frontiers | Frontiers in Physics

TYPE Mini Review
PUBLISHED 12 February 2026
DOI 10.3389/fphy.2026.1765091

Check for updates

OPEN ACCESS

EDITED BY
Alex Hansen,
NTNU, Norway

REVIEWED BY
David Cussans,
University of Bristol, United Kingdom

*CORRESPONDENCE
Alessio Figalli,
alessio.figalli@math.eth.ch

RECEIVED 10 December 2025
REVISED 10 January 2026
ACCEPTED 14 January 2026
PUBLISHED 12 February 2026

CITATION
Figalli A, Qasim SR, Owen P and Serra N (2026) Designing particle physics experiments with artificial intelligence.
Front. Phys. 14:1765091.
doi: 10.3389/fphy.2026.1765091

COPYRIGHT
© 2026 Figalli, Qasim, Owen and Serra. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Designing particle physics experiments with artificial intelligence

Alessio Figalli[1]*, Shah Rukh Qasim[2,3], Patrick Owen[2] and Nicola Serra[2,3]

[1]Department of Mathematics, ETH Zürich, Zürich, Switzerland, [2]Department of Physics, Universität Zürich, Zürich, Switzerland, [3]Department of Mathematical Modeling and Machine Learning, Universität Zürich, Zürich, Switzerland

The design of modern particle physics detectors can become a strong benchmark for contemporary machine-learning techniques. It offers a realistic large-scale optimization task grounded in well-understood physics and reliable simulations, providing a controlled setting to test methods aimed at complex real-world problems; the proposed Future Circular Collider is a prominent example of the scale and ambition involved. This review introduces the detector-optimization problem and discusses the growing interest in applying AI methods to detector design, providing a comparative perspective on various methodologies. We show how a specific version of the detector-optimization problem can, and has been, tackled with Bayesian optimization and gradient-based methods, while reinforcement learning addresses a more general formulation that includes sequential and combinatorial structure. The substantial computational burden of Monte Carlo simulation remains a central obstacle, for which we outline how generative machine-learning approaches offer effective mitigation. We also discuss how uncertainty, arising from stochastic detector response, systematic shifts in physics modelling and reconstruction, and long-term operating conditions, can be incorporated into the design process. In particular, we discuss how distributional and distributionally robust reinforcement learning, together with optimal-transport–based ambiguity sets, provides a principled way to capture plausible deviations from nominal assumptions and to search for designs that maintain reliable performance across varied scenarios.

KEYWORDS

differential programming, generative AI, instrument design, machine learning, particle physics, reinforcement learning, robust optimisation

1 Introduction

A core activity in the field of particle physics is the design and construction of large detectors that record the signals of sub-atomic particles. Modern particle physics experiments are among the most complex machines ever built. The Large Hadron Collider (LHC) at CERN illustrates the scale involved: constructing its 27 km ring of superconducting magnets required roughly 7.5 billion, while its major detectors—ATLAS and CMS—each added budgets of about 400–500 million and collaborations of several thousand people. These detectors were conceived in the 1980s and then underwent multiple design iterations lasting decades.

In the future, the proposed Future Circular Collider (FCC) pushes these demands to an even more formidable scale. The current design envisions a 90–100 km tunnel that would host a machine reaching collision energies around three times larger than at the LHC. Such ambitions hinge on advances that are still at the edge of what is technically achievable, from 16 T-class superconducting magnets and vast cryogenic plants to unprecedented vacuum, alignment, and power-delivery systems. As with the LHC, the FCC design process will span decades and must navigate substantial technical, financial, and scientific uncertainties making systematic design methodologies essential from the outset.

So far, detector design has been largely a manual process, consisting of iterations centered around implementations of manually chosen designs in simulation. This iteration process typically takes many years and a large amount of full time effort. Decisions on the design of an experiment are often factorised into different areas, meaning that a global optimisation is typically not achieved.

With the advent of modern machine learning techniques in particle physics, attention has turned into how one can automate the design of particle physics detectors with the aid of artificial intelligence. Beyond the direct implications of more performant and cost effective detectors, this challenge is also highly valuable as a testing ground to develop machine learning methods as the data is readily generated but the challenges suffer from many complications that are vital to solve in order to apply these methods to real world problems. In particular, possibilities to improve the robustness of the optimisation results via mathematical techniques such as optimal transport would allow both more reliable designs as well as methods can be applied outside of physics. This review will first discuss the challenge of designing a detector, followed by a discussion of the main methodologies that are used. Finally, an outlook of the next developments will be discussed, which main focusses on the introduction of realistic uncertainties into the design procedure.

To enable systematic optimisation of particle-physics detectors, we begin by establishing a structured representation of a detector configuration. We denote a detector by

$$\mathcal{D} = \{C_1, C_2, \ldots, C_N\},$$

where each C_n corresponds to a sub-detector or functional module such as a tracking layer, calorimeter element, or timing plane. The number of components N is flexible and depends on the overall detector concept.

Each component C_n is parameterised by a set of variables

$$C_n = \left[p_1, p_2, \ldots, p_{K_n}\right],$$

which encode physical layout, material choices, and relevant engineering properties. These parameters may be continuous ($p_k \in \mathbb{R}$), discrete (such as layer counts or segmentation steps), or categorical ($p_k \in \mathcal{S}$) representing, for example, the sensor technology or material type.

In realistic detector concepts, components are often not monolithic objects but hierarchical assemblies. A given C_n may itself be composed of multiple sub-components modules, sensors, support structures, electronics each with its own internal parameter set and constraints. This hierarchical decomposition can be extended recursively to arbitrary depth, providing a flexible and detailed description of complex detector systems. For the purposes of this conceptual formulation, however, we treat each C_n as a single, self-contained unit with an associated parameter vector. More detailed structures can be incorporated straightforwardly by expanding or nesting these parametrisations whenever required.

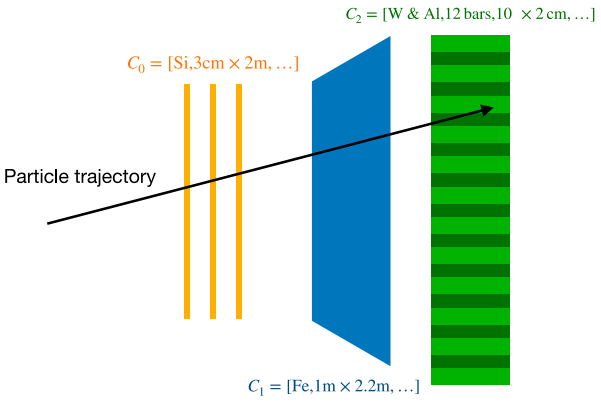

FIGURE 1
Diagram showing an illustration of a particle-physics detector with three sub-detectors. The lists on the top indicate the properties (e.g., material, size) of the sub-detectors that are to be optimised for performance.

A simple illustration is shown in Figure 1, which depicts a simplified detector with three sub-detectors. The lists above each element indicate representative properties such as material or size that must be optimised for performance. Realistic detector descriptions contain many such components, often with correlated parameters and practical constraints arising from mechanical tolerances, service routing, symmetry requirements, and radiation limits.

Given a detector configuration $\mathcal{D}$, we assign a scalar score function

$$S(\mathcal{D}) \in \mathbb{R},$$

which quantifies the quality of the design. Although detector performance is inherently multi-objective, balancing for example, resolutions, efficiencies, particle-identification power, background rejection, cost, and engineering feasibility, it is useful to encode these considerations in a score function $S(\mathcal{D})$ that reflects the intended physics programme. Different physics goals naturally imply different utilities, and $S(\mathcal{D})$ should be understood as encoding prior assumptions about which measurements or signatures are most valuable. The score need not be a linear combination of metrics and may instead capture nonlinear trade-offs or threshold requirements, or the optimisation can be formulated in a multi-objective manner by characterising the Pareto-optimal set of designs. In practice, $S(\mathcal{D})$ depends on four components: simulation, reconstruction, analysis, and cost.

The design task then becomes an optimization problem,

$$\mathcal{D}^\star = \arg\max_{\mathcal{D}} S(\mathcal{D}),$$

subject to all physical, mechanical, and engineering constraints the detector must satisfy. These may impose bounds on material budgets, clearances, power consumption, radiation tolerance, or the

allowed positions and sizes of components. This parametrization provides a concrete and flexible foundation for applying optimization and AI-driven search methods to detector design.

2 The challenge of detector design

The main challenge of detector design is the huge space of possible solutions to explore and evaluate to find the optimum. In order to detect the full suite of potential physics processes, particle physics detectors consist of several sub-detectors that are each optimised for a particular task. For example, one can have tracking detectors which measure particle momentum or calorimeters which measure particle energy. Designing a complete detector therefore requires simultaneously optimising these different sub-detectors, which results in a very large design space. The realities of engineering requirements constrain the design solution space in complex ways. Constraints cut across many subsystems such as power, cooling and radiation tolerance. A small improvement in geometry for physics performance can have implications in calibration procedures and civil-engineering needs. Algorithms or machine learning models that can efficiently explore this space, while navigating real world complicities are therefore an essential component of the automated design procedure.

These sub-detectors are not only specialised for different tasks but often provide overlapping measurements, for example, charged-particle momenta can be inferred from both tracking systems and calorimetry, and particle identification can rely on several complementary technologies. This redundancy is crucial for robust physics performance, yet it also means that different subsystems optimise similar observables in different ways. These subdetectors also must share limited space, material budget and readout infrastructure. As a result, improving the performance of one system can restrict what is feasible for another, turning the overall detector layout design into a multi-objective procedure. The balance between these different objectives makes the performance evaluation depend on the relative importance of different physics goals, which further complicates the process.

As detector design naturally considers instruments that do not yet exist, a detailed simulation of the detector as well as particle physics interactions are essential in order to evaluate the performance of a particular design. Exploring the large space of possible designs requires a large number of performance evaluations, with each evaluation needing billions of particle trajectories to be simulated. Even manual design places stringent requirements on the amount of computational time needed to perform these performance evaluations, making the speed of the simulation one of the main bottlenecks for the design. Therefore techniques to improve this are equally important as those to explore the optimisation space.

Machine learning has become central to this effort because generative models can learn the statistical structure of simulated data and then reproduce it at a fraction of the cost. Fast simulation using modern machine learning methods began with early applications of Generative Adversarial Networks, which showed that high-granularity calorimeter responses could be emulated orders of magnitude faster than traditional simulation [1]. These methods demonstrated that, once the relevant patterns are learned from detailed simulation, the computational burden can be shifted from repeated event generation to a one-time training stage. Building on this idea, more recent approaches use normalizing flows to learn probability distributions with exact likelihoods. This is especially useful for detector design, since flows do not simply generate approximate samples but also provide quantitative control over uncertainties. Calorimeter response models built with flows can reproduce the fine structure of showers with high fidelity while remaining efficient enough to evaluate many design variants [2]. Because flow models are invertible maps, they also allow the conditioning on geometry, materials, or particle type to be handled in a principled way, which is important when exploring large configuration spaces. Most recently, diffusion models have entered the field and are showing considerable promise [3]. These models learn to reverse a gradual corruption process and can generate detailed multi-scale detector responses with high stability. Diffusion models also tend to avoid some of the training instabilities that affect adversarial methods, which is advantageous when performing systematic design scans.

In parallel, traditional transport is also being accelerated on modern hardware. AdePT [4] integrates with Geant4 to offload electromagnetic showers to GPUs, while Celeritas [5, 6] takes a GPU-first approach that achieves order-of-magnitude speedups with physics preserved within statistical uncertainties. In another work, histogram sampling was employed to increase the speed of muon transport by several orders of magnitude by running the operations on a GPU [7].

Finally, uncertainty affects all aspects of the design procedure. Physics processes and their interactions with matter are imperfectly known, as well as the operational conditions in the real experiment. Furthermore, the fast surrogates that are trained to accelerate simulation introduce their own uncertainties that need to be considered. Robust design is therefore essential, where the goal is not only to optimize nominal the performance, but also to favour solutions that remain valid under realistic simulation inaccuracies and operational reality.

3 Methodologies

Design optimisation problems across the physical sciences have long been approached with heuristic and numerical strategies, such as manual tuning or parameter scans. In particle physics, detector optimisation has traditionally relied on expert intuition and low-dimensional scans, with limited adoption of formal optimisation techniques due to the cost of realistic simulation. Motivated by the broader success of modern machine learning, the field is now turning toward more systematic methodologies. Current research has converged on three families that show particular promise for detector design: surrogate-based search, gradient-based optimisation, and reinforcement-learning frameworks.

3.1 Bayesian optimization

Bayesian optimization (BO) is well suited to detector-design problems in which the configuration

$$\mathcal{D} = [C_1, \ldots, C_N]$$

is parameterized by continuous variables and each evaluation of the score function requires a computationally expensive simulation. In such settings, BO offers a principled black-box strategy that combines a probabilistic surrogate, typically a Gaussian process, with an acquisition function that proposes new designs by balancing exploration and exploitation. This enables global search while keeping the number of simulator calls small and without relying on gradient information.

The surrogate-driven nature of BO makes it highly sample-efficient and particularly effective when the design landscape is nonlinear, multimodal, or sensitive to physics thresholds. A prominent demonstration is the optimization of the active muon shield for the SHiP experiment, where BO was used to tune magnet shapes, currents, and spacings under tight engineering constraints. Coupled with fast field calculations and large-scale cloud-based simulation, BO produced shield configurations with substantially reduced downstream muon flux and realistic geometries, outperforming manual tuning [8, 9]. Multi-objective BO pipelines have likewise been incorporated into the ePIC optimization framework for the Electron–Ion Collider, enabling Pareto analyses across detector technologies and competing design variants [10].

These advantages come with important limitations. Gaussian-process surrogates become expensive and less accurate as dimensionality increases, and the very sample efficiency that gives BO its power also constrains the breadth of design-space exploration. Consequently, BO is most effective for low-to moderate-dimensional continuous parameterizations where global structure matters but the design space does not require large-scale combinatorial flexibility.

3.2 Gradient-based optimization

Gradient-based optimization seeks to update detector parameters by following the gradient of a performance objective. If a detector configuration $\mathcal{D}^{(t)}$ can be differentiated with respect to its parameters, an update of the form

$$\mathcal{D}^{(t+1)} = \mathcal{D}^{(t)} + \eta_t \nabla_{\mathcal{D}} S\left(\mathcal{D}^{(t)}\right)$$

mirrors standard gradient-descent procedures in machine learning and, in principle, offers a sample-efficient way to navigate continuous design spaces.

The first concrete attempt to apply this philosophy to detector design was the Local Generative Surrogate Optimization (L-GSO) method introduced by [11]. L-GSO proposed learning a local generative surrogate of the simulator around the current design point and then differentiating through this surrogate to update the detector parameters. Although the method used a GAN, the choice of generative model was not fundamental, any differentiable surrogate capable of capturing the local expectation would have sufficed. Recent work has demonstrated end-to-end gradient-based optimization using diffusion-model surrogates for calorimeter optimisation [12].

Building on this motivation, subsequent work has attempted to make increasingly large portions of the simulation and reconstruction analysis chain differentiable. Physics simulation includes discrete branching processes such as bremsstrahlung and pair production; reconstruction algorithms rely heavily on thresholding, clustering, and categorical decisions; and analyses introduce further discontinuities. These features make end-to-end differentiation extremely challenging. Efforts within the MODE collaboration [13] pursue differentiable approximations using hybrid Monte Carlo gradient estimators: pathwise derivatives handle smooth parametric components, while score-function terms are used for stochastic branching [14, 15]. In particular, differentiable design techniques developed for muon-tomography applications [16]. For example, fast electromagnetic-shower surrogates augmented with pathwise derivatives allow gradient-based adjustments to sampling-calorimeter layouts [17]. However, these gradients remain tied to reduced models and inherit their limitations.

Despite these developments, substantial obstacles remain. Surrogate-based gradients are only as reliable as the neighborhood in which the surrogate is trained, and their accuracy deteriorates rapidly as one moves away from that region. Differentiable physics approximations often require simplifying assumptions that compromise realism, and noisy or biased gradients can quickly destabilize optimization. Moreover, detector-parameter spaces are typically low dimensional and highly structured, yielding nonconvex loss landscapes that lack the benign optimization properties often seen in large neural networks [18–21].

3.3 Reinforcement learning

Reinforcement learning (RL) provides a framework for optimizing sequential decisions when only a scalar reward signal is available. An agent interacts with an environment, observes states, takes actions, and ultimately receives a reward that reflects the quality of the completed design (Figure 2). A central ingredient in modern RL is temporal-difference (TD) learning, which allows value estimates to be updated from partial trajectories rather than waiting for full returns. This is especially important in settings where each evaluation is expensive: TD learning can reuse intermediate predictions to improve sample efficiency and reduce the number of full simulations required.

These properties have made RL successful in several design domains with combinatorial or constrained action spaces, including molecular generation [23], chip placement [24, 25], and applications in materials [26] and drug design [27]. In these settings, long decision chains, sparse rewards, and high simulation costs make RL particularly attractive.

In particle physics, RL has been explored for accelerator control [28–30], online data-quality monitoring [31], and event-level reconstruction and interpretation [32, 33]. For detector design, RL is especially appealing because a detector can be constructed sequentially: the state encodes a partial configuration, actions modify or extend it, and the episode terminates once the full design is specified. The reward is the score assigned by the simulation–reconstruction–analysis pipeline. Because the reward is terminal, TD learning plays a crucial role in propagating information back through the design process, enabling the agent to identify sequences of decisions that lead to strong performance.

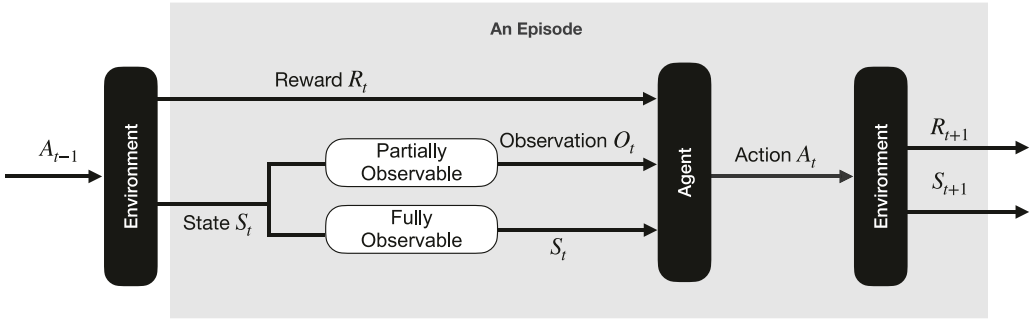

FIGURE 2
Agent-environment interaction in Reinforcement Learning, reproduced from [22].

RL naturally accommodates discrete, structural, and continuous choices, making it well suited to the combinatorial nature of detector design where the number, types, and arrangement of components are not fixed in advance. It also places no differentiability requirements on the simulator. Recent studies have shown that RL can optimize longitudinal calorimeter segmentation and joint tracking–calorimetry layouts in spectrometer-style geometries [22], often discovering design families that would be difficult to reach with local or gradient-based methods.

4 Discussion

Detector design lives under deep uncertainty, from incomplete knowledge of the underlying physics (cross sections, pile-up, underlying-event models) to manufacturing tolerances, alignment and calibration drifts, and rare backgrounds that appear only after long running periods. Historically, these effects are handled by expert-driven stress tests: identify a few plausible worst-case scenarios and verify that the baseline design can cope. While invaluable, such workflows can miss hard-to-anticipate combinations and do not use the optimisation loop itself to discover which scenarios matter most. If one aims to exploit the full potential of AI-assisted design, it is natural to elevate uncertainty to a first-class design parameter. The goal is to find solutions whose performance is strong on average and degrades gracefully under model misspecification and operational shifts.

It is convenient to distinguish, following standard particle-physics practice, between aleatoric variability that is well described statistically (for example, stochastic detector response or shot noise) and systematic shifts that reflect limited or biased knowledge (for example, misalignments or out-of-distribution rare processes).

In optimisation terms, the first category can be treated as expectations under a known or estimated distribution, whereas the second calls for robust objectives that hedge against an ambiguity set of distributions or scenarios.

Recent developments in RL (we follow standard notation; see Sutton and Barto [34]) provide a natural place to incorporate both. Stochasticity appears in the distribution of returns and can be handled with distributional RL [35], while model misspecification can be addressed with distributionally robust RL [36].

In RL language, the return is a random variable. Distributional RL therefore models its full distribution $Z_\pi(s,a)$ rather than only its mean $Q_\pi(s,a) = \mathbb{E}[Z_\pi(s,a)]$. Here $Z_\pi(s,a)$ denotes the distribution of the return under policy π, while $Q_\pi(s,a) = \mathbb{E}[Z_\pi(s,a)]$ is its mean. This enables optimising tail-sensitive criteria such as entropic risk or conditional value at risk (CVaR), which explicitly target lower-quantile performance.

More challenging are uncertainties arising from systematic distributional shifts, where the nominal transition model may be biased or incomplete. Optimal transport methods provide a natural framework for this situation. In distributionally robust RL, given $\varepsilon > 0$, the allowed mismatch between the true and nominal dynamics, one specifies an ambiguity set $\mathcal{U}_\varepsilon$ of plausible transition kernels or reward models, commonly taken as a Wasserstein ball around the nominal dynamics [37]. Intuitively, $\mathcal{U}_\varepsilon(s,a)$ collects all transition models that differ from the nominal simulator $P_0(\cdot \mid s,a)$ by at most ε, measured using a geometry-aware notion of distance between probability distributions.

For a nominal transition kernel $P_0(\cdot|s,a)$, a typical choice is

$$\mathcal{U}_\varepsilon(s,a) = \{P(\cdot|s,a) \quad : \quad W_c(P(\cdot|s,a), P_0(\cdot|s,a)) \le \varepsilon\},$$

where W_c denotes the optimal-transport cost induced by the cost function c. Then, the corresponding robust Bellman operator acts on a value function $V(s)$ and incorporates an adversarial inner optimisation:

$$(T_{\text{rob}}V)(s) = \max_a \min_{P\in\mathcal{U}_\varepsilon(s,a)} \left\{ r(s,a) + \gamma \mathbb{E}_{s'\sim P(\cdot|s,a)}[V(s')] \right\},$$

producing policies that remain stable under model misspecification, alternative physics tunes, or calibration envelopes. This update can be interpreted as a two-player game in which the optimiser selects the action a, while an adversary chooses, within $\mathcal{U}_\varepsilon(s,a)$, the most unfavourable plausible dynamics, leading to policies that are robust to model misspecification.

Another aspect that will be important to consider will be the fact that many performance metrics depend on the combination of two or more detectors, for example, particle identification that combines momentum from a spectrometer with specialised particle identification detectors. Conceptually, RL can accommodate this by expanding the state to include the configuration of all relevant subsystems and allowing actions that modify any of

them; however, the coupled nature of the reward makes scalable optimisation nontrivial. Practical strategies include (i) an end-to-end optimisation over the full layout and all subsystem parameters, (ii) a two-stage procedure where a high-level search chooses the layout/technology mix and a lower-level optimiser tunes subsystem parameters, or (iii) multi-agent RL with subsystem-level policies trained under a shared global reward. In all cases, the bottleneck is maintaining a consistent multi-subdetector simulation so that correlations relevant for the global performance are preserved.

Looking ahead, the full potential of AI-assisted detector design will only be realised once realistic engineering constraints are embedded directly into the optimisation loop. These ingredients introduce sharp nonlinearities, discrete feasibility boundaries, and complex dependencies that are difficult to treat with standard black-box or gradient-based methods. RL is particularly well suited to this regime because it handles sequential design decisions, integrates constraint checks through the environment dynamics, and remains effective even when the action space mixes continuous and discrete choices. As these tools mature, the optimisation process can move from tuning idealised detector concepts to exploring realistic design families consistent with engineering practice and long-term operability.

Detector design is becoming a powerful proving ground for advanced optimisation and learning methods. The main difficulties it presents, such as complex objectives, strict engineering constraints, uncertainty and long decision chains, are shared with many real-world applications, yet in particle physics they can be explored in a more controlled setting thanks to reliable simulations and well-understood physical models. The proposed FCC, with its long design horizon and unprecedented complexity, would provide an ideal environment to develop and test these approaches at scale. In this perspective, the evolution of detector optimisation methods, the emergence of RL approaches that handle distributional shifts through optimal transport, and the design effort for future facilities such as the FCC form a mutually reinforcing cycle, in which progress in one area accelerates progress in the others.

Author contributions

AF: Writing – original draft, Writing – review and editing. SQ: Writing – original draft, Writing – review and editing. PO: Writing – review and editing, Writing – original draft. NS: Writing – original draft, Writing – review and editing.

Funding

The author(s) declared that financial support was received for this work and/or its publication. N.S. acknowledges support from the Swiss National Science Foundation through the Advance Grant TMAG-2_209263. This work was supported by the Swiss High Energy Physics initiative for the FCC (CHEF), with funding provided specifically by SERI and the University of Zurich. Open access funding by ETH Zurich.

Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declared that generative AI was not used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

1. Paganini M, Oliveira Lde, Nachman B. Calogan: simulating 3d high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks. *Phys Rev D* (2018) 97(1). ISSN 2470-0029. doi:10.1103/PhysRevD.97.014021

2. Krause C, Shih D. Fast and accurate simulations of calorimeter showers with normalizing flows. *Phys Rev D* (2023) 107(11):113003. doi:10.1103/PhysRevD.107.113003

3. Mikuni V, Nachman B. Caloscore v2: single-shot calorimeter shower simulation with diffusion models. *J Instrumentation* (2024) 19(02):P02001. ISSN 1748-0221. doi:10.1088/1748-0221/19/02/P02001

4. Amadio G, Apostolakis J, Buncic P, Cosmo G, Dosaru D, Gheata A, et al. Offloading electromagnetic shower transport to gpus. *J Phys Conf Ser* (2023) 2438(1):012055. doi:10.1088/1742-6596/2438/1/012055

5. Tognini SC, Canal P, Evans TM, Lima G, Lund AL, Johnson SR, et al. Celeritas: Gpu-accelerated particle transport for detector simulation in high energy physics experiments (2022). Available online at: https://arxiv.org/abs/2203.09467 (Accessed February 3, 2026).

6. Lund AL, Esseiva J, Johnson SR, Biondo E, Canal P, Evans T, et al. Accelerating detector simulations with celeritas: profiling and performance optimizations. *EPJ Web of Conferences* (2025). 337:01292. doi:10.1051/epjconf/202533701292

7. Cattelan LFP, Rukh Qasim S, Owen PH, Serra N. Ultra-fast muon transport via histogram sampling on gpus (2025). Available online at: https://arxiv.org/abs/2512.10520 (Accessed February 3, 2026).

8. Kurbatov E, Ratnikov F, Ursov E. Multi-parameter detector optimization: ship case. In *EPJ Web Conf EPJ Web of Conferences (CHEP 2023 Proceedings)* (2025). 295, 03023. doi:10.1051/epjconf/202429503023

9. Baranov A, Burnaev E, Derkach D, Filatov A, Klyuchnikov N, Lantwin O. Optimising the active muon shield for the SHiP experiment at CERN. *J Phys Conf Ser*. Bristol, United Kingdom: IOP Publishing (2017). doi:10.1088/1742-6596/934/1/012050

10. Diefenthaler M, Fanelli C, Gerlach L, Guan W, Horn T, Jentsch A, et al. AI-assisted detector design for the EIC (AID(2)E). *JINST* (2024) 19(07):C07001. doi:10.1088/1748-0221/19/07/C07001

11. Shirobokov S, Belavin V, Kagan M, Ustyuzhanin A, Baydin AG. Black-box optimization with local generative surrogates. In: H Larochelle, M Ranzato, R Hadsell, MF Balcan, H Lin, editors. *Advances in neural information processing systems*, 33. Red Hook, NY: Curran Associates, Inc. (2020). p. 14650–62. Available online at: https://proceedings.neurips.cc/paper_files/paper/2020/file/a878dbebc902328b41dbf02aa87abb58-Paper.pdf (Accessed February 3, 2026).

12. Schmidt K, Kota KN, Kieseler J, De Vita A, Klute M, Abhishek MA, et al. End-to-end detector optimization with diffusion models: a case study in sampling calorimeters. *Particles* (2025) 8(2):2571–712X. doi:10.3390/particles8020047

13. Güneş Baydin A, Cranmer K, De Castro Manzano P, Delaere C, Derkach D, Donini J, et al. Toward machine learning optimization of experimental design. *Nucl Phys News* (2021) 31(1):25–8. doi:10.1080/10619127.2021.1881364

14. Dorigo T, Giammanco A, Vischia P, Aehle M, Bawaj M, Boldyrev A, et al. Toward the end-to-end optimization of particle physics instruments with differentiable programming. *Rev Phys* (2023) 10:100085. doi:10.1016/j.revip.2023.100085

15. Aehle M, Arsini L, Barreiro RB, Belias A, Boldyrev A, Bury F, et al. Progress in end-to-end optimization of fundamental physics experimental apparata with differentiable programming. *Rev Phys* (2025) 13:100120. doi:10.1016/j.revip.2025.100120

16. Strong GC, Lagrange M, Orio A, Bordignon A, Bury F, Dorigo T, et al. Tomopt: differential optimisation for task- and constraint-aware design of particle detectors in the context of muon tomography. *Machine Learn Sci Technology* (2024) 5(3):035002. doi:10.1088/2632-2153/ad52e7

17. Aehle M, Novák M, Vassilev V, Gauger NR, Heinrich L, Kagan M, et al. Optimization using pathwise algorithmic derivatives of electromagnetic shower simulations. *Computer Phys Commun*, 309:109491 (2025) doi:10.1016/j.cpc.2024.109491

18. Chizat L, Colombo M, Fernández-Real X, Figalli A. Infinite-width limit of deep linear neural networks. *Commun Pure Appl Mathematics* (2024) 77(10):3958–4007. doi:10.1002/cpa.22200

19. Mei S, Montanari A, Nguyen P-M. A mean field view of the landscape of two-layer neural networks. *Proc Natl Acad Sci* (2018) 115(33):E7665-E7671. doi:10.1073/pnas.1806579115

20. Shaolei Du S, Zhai X, Póczos B, Singh A. Gradient descent provably optimizes over-parameterized neural networks. *ArXiv, abs/1810.02054* (2018). Available online at: https://api.semanticscholar.org/CorpusID:52920808 (Accessed February 3, 2026).

21. Nguyen Q, Chandra Mukkamala M, Hein M. On the loss landscape of a class of deep neural networks with no bad local valleys. In: *International conference on learning representations* (2019).

22. Rukh Qasim S, Owen PH, Serra N. Physics instrument design with reinforcement learning. *Machine Learn Sci Technology* (2025). doi:10.1088/2632-2153/adf7ff

23. Olivecrona M, Blaschke T, Engkvist O, Chen H. Molecular *de-novo* design through deep reinforcement learning. *J Cheminformatics* (2017) 9(1):48. doi:10.1186/s13321-017-0235-x

24. Mirhoseini A, Goldie A, Yazgan M, Jiang JW, Songhori E, Wang S, et al. Author correction: a graph placement methodology for fast chip design. *Nature* (2022) 604(7906):E24. doi:10.1038/s41586-022-04657-6

25. Schulman J, Wolski F, Dhariwal P, Radford A, Klimov O. Proximal policy optimization algorithms (2017). Available online at: https://arxiv.org/abs/1707.06347 (Accessed February 3, 2026).

26. Guo K, Yang Z, Yu C-H, Buehler MJ. Artificial intelligence and machine learning in design of mechanical materials. *Mater Horizons* (2021) 8(4):1153–72. doi:10.1039/D0MH01451F

27. Popova M, Isayev O, Tropsha A. Deep reinforcement learning for *de novo* drug design. *Sci Adv* (2018) 4(7):2375–548. doi:10.1126/sciadv.aap7885

28. Kain V, Hirlander S, Goddard B, Velotti FM, Porta GZD, Bruchon N, et al. Sample-efficient reinforcement learning for cern accelerator control. *Phys Rev Accel Beams* (2020) 23:124801–12. doi:10.1103/PhysRevAccelBeams.23.124801

29. Simon H, Bruchon N *Model-free and Bayesian Ensembling model-based deep reinforcement learning for particle accelerator control demonstrated on the FERMI FEL*, 12 (2020). doi:10.48550/ARXIV.2012.09737

30. Ibrahim A, Derkach D, Petrenko A, Ratnikov F, Kaledin M. Optimization of the accelerator control by reinforcement learning: a simulation-based approach. *Phys Particles Nuclei* (2025) 56(6):1476–81. ISSN 1531-8559. doi:10.1134/S1063779625700716

31. Jullian Parra O, Pardiñas JG, Pianta Pérez LD, Janisch M, Klaver S, Lehéricy T, et al. Human-in-the-loop reinforcement learning for data quality monitoring in particle physics experiments. *Machine Learn Sci Technology* (2025) 6(4):2632–153. doi:10.1088/2632-2153/ae1563

32. Kortus T, Keidel R, Gauger NR. And Jan Kieseler. Constrained optimization of charged particle tracking with multi-agent reinforcement learning (2025). Available online at: https://arxiv.org/abs/2501.05113 (Accessed February 3, 2026).

33. Dillon BM, Spannowsky M. Theory-informed neural networks for particle physics. 7 (2025). doi:10.48550/ARXIV.2507.13447

34. Sutton RS, Barto AG. *Reinforcement learning: an introduction*. 2 edition. Cambridge, MA: MIT Press (2018). Available online at: http://incompleteideas.net/book/the-book-2nd.html (Accessed February 3, 2026).

35. Bellemare MG, Dabney W, Munos R. A distributional perspective on reinforcement learning. In: D Precup, YW Teh, editors. *Proceedings of the 34th international conference on machine learning, volume 70 of proceedings of machine learning research*. Netherlands: PMLR (2017). p. 449–58. Available online at: https://proceedings.mlr.press/v70/bellemare17a.html (Accessed February 3, 2026).

36. Moos J, Hansel K, Abdulsamad H, Stark S, Clever D, Peters J. Robust reinforcement learning: a review of foundations and recent advances. *Machine Learn Knowledge Extraction* (2022) 4(1):276–315. ISSN 2504-4990. doi:10.3390/make4010013

37. Mohajerin Esfahani P, Kuhn D. Data-driven distributionally robust optimization using the wasserstein metric: performance guarantees and tractable reformulations. *Math Programming* (2017) 171(1–2):115–66. ISSN 1436-4646. doi:10.1007/s10107-017-1172-1