Searching high and low: precision measurement, machine learning, and experimental discovery in particle physics

Adam Koberinski* †

Abstract

Discoveries in particle physics have traditionally occurred through direct detection, usually by producing new particles in particle accelerators. However, the lack of new discoveries from the Large Hadron Collider at CERN has undermined confidence that direct discovery is likely at technologically accessible energies. In this paper, I provide an epistemic analysis of more and less direct search strategies in current-day particle physics. I make explicit the epistemic features of these search strategies in what is has traditionally been called the process of scientific discovery. I describe two promising methods of indirect detection of new physics, possible to conduct at comparatively lower energies and at significantly reduced cost: precision measurement and anomaly detection via unsupervised machine learning. I argue that indirect detections can reliably lead to discovery, and that direct versus indirect searches have complementary epistemic and practical advantages. I provide an account of indirect discovery that proceeds from discrepancies, to anomalies, and finally to (theory-mediated) discovery. I further argue that these methods are capable of leading to new discoveries when mediated through the generalized effective field theory framework. Indirect searches should not replace direct detection efforts; instead, they provide complementary epistemic advantages and jointly make for a more exhaustive search for physics beyond the Standard Model.

Acknowledgements

I am grateful to Chris Smeenk and Juliusz Doboszewski for helpful discussions on earlier versions of this paper. I am also grateful to two anonymous referees for their constructive feedback. This work is partially supported by a Social Sciences and Humanities Research Council of Canada Insight Grant for the project "The Structure of Effective Theories".

*Rotman Institute of Philosophy, University of Western Ontario, London, ON, Canada
†akoberin@uwo.ca

1 Introduction

Since the Large Hadron Collider (LHC) at CERN started running, it promised to find the key missing piece of the Standard Model of particle physics—the Higgs boson—and open a window into physics beyond that accounted for in the Standard Model. While it famously fulfilled the first promise between 2012-2014, it has thus far failed to deliver on the second. To make matters worse, its failure to discover new physics near the energy required to produce Higgs bosons has undermined our most plausible theoretical expectations constraining the energies at which new physics should occur (Giudice 2017). We are no longer sure that a new collider operating at energies one or two orders of magnitude higher than the LHC will find new physics. Even more pessimistically, we have no good reason to expect new physics to be detectable even at energies several orders of magnitude higher than currently accessible, straining the plausible physical limits on constructing new particle accelerators.[1] Given both the cost and technological feasibility of constructing bigger, more powerful accelerators, one might worry that the expected epistemic payoff is too low to justify funding for a new generation of accelerators. Does this mean that experimental discovery in particle physics might come to an end?

Traditionally, experimental discovery in particle physics has come in the form of what physicists call *direct detection* (Sec. 2). The history of particle physics post-WWII has been dominated by the construction of larger, more powerful particle accelerators, designed to impart subatomic particles with high energies before colliding them. Direct detection of new particles comes from colliding known particles together with high enough energies to produce the new particles in the lab. These particles are often short lived, and rapidly decay into other lighter products, until only stable particles remain. Some knowledge of particle physics is thus needed to reconstruct event histories from stable decay products, even for direct detection.

This type of direct detection has come in two stages: first was largely theory-agnostic direct detection. In the days before the construction of the Standard Model, experiment typically led the way, with many new particles discovered without a clear theoretical understanding of their underlying structure. Theorists were devising principles, concepts, and phenomenological models on the basis of experiment.[2] This work culminated in two quantum field theories (QFTs) describing the fundamental electroweak and strong forces. These theories describe a basic set of fundamental particles and their interactions; from these, all other known particle phenomena were reconstructed as composites of the fundamental particles, interacting via the fundamental forces.[3] After the construction of the Standard Model, experimental discovery became far more theory-mediated. The Standard Model itself was used in the design of new experiments, specifically those searching for the missing pieces of the model. Since the late 1970s, all new particle discoveries have confirmed expectations of the Standard Model, culminating in the Higgs discovery in 2012. However, excitement was high in 2012 for the LHC for more than just the Higgs. Physicists were hopeful for a new era of discovery in particle physics. Given the modern understanding of the structure of the Standard Model as an effective field theory (EFT), the Higgs mass plays an epistemically privileged role in that its value is proportional to the scale at which new physics beyond the Standard Model is supposed to become important (Koberinski 2025). Thus, physicists expected that the energy scale at which the Higgs could be reliably produced would be close to the energy scale at which other new physics effects—new forces, new particles, new symmetries, etc.—would be observed. The generalized framework of EFT provided meta-theoretic resources to suggest that

[1]Of course, we similarly do not have any reason to expect that new physics *won't* be there; but given the cost of building new accelerators, it seems likely that a more concrete expectation of payoff would be needed. To give a rough estimate of the size of accelerator needed to reach significantly higher energies than the LHC, a circular accelerator wrapped around the surface of the earth would be capable of creating a collision energy of around 10,000–20,000 times what is possible at the LHC, given substantial but feasible improvements to the strength of the magnets used to deflect the particle beams (i.e., less than an order of magnitude). Despite the implausibility of such an accelerator, this is only an improvement of 4 orders of magnitude. The gap between LHC energies and the Planck energy—where we have strong theoretical grounds to expect new physics—is 18 orders of magnitude. Thus, even outside of the cost of a new accelerator, the ability to probe meaningfully higher energies seems severely limited.

[2]This interplay of experimental discovery, followed by phenomenological models to classify and predict features of the new particles, all informed by high-level principles from the theoretical framework is what led to the construction of the strong and electroweak models forming the Standard Model (cf. Karaca 2013a; Koberinski 2019; Koberinski 2021).

[3]The word "fundamental" here is used in a comparative sense. According to the Standard Model, the most fundamental particles are those given in its Lagrangian. Comparatively less fundamental particles are treated as composites of the fundamental particles. But what is fundamental according to the Standard Model is not necessarily fundamental in nature. Given the attempts to formulate theories more fundamental than the Standard Model immediately after its construction, it is fair to assume that usage of the term "fundamental" in particle physics is both comparative and theory-relative.

the standard operating procedure of direct detection would continue on after the Higgs, at nearby energies. Further, several specific candidate beyond Standard Model theories were guiding the search, in particular models that included supersymmetry. These two prongs of theory mediation both supported the search, providing grounds to suppose that another direct detection would happen near the Higgs scale. With the failure to detect even indirect evidence for new physics has come a crisis in our understanding of naturalness and the Standard Model as an EFT (Giudice 2017; Wallace 2022; Koberinski 2025). While direct detection might be possible at the next generation of accelerators, the crisis has prompted physicists to look for new physics via means other than theory-mediated direct detection.

In this paper, I provide an epistemic analysis of more and less direct search strategies in current-day particle physics. I make explicit the epistemic features of these search strategies in what is has traditionally been called the process of scientific discovery[4] I argue that there are several promising avenues for *indirect detection* of new physics in particle physics, and that these provide complementary epistemic and practical benefits to standard efforts for high-energy, direct detection. Indirect detection is often comparatively cheaper and easier than high-energy accelerator experiments seeking direct detection. Indirect detection therefore presents a promising and reliable, though indirect, path to discovery. After providing a conceptual analysis of direct detection and theory-agnostic searches for new physics in Sec. 2, in Sec. 3 I provide an account of how one can go from a discrepancy between Standard Model prediction and experiment to a discovery claim, even when only indirect evidence of some physical source is found. When one has a guiding framework, the structure of the inference pattern is relatively simple. In the context of modern day particle physics, the EFT framework provides that structure, allowing for inferences to stable sources in the form of non-renormalizable terms added to the Standard Model EFT extension. I describe two promising avenues for indirect detection: (comparatively) low-energy precision experiments (Sec. 4.1) and anomaly detection using machine learning algorithms (Sec. 4.2). The two avenues differ, but both offer indirect paths to discovery. In Sec. 4.1, I outline how precision measurement of known quantities can be used to find anomalies, and argue that these anomalies can be structured within an EFT framework to lead to eventual discoveries. Precision of a measured quantity can be increased in two ways: through reducing systematic uncertainty or reducing statistical uncertainty. Depending on the quantity of interest, both can play an important role in anomaly detection. In Sec. 4.2, I outline how unsupervised machine learning can be used as a tool to search old datasets for candidate anomalies. Interestingly, many of the known conceptual problems for discovery via unsupervised machine learning are turned into advantages in this particular context. In Sec. 5 I argue that indirect searches provide complementary epistemic advantages to high-energy direct searches; the weaknesses involved in one strategy are complemented by the strengths of the other. When pursued together, they provide broad coverage possibility space for theories going beyond the Standard Model. While this is an analysis of the epistemology of experimental search strategies in the context of particle physics, it depends only on the more general features of direct versus indirect detections, and the relative theory-agnosticism of the latter. Insofar as these features can generalize to other scientific disciplines, so too would the conclusions of Sec. 5. Sec 6 concludes the paper.

2 Directness and theory-mediation in particle physics experiments

Particle physicists often advertise new experimental search techniques as being able to provide 'direct detections' of some target phenomena (LHCb Collaboration et al. 2023; Misiaszek and Rossi 2024), or allowing for 'theory-independent' or 'model-independent' characterizations of the evidence (Ghosh et al. 2020). These are meant to be attractive features of a given setup, and supposedly contrast with other search techniques that are not direct, or theory-dependent, respectively. In this section I provide a conceptual analysis of these two features of experimental searches, by highlighting features of *direct* versus *indirect* detections of particle physics phenomena and *theory-mediated* versus *theory-agnostic* experimental searches.[5] The latter

[4] Though it is outside the scope of this paper, I believe that the distinction between contexts of discovery and justification is not a useful one. Science is a process of building new theories, frameworks, and concepts for understanding the world, and subjecting them to critical test via contact with empirical evidence. The two contexts are thus always linked in practice.

[5] I use the modifier 'agnostic' rather than 'independent' to indicate that these are not completely free of theoretical commitments, as some background theory is needed to interpret almost any experimental setup and result (Smith 2014). I further use 'theory-mediated' rather than the more familiar 'theory-laden' to avoid the negative connotation often attributed to theory-ladenness. I find this usage clearer, despite its departure from standard philosophical debates.

distinction has to do with the degree to which the experimental setup depends on the particular theories being tested, while the former distinction tracks the ability to infer from a positive experimental result to a given physical source. Note that these distinctions are matters of degree: detections can be more or less direct, and can be dependent on some features of a given theory but agnostic to others. Nevertheless, they help illuminate the different epistemic and pragmatic features of high-energy collider experiments on the one hand, and precision measurements and machine learning searches on the other. Further, they are not simply features of an experiment independent of the aims and context in which it is conducted; these characterizations depend on the experimental setup itself, the theoretical context in which the experiment is conducted, and the pragmatic aims of researchers conducting the experiment.[6]

As a first pass on these distinctions, take an experiment that tests against a null hypothesis: a positive result is a result that provides evidence that the null hypothesis is false, while a negative result is consistent with the null hypothesis. An experiment *detects* something new if it supplies a positive result. The experiment must be interpreted in light of the background theoretical context; sometimes the null hypothesis is contrasted directly with a competitor, while other times not. It is also possible that the result of the experiment forces a new hypothesis on us more directly. A *direct detection* is one that allows for a more immediate inference from a positive result to an attribution of the effect to some physical source, given the background theoretical context. An *indirect detection* requires some degree of mediation. The forms of inference, the degree of mediation, and theoretical context will differ widely in different types of experiments; I provide some examples of more and less direct inferences below. Given the background theoretical context, one can interpret a direct detection as a relatively unmediated link between the physical source and the experimental apparatus; in this case, the apparatus directly measures or detects some signal from the source. An ideal direct detection would be to produce an observable entity or phenomenon. Most scientific experiments do not result in an ideal direct detection, but it is still useful to consider experiments that allow for more or less mediated inferences. Thus direct and indirect detection is comparative, and depends on the ability to infer from an experimental result to an underlying source.

Skulberg and Elder (2025) articulate various senses of the term 'direct' in the context of direct imaging of black holes, and direct detection of gravitational waves. Both of these are observational contexts. One example they give as a signal of directness has to do with how the measuring apparatus is modelled, and whether it can be modelled as directly interacting with the desired physical source. This is one general example of providing a more immediate inference from result to physical source. Indirect detections require some additional model of how the data and physical source relate. In my terminology, this is part of the theoretical context of the experiment. The model involving interaction between apparatus and source includes a prescription for what one would expect to observe in the apparatus, should an interaction take place. This allows for a more immediate inference from an observed outcome to an attribution to the source. The inference to the source is more direct to the extent that the theoretical context allows one to infer from a signal to features of the source. In particle physics, a standard for relatively direct detection is the production of some particles whose signatures qualitatively change the distribution of decay products in some way. In searches for beyond Standard Model physics, direct detection often refers to the on-shell production of the new particles in a collider (Misiaszek and Rossi 2024; LHCb Collaboration et al. 2023). One marker of a particle appearing on-shell it that the reconstructed event is modelled as producing the new particle as one of the external lines of a Feynman diagram, not simply that the new particle's contribution to internal lines needs to be factored into the modelling. For the remainder of the paper, I will use the term 'direct detection' to refer to production experiments that produce or detect new particles on-mass-shell.

By contrast, the distinction between theory-mediated and theory-agnostic experiments has more to do with the interpretation of the experiment, prior to any result. An experiment is more *theory-mediated* to the extent that its design and implementation requires assumptions regarding the specific theory or theories one is seeking to test, while it is *theory-agnostic* to the extent that it is free of such assumptions. Again, experiments can come in various degrees of theory-mediation and theory-agnosticism, so these terms are best thought of as comparative. While much has been said about the threat of circularity for theory-laden experiments, there have been several local defences of theory-laden experiments, showing that in successful cases the threat of circularity can be adequately managed (Beauchemin 2017; Karaca 2013b; Ritson and

[6] Beauchemin and Staley (2024) have argued a similar point regarding the classification of experiments as exploratory. Rather than being distinct types of experiments, exploration is a certain attitude one can take toward the act of experimentation, and exploratory experimentation is defined largely by the context and aims of the experimental practice.

Staley 2021). I take these to provide strong support for the idea that theory-ladenness is not a problem on its own; one must also consider the overall theoretical and experimental context.

First, note that both of these distinctions apply to an experiment in a certain theoretical context. The exact same experiment can lead to direct or indirect detections, depending on the theoretical and pragmatic contexts in which it is conducted. An experiment can be mediated by one relevant background theory, but remain agnostic to other theories for which its results might be relevant. Next, the two dimensions here can relate to each other, but need not always go together. For example, a highly theory-mediated experiment might typically make for a more direct-detection: since the setup relies on the theory being tested, it would be more straightforward to infer from a positive result to physical source. But it is possible for a theory-mediated experiment to result in a more indirect detection, if the theory-mediation in experimental design does not tightly constrain the possible outcomes.

Searches are experiments designed to look for some (range of) as yet undetected physical effect(s). Searches might be guided heavily by certain expectations, or may be more exploratory in nature. As such, they may fall under direct or indirect detection, and may be theory-mediated or theory-agnostic. Traditionally, philosophers have argued that experiments mediated by a particular theory cannot be properly exploratory with respect to possibilities beyond that theory, since the range of possible searches is limited by the dependence on theory (Steinle 1997). Elliott (2007) and Karaca (2017), among others, have refuted this point, arguing that exploration does not necessitate a lack of theoretical guidance . Beauchemin and Staley (2024) argue that exploration is not defined by a distinctive set of exploratory procedures or characteristics, but is instead inherent in the context of inquiry. I take for granted Beauchemin and Staley's point that experiments—including indirect searches—can be theory-mediated and exploratory, depending on the overall context in which they are conducted.[7]

Particle physics is a discipline whose subject matter is highly epistemically remote: the subatomic particles and their interaction forces are far too small and short-lived to be directly observable. This means that all experiments in particle physics rely on the use of instruments to detect the effects of particle physics phenomena and transform them into a form that is observable and interpretable to humans. The instruments themselves are modelled as physical systems, and scientists use theory to model the interactions between subatomic particles and the instruments. This modelling often depends on assumptions regarding the Standard Model of particle physics. We do not directly observe the sources of particle physics phenomena, nor are we able to interpret what we do observe without the guidance of at least some scientific theories. In this sense, *all* particle physics experiments provide *indirect* evidence of underlying sources, and all are *theory-mediated* to some extent.

Moreover, most high-energy collision experiments in contemporary particle physics are not just somewhat theory-mediated, but highly so. Some simple features of new particles can be reconstructed with minimal theoretical input, such as the particle mass from a resonance in the differential cross-section for particle scattering, while others require further reconstruction and input from theory. Several recent studies have examined the degree to which the Standard Model is integral to the design of experiments, the functioning and calibrating of the instrumentation, the choice of hardware and software cuts, as well as the separation of signal from background (Karaca 2013b; Beauchemin 2017; Mättig and Stöltzner 2020). The same studies have all argued that this theory-mediation is not viciously circular, but it remains true that contemporary high-energy experiments are expensive to run and highly theory-mediated.

In the history of particle physics, experimental practice started as mostly theory-agnostic through the 1940s-early 1960s, and has since become heavily mediated by expectations from the Standard Model (cf. Brown, Dresden, and Hoddeson 1989; Cushing 1990; Hoddeson et al. 1997; Cao 2010). Claims of discovery of new physics—new particles and new interactions between them—have been reserved for the most direct possible detections. This typically occurs by creating conditions under which the desired particle is produced, and reconstructing the event from the resulting decay products. While early instances of direct detection came from cosmic rays, current direct detection involves producing the sought-after particles in an accelerator. Indirect detections have played a secondary role in discovery, as only taken to provide evidence for a new phenomena. For example, there was indirect evidence that the weak force violated parity symmetry, but this

[7] See Sec.'4.1 and Koberinski (2023a) for ways that the theoretical framework can be modified to turn precision determinations of constants—something taken even by Karaca as a paradigm example of a non-exploratory type of experiment—into explorations of a more general theory space. I take this to reinforce the point that the theoretical context in which the experiment is conducted plays a key role in understanding its possible implications.

was not conclusive until the direct detection of interactions that did not conserve parity Wu et al. (1957). Up to and including the discovery of the Higgs boson, the norms of particle physics dictated that experimental discovery required direct detection, and this was usually accomplished by producing the sought after particles in an accelerator. Direct detection is still the ideal standard, but pragmatic concerns involving technological and cost limitations, as well as unresolved theoretical puzzles have made the prospects for further direct detection seem dim. Koberinski (2023a) and Koberinski (2025) note that standard interpretations of the Standard Model as an EFT indicated that some signs of new physics—whether direct or indirect—should be observed at the same energy scale needed to produce Higgs bosons. This has not happened, and there are few proposals replacing our current understanding of EFT that make concrete predictions for where new particles might be found. Given the pragmatic limitations of direct detection, it may take decades for a new accelerator to be built and put online. Even then, we have no guarantee that the next generation of accelerators will be capable of directly detecting new physics. It is therefore highly desirable to find other ways to explore the space of possibilities beyond the Standard Model.

3 From discrepancies to discoveries

When searching for signals of new, unknown physics beyond the Standard Model, the experimental goal is to generate positive results against the null hypothesis that the physics of the experiment is accurately captured by the Standard Model. A positive result is a discrepancy between the Standard Model prediction and the experimental result. In this context, the design of these searches all presuppose the Standard Model to some extent. Therefore the relevant dimension of theory-mediation is the extent to which assumptions about *beyond Standard Model physics* play an essential role in the design and setup of a given experiment. To see how direct versus indirect searches lead to discovery, we must analyze the process of inference from a discrepancy with a Standard Model prediction to a discovery claim, and how the underlying physical source can be conclusively established.

Start with the paradigmatic direct detection: producing a resonance at a given energy in a differential scattering cross-section. The resonance would appear as an excess of events at that energy, compared to the expectations from the Standard Model. This discrepancy needs to meet a threshold of statistical significance—5σ in particle physics—as a necessary condition for a discovery claim.[8] The inference is from a resonance at a given energy to the existence of a new particle whose mass is directly determined from that resonance. A discovery like this is guided by relatively theory-agnostic expectations about how scattering cross-sections are influenced by the production of particles. This inference does not depend on any particular candidate theory going beyond the Standard Model, though the experimental design and targeted search area might still be mediated by specific theories. This is the ideal for direct detection, and has been the driving force for building increasingly more powerful accelerators. The inference from an excess of events to the existence of a new particle is very direct, regardless of whether the experimental design was guided by a particular theory of beyond Standard Model physics or not. Thus, the advantage of direct detections like this is the ability to infer directly from a discrepancy with the Standard Model prediction to a physical source. In principle, direct detection need not be heavily theory-mediated but in practice the possible regimes in which new physics could occur are too vast to search unconstrained.

How does the inference pattern from discrepancy to discovery work for indirect detection? As I have defined the term, an indirect detection does not allow for a direct inference to some source, so the structure must be mediated in some way. In general, indirect detections can only allow one to infer the existence of a discrepancy for the background theory. Persistent discrepancies signal an anomaly[9]: our current theory is inadequate in that domain, and therefore that new physics is needed. Some amount of theoretical context is needed to explain the presence and magnitude of the anomaly. Anomalies play an important role in the

[8]Reporting a tension of $n\sigma$ means that the discrepancy between the measured and predicted value differ by n times the calculated standard deviation of the difference, on the assumption that the calculated uncertainties for each quantity are normally distributed. While 5σ is the standard for discovery in particle physics, Lyons (2013) argues that this threshold should be more nuanced. In particular, when it comes to claiming a discovery of new, beyond Standard Model physics, he argues that a higher threshold should be posited. The more unexpected the result, and the more it clashes with currently accepted physics principles, the higher the statistical threshold. However, standard particle physics practice still conforms to the 5σ standard.

[9]The term ‘anomaly’ might come burdened with Kuhnian undertones, but here I mean only that an anomaly is a persistent discrepancy between experimental results and theoretical prediction.

process of theory construction, since they signal physical sources not accounted for in the current theory. They provide touchstones that any successor theory should aim to capture if it is to succeed the current theory. However, anomalies do this indirectly; there is no algorithm or recipe allowing the inference from an anomaly to some new theory explaining it away. An anomaly is simply a persistent difference in the value of some quantity. In order to resolve the anomaly, a new theory is needed, within which there is now agreement between theory and our evidence. Anomalies get attributed to some underlying physical source only in the context of the new theory. For example, the anomalous motion of Mercury's perihelion was an anomaly for Newtonian gravity. The anomaly was explained from the point of view of general relativity as being due to spacetime curvature effects in strong gravitational fields. The source of the anomaly was the curvature of spacetime in the theoretical context of general relativity. Therefore one needs more than the detection of an anomaly to establish some new physical source.

But how do we first establish the presence of an anomaly? One must start with a discrepancy between expectations from theory and experimental results, but on its own this is insufficient. A discrepancy is simply a statistically significant difference in the value of a quantity determined by different means. In the case of experimental discovery, that discrepancy would be between a theoretical prediction and best value for a measured quantity. Given a significant discrepancy, work must be done to attempt to resolve it from *within* the current theory. Perhaps approximations made in the theoretical calculation can be refined, computational techniques improved, or perhaps systematic experimental errors were underestimated.[10] The former two points are of particular concern for quantities predicted by the Standard Model, where perturbative methods and other approximation techniques are central to the very definition of the quantum field theories that comprise it (J. D. Fraser 2020; Miller 2021; Dougherty 2023). It is rather common for advances in precision of theoretical predictions to come from better approximation techniques or better numerical methods for solving complex equations. When a discrepancy can be resolved in this way, it does not rise to the level of anomaly, but instead plays a central role in the elaboration and development of secure, empirically grounded facts based on the current theory (cf. Smith 2014; Koberinski and Smeenk 2020; Koberinski, Falck, and Smeenk 2023).

It is only when a discrepancy persists despite attempted resolution that it rises to the level of anomaly, and therefore becomes a candidate for an indirect discovery. While searches for new physics are heavily mediated by the Standard Model, many indirect searches remain largely independent of any particular successor theory that one might propose. Given the strong degree of model-dependence required to perform experiments, conduct data analysis, and interpret the results of a particle physics experiment, heavy dependence on some set of theories is to be expected. But since the search is for physics that *does not* conform to the Standard Model, we can treat the Standard Model itself as a background theory that is only indirectly guiding the search. Anomaly detection therefore has the potential to be theory-agnostic in the relevant sense, with the advantage that one is only searching for quantities that are not accurately predicted by the Standard Model.[11] Whether a given search is theory-agnostic or not will depend on the details of its design and setup. The indirect search strategies I describe in the next section are relatively theory-agnostic, with precision measurement being more agnostic than the machine learning methods.

The disadvantage to theory-agnostic searches is that, by remaining agnostic as to the features of a successor theory, the search is highly unconstrained. Further, the inference from a given anomaly to a new theory—or even to a new phenomenon—is far from obvious. Finding an observable quantity that differs in value from the Standard Model's prediction gives almost no insight into what new physics is responsible. Koberinski (2023a) has argued that treating the Standard Model as an EFT can provide a theory-agnostic scaffolding for organizing indirect searches; the experimental side of this approach will be discussed in Sec. 4.1. A second strategy for structuring searches is to use machine learning to flag anomalous events, this will be the subject of Sec. 4.2. While a direct detection of a new particle would be the strongest evidence for beyond Standard Model physics, indirect theory-agnostic searches can provide the first steps of evidence of where beyond Standard Model physics is needed. In Sec. 5 I further argue that indirect searches provide

[10]All of these have been found to be sources of a seemingly promising discrepancy between empirical findings and the Standard Model at some point in the last 15 years.

[11]These are further along the theory-agnostic dimension than searches that presuppose a particular family or class of beyond Standard Model theories. Examples of this latter form include searches for proton decay—a signature of a large class of grand unified theories (Bajc et al. 2016)—or searches using so-called simplified models derived from supersymmetric extensions of the Standard Model (McCoy and Massimi 2018). However, even in indirect detections, one might set expectations for where to look on the basis of beyond Standard Model theories.

complementary epistemic benefits to standard high-energy direct detection efforts.

4 Indirect methods of discovery

Experimental practice in particle physics is highly complex, due to the difficulties in reconstructing short-lived subatomic particle phenomena from decay products detected by highly sophisticated networks of particle detectors. Production of higher-mass particles requires larger accelerators that can steer and accelerate stable particles (electrons, protons, etc.) to speeds very near the speed of light. Resulting collisions occur sporadically, producing unstable particles that rapidly decay into lighter products until only stable particles remain. Collision products can only be predicted probabilistically from the incoming particles. The challenge in such collision experiments is to reconstruct an event—a single collision—from potentially far-flung decay products whose properties can only be constrained by the detector outputs (Karaca 2017; Beauchemin 2017). High-energy collision experiments are expensive and can only be conducted where large accelerators exist. Building new accelerators is even more prohibitively expensive. While high-energy collision experiments must still play a central role in the experimental practice of particle physics moving forward, other experimental practices should be pursued in parallel. In particular, comparatively low-energy "precision" experiments provide a cheaper option in the search for anomalies. Some of these can be conducted in smaller labs using cheaper equipment, and carried out at a faster rate. A second approach is to re-analyze data produced in previous runs of high-energy accelerators like the LHC, to look for events whose signatures are an ill-fit with expectations from the Standard Model. Given that the Standard Model only makes probabilistic predictions regarding the outcome of a scattering event, flagging individual events is somewhat tricky. One method that I will focus on here uses unsupervised machine learning algorithms—specifically autoencoders—trained on data known to fit well with Standard Model predictions to analyze and reconstruct a broader range of past events. Poorly reconstructed events are events that the autoencoder takes to be sufficiently different from the training data, and are therefore candidates for Standard Model anomalies.

4.1 Precision measurement

There are two types of precision measurement one can use in particle physics for the purposes of discovery, and the distinction can be made in terms of a qualitative recourse to EFT methods. As argued in Koberinski (2023a), the EFT framework serves as a phenomenological formalism to structure the search for new physics. Within the EFT framework one expects indirect signals of new physics in the form of higher-mass dimension operators factoring into observables. The resulting contributions to observables are suppressed by the upper limit cutoff scale Λ for the EFT, in the form $\mathcal{A}_n \propto \left(\frac{\mu}{\Lambda}\right)^n$, where μ is the energy scale at which the experiment takes place, n a natural number, and the full observable will be a summed expansion of the form $\mathcal{A} = \sum_n \mathcal{A}_n$ (more on this below). Under a reading that takes the cutoff scale to have physical significance—it is the scale at which the EFT breaks down and new physics takes over—there are only two ways to increase sensitivity to these higher-order terms. First, at a fixed energy scale μ, increase the experimental precision ϵ, such that $\epsilon < \left(\frac{\mu}{\Lambda}\right)^n$. Second, for fixed experimental precision, increase the energy for the experiment such that $\left(\frac{\mu}{\Lambda}\right)^n$ gets larger. However, increased energies often result in a decrease in experimental precision[12]: what is needed is a targeted increase in energy while maintaining as much precision as possible: $\epsilon(\mu) < \left(\frac{\mu}{\Lambda}\right)^n$.

The first type of precision increase consists of cheaper, low-energy high-precision measurements. Many of these are experiments that don't require the use of a full-scale particle accelerator, and may therefore provide a promising alternative or supplement to the usual experimental practice in particle physics. However, because of this, they are also limited in scope. Certain particle physics phenomena are only reliably produced in high-energy accelerators, and many of the most promising regions of the Standard Model parameter space are concerned with these higher-energy effects. Nevertheless, for some phenomena, low-energy experiments are able to reach an impressive degree of precision. As discussed in Koberinski and Smeenk (2020), precision measurements of the electron's anomalous magnetic moment using a Penning trap have become precise

[12]However, this need not always be true. While experimental uncertainty is often more difficult to control in experiments done at higher-energies, one can still increase precision in high-energy experiments (as discussed below). On the theory side, precision is based on the ability to calculate approximations well; this is often accomplished in regimes where perturbative methods work well. For quantum chromodynamics, for example, perturbative methods work best in the high-energy limit. I am grateful to an anonymous reviewer for pressing me on this point.

enough to report the value to 11 significant digits. Importantly, calculating the magnetic moment to this level of precision requires one to go beyond "pure" quantum electrodynamics, and to include indirect dynamical effects from the strong and weak sectors of the Standard Model. In a counterfactual scenario where the rest of the Standard Model was not discovered before this level of precision was reached, precision tests would have revealed a persistent discrepancy between pure quantum electrodynamics and observation, resulting in an anomaly for the theory. This shows that high-precision measurements of low-energy quantities predicted by the Standard Model are able to find anomalies, even in very low-energy phenomena.

Part of the challenge in calculating the values of observables like the anomalous magnetic moment of the electron is due to the mathematical form of interacting QFTs making up the Standard Model. Observables are generally calculated as expectation values of products of field variables at various spacetime points (or, equivalently, at different four-momenta), and these are in turn calculated from the theory's generating functional $\mathcal{Z}$. Importantly, the measure for the path integral in $\mathcal{Z}$ is undefined on a continuum Minkowski spacetime. Even when one ignores this point (for example, by switching to a Euclidean space instead), $\mathcal{Z}$ has no known exact solution for realistic interacting theories. For theories that have interaction terms that are small—quantum electrodynamics is a prime example—one can calculate observables by treating the interaction term as a small perturbation to the non-interacting theory. The non-interacting theory has exact solutions, so one takes a truncated perturbative expansion in powers of the coupling constants to be an approximation to the true interacting solution.[13] For theories with stronger interaction terms (like quantum chromodynamics at low energies), one makes the idealization that the field theory is defined on a lattice, and numerical simulations are used to generate predictions for observables. Both cases necessarily involve using approximation methods in the absence of some underlying set of closed-form equations that we are knowingly trying to approximate.

For quantities like the anomalous magnetic moment of the electron, one can therefore make more precise predictions by increasing the sophistication of the approximation methods used. This includes calculating to higher orders in perturbation theory, including sub-dominant interactions (like weak and strong nuclear force effects), and developing better numerical approximation techniques for solving the complicated integrals that arise in a given calculation (cf. Aoyama et al. 2012; Koberinski and Smeenk 2020; Huggett and Schneider 2025). On the experimental side, precision is increased by designing apparatus that allow for greater experimental control, reducing statistical uncertainty, and reducing the magnitude of sources of systematic uncertainty. By increasing precision experimentally *and* theoretically, the Standard Model is being put to an increasingly more stringent test, increasing the odds of finding a discrepancy.

While no discrepancies with the full Standard Model have been found for the electron, a comparable discrepancy in the muon's anomalous magnetic moment exists. While the experimental conditions needed to isolate the muon and probe its magnetic moment involve far higher energies than similar probes for the electron, this still qualifies as a relatively low-energy precision measurement, compared to the energies reached in state-of-the-art accelerators. While these higher-energies are accompanied by a decrease in precision for measurement—a composite measured value of the muon's anomalous magnetic moment is only measured to 9 significant digits—the impact of new physics on muon quantities versus electron quantities is expected within the EFT framework to scale as $(m_\mu/m_e)^2$, with $m_\mu \sim 200 m_e$. Thus, despite the decrease in absolute precision, there is good reason to expect the muon to be more sensitive to new physics. The discrepancy was increased by recent measurements by Albahri et al. (2021) to a 4.2σ tension between theoretical prediction and measured value, just below the particle physics 5σ standard threshold for claiming a discovery.

However, the recently discovered discrepancy has not yet become an anomaly. Just before the Fermilab announcement of the 2021 results, Borsanyi et al. (2021) recalculated a component of the muon's anomalous magnetic moment due to virtual hadron polarization using improved lattice quantum chromodynamics techniques. This calculation led to a correction to the predicted magnetic moment that reduced the difference between measured and predicted value, therefore reducing the tension (Muon $g-2$ Collaboration, Aliberti, et al. 2025). The final round of results has since been released, with FermiLab's new value for the anomalous

[13] I am glossing over several interpretive issues that complicate the current discussion. For example, Haag's theorem states that we are not licensed to treat the full interacting theory as a perturbation to the free theory (Earman and D. Fraser 2006; Miller 2018; Koberinski 2023b). The perturbative expansion also has zero radius of convergence, and individual terms diverge as well (J. D. Fraser 2020). All of these points only make the approximate nature of QFT even stronger, however, so glossing over them here is simply for ease of presentation.

magnetic moment of the muon being

$$a_{\mu,exp} = [116,592,070.5 \pm 11.4 \times (stat.) \pm 9.1(syst.)] \times 10^{-11}, \tag{1}$$

with an overall precision of 127 parts per billion (Muon $g-2$ Collaboration, Aguillard, et al. 2025). With this new result, and the updated theoretical prediction, the tension has entirely disappeared: $a_{mu,exp} - a_{\mu,SM} = 38 \pm 63 \times 10^{-11}$. The discrepancy was still important, even if it did not result in a persistent anomaly. Discrepancies bring to light tensions between incomplete theoretical models and experimental results. It will often be the case that a discrepancy is resolved in a consistent manner by updating the approximations and idealizations that go into the construction of the theoretical model. In this case, the tension led to improvements in calculating hadronic effects with lattice quantum chromodynamics, and to a more precise measurement of $a_{\mu,exp}$. When this is done, we learn more about the quantitative implications of our best theories, and gain tighter constraints on the allowed deviation introduced by theories of beyond Standard Model physics. When it can't be done, we are left with an anomaly.

The second type of precision experiment is one that increases the precision on relatively high-energy measurement results. While absolute precision will ultimately be lower, there are clusters of phenomena in the Standard Model where even modest gains in precision could serve to highlight discrepancies, in a similar manner to the relative sensitivity of the muon and electron to new physics. Now that the LHC has reached peak centre of mass energy, it is shifting into a new phase of high-luminosity collisions, with construction already underway, and completion expected by 2029. The luminosity in a collision experiment is proportional to the collision rate at beam crossing; higher luminosity therefore means more collisions per unit time at a given energy. This is helpful for improving precision on high-energy observables by reducing the statistical uncertainty introduced when the number of events is low. Statistical uncertainty is the uncertainty associated with inferring from a sample of a certain size to the category of events as a whole. For larger sample sizes, this uncertainty is smaller, since one has more confidence that one is dealing with an appropriately representative sample. For rare, higher-energy events, statistical uncertainty can be a major source of overall uncertainty on experimental results.

Before the high-luminosity LHC is up and running, targeted collision experiments can decrease statistical uncertainty in more local ways. For example, if one is studying observables related to weak bosons, statistical uncertainty is reduced if the experiment produces a greater number of events containing the relevant weak bosons, from which the magnitude of the observable can be inferred. By tuning the centre of mass energy to be near the energy required for the most probable pathways to create weak bosons, one increases the likelihood that more of these particles will be produced. Specially designed triggers can help increase sensitivity to desired events, while minimizing the chances of recording irrelevant events. For some observables, this reduction in statistical uncertainty could be sufficient to find persistent discrepancies with Standard Model predictions. Some of the more promising sectors to search are related to the weak (W and Z) bosons and b quark, though here again recent results have served to decrease the tension with many Standard Model predictions (Isidori 2023; Gibney 2024). While the means for increasing precision differ, the inference pattern for both types of precision tests is the same. When the discrepancies are not found to persist, increased precision instead places stricter bounds on possible disagreement between candidate theories of beyond Standard Model physics and the Standard Model predictions. When a discrepancy persists to the point of becoming an anomaly, we find a known set of phenomena for which the Standard Model alone is inadequate.

4.2 Using machine learning to discover new physics

While machine learning has been prevalent in particle physics for decades, the rapid growth of computing power and the ensuing development of deep neural networks has greatly expanded the possibilities for its novel use. Boge and Regt (2025), in particular, have discussed some of the recent advances of machine learning in particle physics, arguing that artificial intelligence may be able to make meaningful scientific discoveries in a theory-neutral manner. For my purposes, I will set aside the question of intelligibility of machine learning discoveries, and focus instead on their use as a tool in the search for discrepancies in particle physics data. It is only once such discrepancies are flagged that questions of robustness arise; I will discuss this possibility at the end of the section. The further goal of autonomous discovery via machine learning is aspirational, and outside the scope of the current argument entirely. What is novel for the context of experimental discovery is that certain machine learning algorithms can re-examine old data without the same expectations of researchers. Further, these algorithms have the ability to flag *individual events* as potentially discrepant, rather than aggregated statistical results. Given that particle physics predictions are probabilistic, this implies that the use of machine learning is only as a first step in flagging events of interest; further investigation is required to even arrive at a discrepancy, let alone an anomaly.

For much of the history of particle physics, machine learning algorithms were used as computational tools, to help produce theoretical estimates of parameters, classification of events, or to aid in statistical analyses of data.[14] They have formed pieces of larger networks of methods all designed to take readings from the electronic sensors in a detector and reconstruct the events that produce the detected decay products. One distinctive feature of experimental particle physics is the heavy theory-mediation in the analysis and reconstruction of events. Another is the overwhelming amount of raw data produced in collision experiments, with limited ability to store and analyze it all. Machine learning has played an important part in managing both of these challenges. In particular, machine learning has been used to more efficiently develop filters that select only a small subset of candidate events for storage and reconstruction, and to do so without explicitly simulating the events using the Standard Model. However, the advent of more complex deep learning algorithms, especially deep neural networks, have allowed for novel applications in data analysis.

One of the strengths of deep neural networks is the ability to search large swaths of data in order to recognize patterns that might be particularly challenging for human scientists to notice. Unsupervised machine learning is especially helpful in this context for discovering novel patterns or features of the data, since the algorithm is not prepared with a set of preconceived data tags already thought to be important. This way, the algorithm can pick out its own features as relevant, and may recognize novel features of the data that have thus far been overlooked. However, philosophers have been quick to point out several epistemic drawbacks to machine learning as a tool for discovery in science. Of particular concern are the problems of opacity and bias to training data (Sullivan 2022; Kauffmann et al. 2024). The problem of opacity has to do with the complexity of machine learning algorithms. While they learn to perform tasks on training data, and we can validate their success on other data sets, it is rarely the case that we know *how* algorithms are performing the task, or what specific patterns in the data they are latching onto (Boge 2022). In general, the larger and more complex the algorithm, the more opaque it is to human understanding. The desire for *intelligibility* of machine learning algorithms is a desire to reduce this opacity. If we don't understand how the machine learning algorithm performs its task, then we have less understanding of why and how it works, undermining confidence that it is using a reliable method that will generalize to novel use cases. The problem of bias to training data is a related problem, leading the machine learning algorithm to abstract the wrong features to generalize based on biases in the training data. For example, if an image classification algorithm is supposed to distinguish house-cats from lions, it might learn instead to classify cat versus lion based on the image background, unless the training data is sufficiently varied to account for this. The two problems are related in that they undermine our confidence in the ability of a machine learning algorithm to succeed in its task in new contexts.[15]

What is interesting about current state-of-the-art uses of machine learning for anomaly detection in particle physics is that these two major issues are largely neutralized, at least to a first approximation. Current searches use a particular type of algorithm architecture called an autoencoder, and intentionally bias the training data to fit expectations of the Standard Model. The structure of the neural network for a typical autoencoder is to take a high-dimensional input layer to "read" the data, greatly reduce the dimensionality in an interior coding layer, and then to decode back out to a reconstruction of the input data. The goal is to have the encoding portion find useful features of the data to abstract into the coding layer, such that the subsequent decoding leads to a minimal loss as defined by a loss function comparing input to output. Once an autoencoder has been trained, it ideally has a good, generalizable set of features to abstract from data similar in kind to the training data. The training ensures that these are sufficiently useful to successfully reconstruct an image from the encoding. When using autoencoders for discovery of beyond Standard Model physics, the training data is a set of events that are highly typical according to the Standard Model. Training is conducted to minimize the loss between input and reconstructed output. Then,

[14] A great resource highlighting the various uses of machine learning in the High-Energy Physics living review on machine learning (Feickert and Nachman 2024).

[15] There are further important questions concerning the ethical implications of the use of machine learning, given the above epistemic concerns (e.g. Dubber, Pasquale, and Das 2020). I do not address these issues here, as the ethical implications in the context of particle physics are minimal, compared to uses in profiling, human data collection, and ownership rights on training data.

datasets consisting of other similarly coded events are fed into the trained autoencoder. Events that it fails to reconstruct well do not share the abstracted features of the sample data on which it was trained; these are flagged as candidates for events with beyond Standard Model signatures. These are the events we are interested in. What is typically recognized as a problem in standard accounts of machine learning discovery is repurposed to become the essential feature used to find candidate discrepancies; the problem of bias to training data is acknowledged and exploited as the means of discovery.

The problem of opacity is also largely neutralized in this context. Opacity is useful in that the autoencoder provides an analysis of events distinct from what could be done via human-led statistical analysis. Since the goal is simply to use the autoencoder as a first step to flag events for further investigation, it is less important to understand its inner workings. In this context it is simply a tool whose outputs are analyzed by human scientists. Human analysis of data requires the use of statistical methods to analyze high numbers of events and abstract features. What the autoencoder can do is flag individual events in the dataset; human scientists must then build a more robust case for the presence of a discrepancy via statistical means.

These autoencoder searches help re-scrutinize data that has already been produced. Once the autoencoder is trained, it is significantly cheaper to re-analyze old events rather than produce new data by running a particle accelerator. When this re-analysis is conducted via unsupervised machine learning, it minimizes the impact of physicists' preconceived expectations on the search, and allows for a more fine-grained analysis of potential anomalies. Previous work has tested anomaly detection by training an autoencoder on jet images containing low-mass quarks, and used top-quark jets as a test for anomalies, with mixed results (Oliveira et al. 2016; Finke et al. 2021). More recently, the ATLAS Collaboration (2024) trained an autoencoder on 1% of the LHC Run2 data, which is taken to be representative of the full dataset, and used this to search for events that may hold signatures of new physics. They limited the search to events containing at least one electron or muon at centre of mass collision energy of 13 TeV, and looked for anomalies in resulting jet production, electron, muon, or photon distribution. In particular, the input data was formatted as a rapidity mass matrix, normalized by the centre of mass energy of the collision event. This matrix encodes several important variables for a given event, including missing transverse energy, transverse masses, transverse energies, two-particle invariant masses, and two-particle invariant rapidities for each of the possible particle combinations. These are variables taken to be salient for finding Standard Model anomalies. The 36×36 entry matrices can be visualized as heat map images each containing 1287 inputs (the entries are not all independent, so 9 inputs are reduced) all normalized to be between 0 and 1. Then the autoencoder is trained to reproduce these rapidity mass matrices in analogy to a QR code image, with input and output layers containing 1287 neurons, and the innermost latent layer containing only 200. While the autoencoder was successfully trained, and found to be able to detect anomalies put in by hand, it did not detect any statistically significant discrepancies on the remaining Run2 data. However, the search was found to increase sensitivity to new physics by allowing for a higher degree of precision in the resulting statistical analysis. This allows for stricter bounds on the possible range of parameter space in which new physics might be found, in a similar manner to negative results of precision measurement.

There are several modelling choices that go into the training and analysis of autoencoder anomaly detection. The choice of input data, including the events to focus on, the variables to select, and the formatting of the input all play an important role in the overall function of the autoencoder. Next, the subsequent analysis of events with high-loss score is carried out independently of the running of the autoencoder. The autoencoder is being used as a *tool* in a theory-agnostic search for signatures of beyond Standard Model physics; it is not detecting anomalies or making discoveries on its own. In this context we are thus far from the question of whether neural networks are capable of making autonomous scientific discoveries, and grounded in more practical questions regarding the efficacy and level of independence from modelling assumptions of such searches. Start with the latter; anomaly detection using an autoencoder is clearly not completely free of modelling assumptions. What is salient in the search for new physics, however, is not that this method be completely assumption-free, but that we minimize dependence on assumptions relating to candidate classes of beyond Standard Model theories. Further, if the modelling assumptions used here are largely independent of assumptions used in searches elsewhere (via precision measurement or high-energy collisions), then one can be confident that the overall search effort is more exhaustive, as several independent methods are used to conduct a search.

Regarding the modelling assumptions, most of these are at the level of creating an appropriate data model, rather than being dependent on any particular theories of beyond Standard Model physics. Choices about which variables to encode, and how to encode them, are at most indirectly informed by expectations of what might be salient for new physics. But even here, the choices are not informed by any particular model or class of models, but by the general EFT framework that encompasses the Standard Model. In this particular instance, the only theory-dependence comes in during the selection of particular anomaly regions to focus on, and on the sample anomalies used to test the trained model. A selection of possible beyond Standard Model theories were analyzed by the ATLAS collaboration in order to select certain loss and intensity thresholds, but these simply set the external constraints on the autoencoder search. Expectations regarding where to look were largely independent of candidate theories, as these expectations were set based on pragmatic considerations of data coding. While the ATLAS analysis has some weak dependence on the particular class of theories chosen, and is not completely theory-agnostic, the mediation from beyond Standard Model theories is minimal. In principle, one could appeal only to the EFT framework to select parameter regions, minimizing the theory-dependence further.

Are the modelling assumptions used here independent of those used in other searches? As described above, precision measurements looking for anomalies are largely informed by the current precision limits on experimental determinations of parameter values and those of the best theoretical predictions from the Standard Model. Given a promising parameter to measure, the modelling assumptions there depend entirely on the Standard Model and particular design choices for increasing precision on the experimental side. While there may be some overlap in the choice of variables/parameters to focus on, especially if EFT is used in both cases, the autoencoder is not pre-trained with any of the expectations used in selecting precision experiments. What the autoencoder finds to be anomalous should therefore be independent from any assumptions guiding precision tests. A similar argument can be made for high-energy collision experiments. While there may be some overlap on using minimal models extended beyond the Standard Model to choose parameter regions to focus on, collision experiments aim for direct detection of new forces or particles, while the anomalies an autoencoder might find are indirect signals of new physics. The major limitation for autoencoder searches is the fact that they can only analyze subsets of data already produced. Whatever assumptions have gone into previous LHC data selection and storage are therefore carried over to anomaly detection in this context, though the analysis of existing data proceeds independently.

Anomaly detection using autoencoders is not as theory-agnostic with respect to physics beyond the Standard Model as precision measurement, but is agnostic in several relevant senses. What of the efficacy of detection? At surface level, the previously described analysis failed to find any statistically significant discrepancies, so in that sense the autoencoder has failed its task. But the failure to detect a significant discrepancy does not mean that the search strategy failed; it could be the case that the analyzed data actually contained no signal of any beyond Standard Model physics. Another standard for assessment could be that the autoencoder search had a greatly increased sensitivity to new physics in several of the target search areas. In this sense, the search was successful in increasing the ability to detect small effects over background, and the negative result therefore more strictly constrains possible deviations from the Standard Model.

In general, there will be no context-independent way to set or assess success conditions of autoencoder searches. This is because of second-order problems of opacity and bias to training data that come back into the picture. It would be much easier to assess the efficacy of a method by some general criterion if the method was intelligible to humans. But by construction, autoencoders learn from training data in a way that is opaque to humans. This is both its advantage and a potential drawback; an advantage in providing "fresh eyes" to examine old data, but a drawback for intelligibility. The ATLAS Collaboration (2024) tested their trained autoencoder on validation datasets, to ensure it could successfully reconstruct the types of events it was trained to reconstruct, and on artificially produced anomalies, to gauge how much loss would typically occur for a range of possible beyond Standard Model signatures. While these benchmark assessments provide some indication that the autoencoder is working as intended, the tests are results-based, and only applicable during the training phase. Without an intelligible method we can abstract from the autoencoder, results-based measures might be the best we can hope for.

One worry related to opacity is a second-order instance of the bias to training data. While anomaly detection uses this bias as a tool for discovery, the opacity of the autoencoder means that we cannot be sure that the autoencoder is biased to the *right features* of the training data. We want the autoencoder to be biased to features of the data that are indicative of Standard Model physics, such that high loss implies beyond Standard Model signatures. But the autoencoder might find more narrow features that are indicative of particular nonessential features of event selection, or more broad features that also generalize to beyond Standard Model effects. For the former, picking a varied training set can help minimize the chance that the autoencoder has abstracted contingent features of the particular dataset. In this case, the Run2 training data was selected to span several years of data collection, where trigger selection procedures vary across these events. By varying the nonessential features of training data, one can minimize the likelihood that the autoencoder will "cheat" by abstracting these features, rather than features that arise as a consequence of Standard Model physics. However, there is no guarantee of success. In the case of jet anomalies, Finke et al. (2021) found that their autoencoder was biased to reproduce images as simple or simpler than the training data. Images that were more complex were reconstructed with high loss, and thus flagged as potentially anomalous. Since there is no guarantee that beyond Standard Model physics will manifest in more complex images of events (in the ATLAS case, more complex rapidity mass matrices), this bias indicates that the autoencoder might still pick up non-projectable features of the training data. This is simply the problem of bias to training data arising again.

One might worry that beyond Standard Model effects are present in the training data, such that the autoencoder abstracts features of beyond Standard Model physics, rather than just Standard Model features. The easiest way to avoid this issue is to train the autoencoder on simulated data, since one can be certain that the simulated data only incorporates known physics. However, this comes with severe drawbacks as well, since the training and test data will differ in significant ways that ought to be inessential to the task at hand. Luckily, there is some evidence that autoencoders and other machine learning algorithms can be trained on contaminated data and still abstract features independent of the noise. As an example Farina, Nakai, and Shih (2020) found that an autoencoder trained to reconstruct background succeeded even if the training data was contaminated with 10% signal. This means that the autoencoder was just as successful at constructing background and failing to reconstruct signal as when it was trained on pure background data. If the training data is chosen such that known features are well-modelled by the Standard Model, then this means that the deviations due to new physics would be small. In this context, some anomalous data in the training set should therefore be expected not to pose a problem for training the autoencoder, as long as the training set is largely in agreement with Standard Model expectations. However, one complicating factor is again the opacity of the autoencoder. The ability to learn background from signal in a set of events might not generalize to the ability to filter out systematic variation due to beyond Standard Model effects. In both cases, one cannot completely eliminate the possibility of unwanted bias, but one can take precautions to limit bias. For the purposes of flagging potentially anomalous events for further scrutiny, bias that increases the false positive rate is preferable to bias that increases the false negative rate. Since further investigation is necessary to find a discrepancy, these false positives could be ruled out during subsequent analysis.

To summarize, assessments of the efficacy of anomaly detection via autoencoder are largely results-driven, and amount to forms of verification and validation testing familiar from the world of simulation (Kleijnen 1995). While we cannot be certain of the efficacy of a trained autoencoder, it can still serve as a useful flagging tool that searches in ways independent of precision testing and high-energy collision experiments. The ability to repurpose old data, and to flag individual events are unique features of indirect discovery via autoencoders. Since the autoencoder abstracts features of the events in a way that is independent of our own theoretical presuppositions, it might be able to abstract patterns from the data that we have not, and therefore flag events as anomalous that we would not classify as such. Coupled to the ability to analyze individual events, there is some promise for qualitatively new discovery. When that fails, the analysis can at least typically place stricter bounds on the possibility of new physics in that sector. At this stage, however, autoencoders can at most point at hints of a discrepancy. Human analysis is needed even to establish a true anomaly, let alone lead to a discovery.

5 Epistemic advantages to searching high and low

The methods outlined in the previous section allow for novel avenues to experimental discovery of discrepancies between the Standard Model and the phenomena it purports to predict. From these, it is an indirect path to the discovery of some new physical source responsible for the discrepancy. It is only when the discrepancy persists despite further refinements that it is elevated to an anomaly. For discrepancies in precision measurement, attempted resolution involves designing new experiments to repeat the measurement with reduced uncertainty, and refining the theoretical model used to generate the prediction. In the case of machine learning, flagged events must first be aggregated and analyzed further—by humans, for now—to determine whether a significant discrepancy is even there. Once this step is completed, then the attempted resolution can occur.[16] While most discrepancies will be resolved at this stage, those that survive are candidates for new physics.

For both precision measurement and anomaly detection via machine learning, the EFT framework plays an important role in structuring inferences involving anomalies. With precision measurement, the EFT framework provides a structure going beyond the Standard Model, within which anomalies are linked to new, non-renormalizable terms in the EFT extension of the Standard Model. For machine learning, the EFT framework is used to help identify regions of parameter space that might contain interesting anomalies. This pre-selection work helps to guide the construction and training of machine learning algorithms. EFT could also be used to help strengthen the case for potential anomalies flagged by the algorithm. In both cases, the move from an anomaly to a discovery is highly indirect.

Of philosophical interest is the epistemic and practical complementarity of indirect searches to more standard direct searches. When deciding whether to conduct a given experiment, one must determine possible epistemic payoffs for the possible experimental outcomes. This is especially pressing in the context of particle physics, where the proposed experiments are typically resource intensive, in terms of both cost and time. Start with direct detection efforts. As described in the introduction, these are becoming highly resource intensive, requiring multi-billion dollar accelerators, decades to build, and thousands of scientists to work on.[17] If direct detection efforts yield a positive experimental result, the epistemic gain is high and relatively direct. We discover some new particle or force with certain characteristics as determined by the outcome of the search. The inference is even stronger if the detection is in close alignment with expectations from some candidate beyond Standard Model theory. But if the result is negative, the epistemic gain is minimal: at best, we can rule out the candidate theory that predicted new physics in that region of parameter space. Even if the search is relatively theory-agnostic, then we have only constrained a small region of parameter space.

By contrast, indirect searches lead to a much weaker epistemic gain if a positive result is found, especially when they are theory-agnostic. When one has a particular candidate theory of new physics, the inference from some new phenomenon to its downstream effects elsewhere is more straightforward; one can typically calculate it from the candidate theory. The inverse problem, however, is much more challenging. Inferring from a downstream effect to the specific form of new physics is far from trivial, especially in light of only one anomalous result. All we gain for sure is knowledge that the Standard Model is inadequate to capture the relevant physics. While weaker than direct detection, knowledge of an anomaly is still epistemically valuable. Anomalies are essential to the process of theory construction beyond the Standard Model as they serve as checks that the newly proposed theory makes correct predictions where the current theory is known to fail. Familiar from other areas of science, a creative leap is often needed to move from anomalies to a new theory or framework.

This doesn't mean that we need to have a specific candidate theory in mind in order to claim a discovery, for either indirect or direct searches. While the goal of experimental discovery is to lead to the development of new theories of physics beyond the Standard Model, new discoveries have often happened in the absence of a new theory to guide expectations; the situation in particle physics post-1975 is exceptional in that regard. Many persistent anomalies can be taken as discoveries of new phenomena; in particle physics, a prominent example is the discovery of parity nonconserving interactions in relation to β decays, in absence of a theory of weak interactions (Lee and Yang 1956; Wu et al. 1957; Koberinski 2019). Currently, the phenomenon of neutrino oscillation serves as a similar example of a discovery of a new phenomenon not accompanied by a theory on the basis of which to expect it (Bilenky, Giunti, and Grimus 1999).

For the types of anomalies that can be established from precision measurement or with the help of deep neural networks, the link to a particular physical source is far less direct than even the discovery of parity nonconservation. For example, a persistent discrepancy in the measured versus predicted values of the W

[16]Recall that my concern here is only with the use of machine learning algorithms to flag potential discrepancies, not with the further question of whether machine learning can make independent scientific discoveries.

[17]This characterization is simplifying things greatly. Accelerators like the LHC can be used for a variety of purposes beyond direct detection. But the contrast is with indirect methods described in this paper, where existing tools can be used without further investing in better accelerators.

boson's mass could rise to the level of anomaly, but the underlying physical source giving rise to this anomaly would be obscure. A similar point could be made for anomalies relating to collision cross-sections as well. This is at least partially due to the complex interactions that the Standard Model posits to give rise to quantities like these, as well as the inherently inexact nature of interacting QFTs. Further, there is a strong underdetermination regarding possible amendments to the Standard Model to resolve a single anomaly of this type. Even using the EFT framework to organize and order potential amendments (as outlined in Koberinski 2023a), there is a considerable amount of freedom in choosing which amendments to make to resolve one anomaly, and therefore considerable flexibility in what new phenomenon is responsible (Bechtle et al. 2022). A proper discovery will likely have to follow the pattern that Smith (2014) calls closing the loop. The newly hypothesized phenomenon responsible for the anomaly must have further consequences for other phenomena within the scope of the theory, and increased precision in measuring quantities related to these other phenomena should reveal the magnitude of the new effects.

What about a negative result in an indirect search? These are actually of high epistemic value, potentially even higher than a positive result. When paired with the EFT framework as a generalized framework for particle physics phenomena, negative results can place tight constraints on the space of beyond Standard Model theories in a much stronger sense than a failed direct detection. Since we are looking for *any* indirect signature of new physics, a negative result for an indirect search is highly constraining, in ways that a failure to detect a new particle is not. When an indirect search is more sensitive to possible discrepancies than previous experiments, a negative result implies that the Standard Model is predictively accurate to an even higher degree of precision than previously thought. This places constraints on the degree to which *any* successor theory can alter the predictions of the Standard Model, at least for the effect of interest. In typical scientific theories, tight constraints on possible predictions in one sector place global constraints on the form of the theory. Therefore, indirect searches have a much more balanced payoff across different possible outcomes. If a discrepancy is first observed, but fails to persist to the point of becoming an anomaly for the Standard Model, we also learn more about the physical details involved in producing the result. As Smith (2014) argues in the context of Newtonian gravity, discrepancies can provide a strong form of evidence in favour of the theory, if they are resolved in a consistent and principled way that allows us to discover more about the world. For example, the resolution of the muon's anomalous magnetic moment discrepancy discussed in Sec. 4.1 involved improvements in the measurement procedure, but also improvements in the calculation of vacuum polarization from virtual hadrons. The discrepancy was resolved within the Standard Model, but we learned more about the physics of the Standard Model along the way.

This epistemic complementarity of direct versus indirect searches is added to the practical complementarity as well: direct searches are highly resource intensive, while indirect searches are comparatively resource-frugal (Huggett and Schneider 2025). On both the epistemic and practical dimensions, direct detection efforts may be thought of as high-risk, high-reward, while indirect searches are low-to-moderate-risk, moderate-reward. A balanced approach to future searches prioritizing both would therefore provide a better expected epistemic payoff.[18]

6 Conclusions

Given the current state of high-energy physics, it seems prudent to widen the search for windows into new physics. I have provided a conceptual analysis of experimental search strategies in particle physics, characterizing them by how *directly* a positive result allows for an inference to specific new physics, and by how heavily *theory-mediated* they are. I argued that highly indirect detection methods are a promising avenue of pursuit due to the complementarity of the epistemic and practical benefits, as compared to high-energy direct detection efforts. My analysis focused on the epistemic dimensions of indirect searches, relying on the cases of lower-cost precision measurements of quantities relevant to particle physics and using machine learning algorithms—particularly autoencoders—for detecting anomalies in pre-existing collision data. Both of these approaches are appealing ways to broaden the search efforts beyond the standard direct-detection efforts, and both do so in a relatively theory-agnostic manner. Precision measurements occur at energies

[18] Perović (2011) has similarly argued for a diversity of experimental search procedures, though the relevant diversity is in automated versus non-automated detection processes. The latter, he argues, are under-pursued despite their epistemic advantage in searching for anomalies. My analysis here should be taken to add to the request for diverse search strategies in particle physics.

below the maximum accessible by the best accelerators at a given time, and aim to increase precision on experimentally determined values of observables by reducing the uncertainty associated with a given experimental setup. Unsupervised machine learning looks with "fresh eyes" at old data, and examines them for features that may not have been the target of their initial construction. When used simply as a tool to flag potential anomalies, many (but not all!) of the drawbacks to unsupervised machine learning are neutralized. Given the relative costs and ease of conducting indirect searches, there is little reason not to pursue these avenues.

However, discovery of new phenomena via indirect detection is, unsurprisingly, a more circuitous process than discovery via direct detection. I have articulated a relatively general account of the process of moving from discrepancies to anomalies, and from anomalies to possible discoveries. Indirect searches reveal discrepancies between the Standard Model and some sets of experimental results. These must be analyzed, and an attempt made at reconciliation before the discrepancies rise to the rank of anomalies. Finally, anomalies must be attributable to some underlying physical source before a discovery claim can be made. In the context of indirect searches in particle physics, the EFT framework serves as a helpful generalized scaffolding for organizing anomalies and interpreting them as possible phenomena to be discovered. Anomalies can be attributed to a set of possible new terms in an EFT extension of the Standard Model. One then searches for the impacts that such terms would have on other observables, and by a process of elimination one arrives at a set of new terms consistent with the empirical data. This program defines an iterative procedure for learning more precisely what phenomena—in the form of new EFT terms—are responsible for what new physical effects, and for constraining the magnitude of those contributions. Indirect searches are epistemically complementary to the standard direct search methods. While positive results in an indirect search are less informative of new physics, negative results provide important constraints on the space of possible alternatives to the Standard Model.

Although the process of discovery via indirect detection is therefore mediated by the theoretical superstructure of EFT, the search for discrepancies and anomalies is relatively theory-agnostic. By definition, discrepancies and anomalies make reference to the *current* theory (i.e., the Standard Model), but the relevant sense of theory-agnostic is with respect to candidate beyond Standard Model theories. Any anomalies established via precision measurement or machine learning would therefore be useful for establishing discoveries on the basis of a different theoretical framework than the EFT framework discussed here. The methods of indirect detect discussed here are relevant in other areas of science whose subjects are epistemically remote, like astrophysics and cosmology (Doboszewski and Elder 2025).

Note that the analysis here focuses on the *epistemic complementarity* or direct and indirect searches. While the current prospects for direct detection of new physics by, e.g., particle production at accelerators are dim, the epistemic value of a direct detection is significantly higher. It is epistemically beneficial that efforts for direct detection will continue in tandem with indirect methods going forward. This offers the advantage of looking through a larger space of possibilities. Hence, there is value in searching high *and* low (energies) for beyond Standard Model physics.

References

Albahri, T. et al. (2021). “Measurement of the anomalous precession frequency of the muon in the Fermilab Muon g- 2 Experiment”. In: *Physical Review D* 103.7, p. 072002.

Aoyama, Tatsumi et al. (2012). “Tenth-order QED contribution to the electron g- 2 and an improved value of the fine structure constant”. In: *Physical Review Letters* 109.11, p. 111807.

ATLAS Collaboration, The (2024). “Search for New Phenomena in Two-Body Invariant Mass Distributions Using Unsupervised Machine Learning for Anomaly Detection at $\sqrt{s} = 13$ TeV with the ATLAS Detector”. In: *Phys. Rev. Lett.* 132 (8), p. 081801. DOI: 10.1103/PhysRevLett.132.081801. URL: https://link.aps.org/doi/10.1103/PhysRevLett.132.081801.

Bajc, Borut et al. (2016). “Threshold corrections to dimension-six proton decay operators in non-minimal SUSY SU (5) GUTs”. In: *Nuclear Physics B* 910, pp. 1–22.

Beauchemin, Pierre-Hugues (2017). “Autopsy of measurements with the ATLAS detector at the LHC”. In: *Synthese* 194, pp. 275–312.

Beauchemin, Pierre-Hugues and Kent Staley (2024). “The epistemological significance of exploratory experimentation: A pragmatist model of how practices matter philosophically”. In: *European Journal for Philosophy of Science* 14.4, p. 59. DOI: 10.1007/s13194-024-00620-6.

Bechtle, Philip et al. (2022). “Bottoms Up: The standard model effective field theory from a model perspective”. In: *Studies in History and Philosophy of Science* 92, pp. 129–143. DOI: https://doi.org/10.1016/j.shpsa.2022.01.014.

Bilenky, Samoil M, C Giunti, and W Grimus (1999). “Phenomenology of neutrino oscillations”. In: *Progress in Particle and Nuclear Physics* 43, pp. 1–86.

Boge, Florian J. (2022). “Two Dimensions of Opacity and the Deep Learning Predicament”. In: *Minds and Machines* 32.1, pp. 43–75. DOI: 10.1007/s11023-021-09569-4.

Boge, Florian J. and Henk W. de Regt (2025). “Machine Learning Discoveries and Scientific Understanding in Particle Physics: Problems and Prospects”. In: *Philosophy of Science for Machine Learning: Core Issues and New Perspective*. Ed. by Juan Durán and Giorgia Pozzi. Synthese Library.

Borsanyi, Sz. et al. (2021). “Leading hadronic contribution to the muon magnetic moment from lattice QCD”. In: *Nature*, pp. 1–5.

Brown, Laurie, Max Dresden, and Lillian Hoddeson, eds. (1989). *Pions to Quarks: Particle Physics in the 1950s*. Cambridge University Press.

Cao, Tian Yu (2010). *From current algebra to quantum chromodynamics: A case for structural realism*. Cambridge University Press.

Cushing, James T (1990). *Theory construction and selection in modern physics: The S matrix*. Cambridge University Press.

Doboszewski, Juliusz and Jamee Elder (2025). *Black Boxes in Black Hole Imaging*.

Dougherty, John (2023). “Effective and Selective Realisms”. In: *The British Journal for the Philosophy of Science* 0.ja, null. DOI: 10.1086/724978. eprint: https://doi.org/10.1086/724978. URL: https://doi.org/10.1086/724978.

Dubber, Markus Dirk, Frank Pasquale, and Sunit Das (2020). *The Oxford handbook of ethics of AI*. Oxford Handbooks.

Earman, John and Doreen Fraser (2006). “Haag’s theorem and its implications for the foundations of quantum field theory”. In: *Erkenntnis* 64.3, pp. 305–344.

Elliott, Kevin C. (2007). “Varieties of Exploratory Experimentation in Nanotoxicology”. In: *History and Philosophy of the Life Sciences* 29.3, pp. 313–336. ISSN: 03919714, 17426316. (Visited on 03/28/2025).

Farina, Marco, Yuichiro Nakai, and David Shih (2020). “Searching for new physics with deep autoencoders”. In: *Phys. Rev. D* 101 (7), p. 075021. DOI: 10.1103/PhysRevD.101.075021. URL: https://link.aps.org/doi/10.1103/PhysRevD.101.075021.

Feickert, Matthew and Benjamin Nachman (2024). *A Living Review of Machine Learning for Particle Physics*. accessed 2024-10-28.

Finke, Thorben et al. (2021). “Autoencoders for unsupervised anomaly detection in high energy physics”. In: *Journal of High Energy Physics* 2021.6, pp. 1–32. DOI: doi.org/10.1007/JHEP06(2021)161.

Fraser, James D. (2020). "The Real Problem with Perturbative Quantum Field Theory". In: *The British Journal for the Philosophy of Science* 71.2, pp. 391–413. DOI: 10.1093/bjps/axx042. eprint: https://doi.org/10.1093/bjps/axx042. URL: https://doi.org/10.1093/bjps/axx042.

Ghosh, Saranya Samik et al. (2020). *General Model Independent Searches for Physics Beyond the Standard Model*. Springer Nature.

Gibney, Elizabeth (2024). "The standard model is not dead: ultra-precise particle measurement thrills physicists". In: *Nature* 633, pp. 745–746. DOI: 10.1038/d41586-024-03042-9.

Giudice, Gian Francesco (2017). "The dawn of the post-naturalness era". In: *arXiv preprint arXiv:1710.07663*.

Hoddeson, Lillian et al., eds. (1997). *The Rise of the Standard Model: A History of Particle Physics from 1964 to 1979*. Cambridge University Press.

Huggett, Nick and Mike D. Schneider (2025). *An apology for the simple and small*. URL: https://philsci-archive.pitt.edu/26204/.

Isidori, Gino (2023). "B anomalies: hopes, disillusions, and future prospects". In: *CERN EP Newsletter*. URL: https://ep-news.web.cern.ch/content/b-anomalies-hopes-disillusions-and-future-prospects.

Karaca, Koray (2013a). "The construction of the Higgs mechanism and the emergence of the electroweak theory". In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 44.1, pp. 1–16. ISSN: 1355-2198. DOI: https://doi.org/10.1016/j.shpsb.2012.05.003. URL: https://www.sciencedirect.com/science/article/pii/S135521981200041X.

Karaca, Koray (2013b). "The Strong and Weak Senses of Theory-Ladenness of Experimentation: Theory-Driven versus Exploratory Experiments in the History of High-Energy Particle Physics". In: *Science in Context* 26.1, pp. 93–136. DOI: 10.1017/S0269889712000300.

Karaca, Koray (2017). "A case study in experimental exploration: exploratory data selection at the Large Hadron Collider". In: *Synthese* 194.2, pp. 333–354.

Kauffmann, Jacob et al. (2024). "The Clever Hans Effect in Unsupervised Learning". In: *arXiv preprint arXiv:2408.08041*.

Kleijnen, Jack PC (1995). "Verification and validation of simulation models". In: *European journal of operational research* 82.1, pp. 145–162.

Koberinski, Adam (2019). "Parity violation in weak interactions: How experiment can shape a theoretical framework". In: *Studies in History and Philosophy of Modern Physics* 67, pp. 64–77.

Koberinski, Adam (2021). "Mathematical developments in the rise of Yang–Mills gauge theories". In: *Synthese* 198, pp. 3747–3777.

Koberinski, Adam (2023a). "Generalized frameworks: Structuring searches for new physics". In: *European Journal for Philosophy of Science* 13.1, p. 3.

Koberinski, Adam (2023b). "What good is Haag's no-go theorem? What axiomatic methods can teach us about particle physics". In: URL: https://philsci-archive.pitt.edu/22807/.

Koberinski, Adam (2025). "The Breakdown of Effective Field Theory in Particle Physics: Lessons for Understanding Intertheoretic Relations". In: *Philosophy of Science* 92, pp. 100–120. DOI: 10.1017/psa.2024.30.

Koberinski, Adam, Bridget Falck, and Chris Smeenk (2023). "Contemporary Philosophical Perspectives on the Cosmological Constant". In: *Universe* 9.3. ISSN: 2218-1997. DOI: 10.3390/universe9030134.

Koberinski, Adam and Chris Smeenk (2020). "Q.E.D., QED". In: *Studies in the History and Philosophy of Modern Physics* 71, pp. 1–13.

Lee, Tsung-Dao and Chen-Ning Yang (1956). "Question of parity conservation in weak interactions". In: *Physical Review* 104.1, pp. 254–258.

LHCb Collaboration, The et al. (2023). "Direct CP violation in charmless three-body decays of $B^{\pm}$ mesons". In: *Phys. Rev. D* 108 (1), p. 012008. DOI: 10.1103/PhysRevD.108.012008. URL: https://link.aps.org/doi/10.1103/PhysRevD.108.012008.

Lyons, Louis (2013). "Discovering the Significance of 5 sigma". In: *arXiv preprint arXiv:1310.1284*.

Mättig, Peter and Michael Stöltzner (2020). "Model landscapes and event signatures in elementary particle physics". In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 69, pp. 12–25. ISSN: 1355-2198. DOI: https://doi.org/10.1016/j.shpsb.2019.07.003. URL: https://www.sciencedirect.com/science/article/pii/S1355219818302119.

McCoy, CD and Michela Massimi (2018). "Simplified models: a different perspective on models as mediators". In: *European Journal for Philosophy of Science* 8.1, pp. 99–123. DOI: 10.1007/s13194-017-0178-0.

Miller, Michael E (2018). "Haag's theorem, apparent inconsistency, and the empirical adequacy of quantum field theory". In: *The British Journal for the Philosophy of Science*.

Miller, Michael E (2021). "Worldly imprecision". In: *Philosophical Studies* 178.9, pp. 2895–2911. DOI: 10.1007/s11098-020-01591-z.

Misiaszek, Marcin and Nicola Rossi (2024). "Direct Detection of Dark Matter: A Critical Review". In: *Symmetry* 16.2. ISSN: 2073-8994. DOI: 10.3390/sym16020201. URL: https://www.mdpi.com/2073-8994/16/2/201.

Muon $g-2$ Collaboration, The, D. P. Aguillard, et al. (2025). *Measurement of the Positive Muon Anomalous Magnetic Moment to 127 ppb*. arXiv: 2506.03069 [hep-ex]. URL: https://arxiv.org/abs/2506.03069.

Muon $g-2$ Collaboration, The, R. Aliberti, et al. (2025). *The anomalous magnetic moment of the muon in the Standard Model: an update*. arXiv: 2505.21476 [hep-ph]. URL: https://arxiv.org/abs/2505.21476.

Oliveira, Luke de et al. (2016). "Jet-images—deep learning edition". In: *Journal of High Energy Physics* 2016.7, pp. 1–32. DOI: 10.1007/JHEP07(2016)069.

Perović, Slobodan (2011). "Missing experimental challenges to the Standard Model of particle physics". In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 42.1, pp. 32–42. ISSN: 1355-2198. DOI: https://doi.org/10.1016/j.shpsb.2010.12.003. URL: https://www.sciencedirect.com/science/article/pii/S1355219810000821.

Ritson, Sophie and Kent Staley (2021). "How uncertainty can save measurement from circularity and holism". In: *Studies in History and Philosophy of Science Part A* 85, pp. 155–165. ISSN: 0039-3681. DOI: https://doi.org/10.1016/j.shpsa.2020.10.004. URL: https://www.sciencedirect.com/science/article/pii/S0039368120301886.

Skulberg, Emilie and Jamee Elder (2025). "What is a 'direct' image of a shadow?: A history and epistemology of 'directness' in black hole research". In: *Forthcoming in Centaurus Special Issue: Shaping a Multi-Messenger Universe: Historical Perspectives on the Changing Skyscape of Astronomical Inquiry*.

Smith, George E. (2014). "Closing the loop". In: *Newton and empiricism*. Ed. by Zvi Biener and Eric Schliesser. Oxford University Press, USA, pp. 262–352.

Steinle, Friedrich (1997). "Entering New Fields: Exploratory Uses of Experimentation". In: *Philosophy of Science* 64.S4, S65–S74. DOI: 10.1086/392587.

Sullivan, Emily (2022). "How Values Shape the Machine Learning Opacity Problem". In: *Scientific Understanding and Representation*. Routledge, pp. 306–322.

Wallace, David (2022). "Quantum gravity at low energies". In: *Studies in History and Philosophy of Science* 94, pp. 31–46. DOI: https://doi.org/10.1016/j.shpsa.2022.04.003.

Wu, Chien-Shiung et al. (1957). "Experimental test of parity conservation in beta decay". In: *Physical review* 105.4, p. 1413.