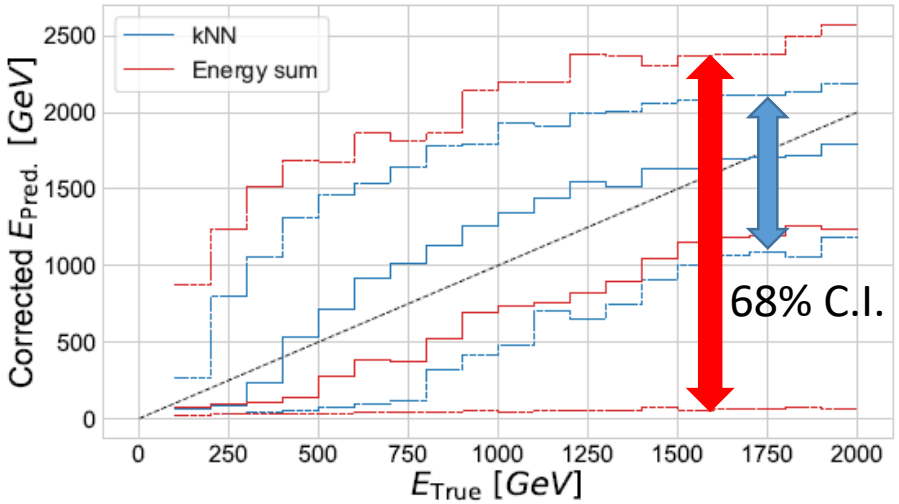

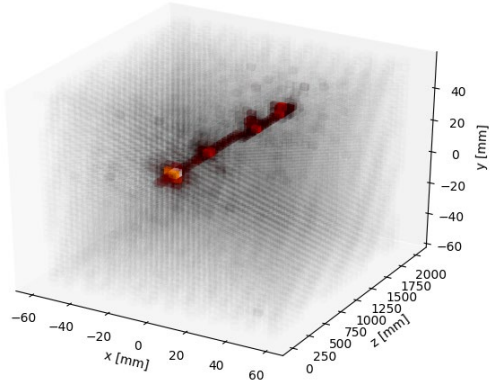

Muon Energy Measurement from Radiative Losses in a Calorimeter for a Collider Detector

Tommaso Dorigo[1], Jan Kieseler[2], Lukas Layer[3], and Giles C. Strong[4]

[1,3,4]INFN, Sezione di Padova
[2]CERN
[3]Università di Napoli "Federico II"
[4]University of Padova

August 25, 2020

Abstract

The performance demands of future particle-physics experiments investigating the high-energy frontier pose a number of new challenges, forcing us to find new solutions for the detection, identification, and measurement of final-state particles in subnuclear collisions. One such challenge is the precise measurement of muon momenta at very high energy, where the curvature provided by conceivable magnetic fields in realistic detectors proves insufficient to achieve the desired resolution.

In this work we show the feasibility of an entirely new avenue for the measurement of the energy of muons based on their radiative losses in a dense, finely segmented calorimeter. This is made possible by the use of the spatial information of the clusters of deposited photon energy in the regression task. Using a homogeneous lead-tungstate calorimeter as a benchmark, we show how energy losses may provide significant complementary information for the estimate of muon energies above 1 TeV.

How good is that, BTW?

Measuring multi-TeV muons has been a Group-2 issue before LHC experiments started to consider it

The resolution of muons traversing 1.5km of ice (=3850 X_0) in IceCUBE has been determined with three different methods in[xx].

Although of course the problem is very different, I have not resisted the temptation to overlay to the graph on the right the ballpark of the resolution we achieve with a 2m-long lead tungstate calorimeter (=225 X_0) + CNN reconstruction

→ 3x better