

Submitted to the Proceedings of the US Community Study
on the Future of Particle Physics (Snowmass 2021)

Machine Learning and Cosmology

Cora Dvorkin^{*1}, Siddharth Mishra-Sharma^{†1,2,3}, Brian Nord^{‡4,5,6}, V. Ashley Villar^{§7,8,9}, Camille Avestruz¹⁰, Keith Bechtol¹¹, Aleksandra Ćiprijanović⁴, Andrew J. Connolly¹², Lehman H. Garrison¹³, Gautham Narayan¹⁴, and Francisco Villaescusa-Navarro^{15,16}

¹Department of Physics, Harvard University, 17 Oxford Street, Cambridge, MA 02138, USA

²The NSF AI Institute for Artificial Intelligence and Fundamental Interactions

³Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁴Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

⁵Department of Astronomy and Astrophysics, University of Chicago, IL 60637, USA

⁶Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

⁷Department of Astronomy & Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA

⁸Institute for Computational & Data Sciences, The Pennsylvania State University, University Park, PA, USA

⁹Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

¹⁰Department of Physics; University of Michigan, Ann Arbor, MI 48109, USA

¹¹Physics Department, 2320 Chamberlin Hall, University of Wisconsin-Madison, 1150 University Avenue Madison, WI 53706-1390

¹²Department of Astronomy, University of Washington, Seattle, WA, 98195, USA

¹³Center for Computational Astrophysics, Flatiron Institute, Simons Foundation, 162 Fifth Ave., New York, NY 10010, USA

¹⁴University of Illinois at Urbana-Champaign Urbana, IL, USA

¹⁵Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY, 10010, USA

¹⁶Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton NJ 08544, USA

Abstract

Methods based on machine learning have recently made substantial inroads in many corners of cosmology. Through this process, new computational tools, new perspectives on data collection, model development, analysis, and discovery, as well as new communities and educational pathways have emerged. Despite rapid progress, substantial potential at the intersection of cosmology and machine learning remains untapped. In this white paper, we summarize current and ongoing developments relating to the application of machine learning within cosmology and provide a set of recommendations aimed at maximizing the scientific impact of these burgeoning tools over the coming decade through both technical development as well as the fostering of emerging communities.

^{*}cdvorkin@g.harvard.edu

[†]smsharma@mit.edu

[‡]nord@fnal.gov

[§]vav5084@psu.edu

1 Introduction

The interplay between models and observations is a cornerstone of the scientific method, aiming to inform which theoretical models are reflected in the observed data. Within cosmology, as both models and observations have substantially increased in complexity over time, the tools needed to enable a rigorous comparison have required updating as well. With an eye towards the next decade in cosmology, the vast data volumes to be delivered by ongoing and upcoming surveys, as well as the ever-expanding theoretical search-space, motivate a re-thinking of the statistical machinery used. In particular, we are now at a crucial juncture where we may be limited by the statistical and data-driven tools themselves rather than the quality or volume of the available data.

Methods based on artificial intelligence (AI) and machine learning (ML) have recently emerged as promising tools for cosmological applications, demonstrating the ability to overcome some of the computational bottlenecks associated with traditional statistical techniques. Machine learning is starting to see increased adoption across different sub-fields of and for various applications within cosmology. At the same time, the nascent and emergent nature of practical artificial intelligence motivates careful continued development and significant care when it comes to their application in the sciences, as well as cognizance of their potential for broader societal impact.

In this white paper, we provide an overview of some of the ways machine learning methods are becoming increasingly central to the way cosmological data is collected, analyzed, and interpreted. Along the way, we highlight our vision for necessary developments, framing these as recommendations—both technological as well as sociological—for the widespread safe and equitable adoption of machine learning methods within cosmology in the coming decade.

2 Examples of science cases

2.1 Cosmic Probes

Cosmology is the study of the content and evolution of our Universe. In the Λ CDM standard model of cosmology, the Universe has three primary components: dark energy (the source of the accelerated expansion of our Universe), dark matter (comprising the majority of the mass density in our Universe, whose presence thus far has been inferred through gravitational interactions), and ordinary visible matter. While the observational data has thus far remained largely consistent with the standard model of cosmology, fundamental physics questions remain unanswered. Unknowns include the particle nature of dark matter, the source of the accelerated expansion of the Universe, and the physics that seeded the first structures in the Universe.

There is an exciting prospect for many of the outstanding problems in cosmology to be solved, or substantial progress towards their solution to be made, using the statistical power offered by current and upcoming cosmological experiments. Maximizing the scientific return of forthcoming data will require methods that can extract as much information

as possible from observations while controlling for systematics associated with measurements as well as theoretical models.

In multiple subfields of cosmology, the pathway to cosmological constraints generally encompasses a source identification step, a summary measurement, and a comparison of the observation to theoretical models for parameter inference. In the rest of this section, we outline examples of cosmic probes where machine learning has already made a significant impact, with continued progress expected over the next decade as observations continue to increase in sensitivity. We note that the directions and works highlighted in this as well as subsequent sections are meant to be representative, rather than exhaustive.

Machine learning techniques will be crucial for detecting and classify cosmological sources, extract information from images, and optimize observing strategies. Examples of probes of cosmology include galaxy clustering, supernovae, strong and weak gravitational lensing, and the cosmic microwave background. Cosmology analyses that incorporate multiple probes have increased constraining power for two primary reasons. First, different probes span model space parameters in complementary ways, allowing for tighter constraints of cosmological parameters. Second, each observable is impacted by its own set of astrophysical processes and observational systematics. Machine learning methods can be used to fully realize the potential of multi-probe cosmology by decorrelating the effect of systematics and optimally combining information from multiple surveys [1].

Many scientific analyses that aim to extract cosmological information rely on astronomical catalogs. These catalogs are constructed from astronomical images, with their contents usually reflecting a “best-fit” model of the underlying image constituents (rather than a distribution over possible constituents). The associated information loss folds in information about uncertainty that cannot be retained in a downstream analysis. Machine learning has the potential to fully realize probabilistic cataloging [2, 3] which aims to detect and characterize constituents (e.g., stars and galaxies) from astronomical images in a fully probabilistic manner. These capabilities were recently demonstrated in the context of deblending crowded stellar fields [4] and can significantly enhanced the scientific output of cosmological surveys in the forthcoming decade.

The detection and characterization of strong gravitational lenses has emerged as a promising application of machine learning methods. An early application included a community-driven challenge on detecting strong lenses (see Ref. [5] for a summary), in which machine learning methods were shown to outperform traditional statistical techniques. Detection algorithms for strong lenses were subsequently applied to data from the Dark Energy Survey (DES) [6, 7] and the Dark Energy Spectroscopic Instrument (DESI) [8, 9]. Further work [10–12] proposed applying parameter inference and uncertainty quantification methods in order to characterize the properties of lensed sources and lensing galaxies. More recently, with an eye towards the large sample of gravitational lenses that will be imaged by forthcoming cosmological surveys like *Euclid* and LSST, there has been significant effort towards understanding how to utilize machine learning to optimally exploit this data towards source/lens characterization [13–16], Hubble constant inference [17], and characterization of dark matter substructure within the lensing galaxies [18–29] in a scalable manner.

Weak gravitational lensing is another area where machine learning has shown significant advantages over traditional methods. In particular, machine learning methods allow

for cosmological parameter inference via weak lensing measurements by leverage information beyond conventional one- and two-point statistics [30–36]. Additionally, machine learning methods have demonstrated the ability to outperform traditional estimators in mass mapping—reconstructing the weak lensing convergence maps from observed galaxy ellipticities [37–40].

Although the primary Cosmic Microwave Background (CMB) signal in the standard cosmological scenario can be statistically described as a Gaussian random field and efficiently analyzed with angular power spectrum estimators, machine learning methods have demonstrated superior performance on CMB analysis tasks including CMB lensing reconstruction [41], foreground separation [42, 43] and inference with polarization maps [44, 45]. Given the anticipated complexity of data, these techniques can significantly enhance the deliverable output of forthcoming surveys like Simons Observatory [46], CMB-S4 [47], and proposed high-resolution experiments like CMB-HD [48]. Analysis of CMB maps using machine learning techniques has also been explored for characterizing galaxy clusters via Sunyaev–Zel’dovich and CMB lensing signatures [49–52].

Machine learning has also been proposed as a way to confront the complexity of 21-cm neutral hydrogen data probing the epoch of reionization, with applications explored for parameter estimation [53–57] and signal extraction [58–60].

2.2 Time Domain & Multi-messenger Astrophysics

In the time domain, machine learning techniques have a key role to play in the era of wide-field surveys across the electromagnetic and gravitational wave spectrum. Of these surveys, the *Vera C. Rubin Observatory*’s Legacy Survey of Space and Time (LSST) will produce \sim 10 million alerts from time-domain phenomena every night—higher than can conceivably be inspected visually by the entire astronomical time-domain community in a lifetime. LSST will be joined by the next-generation Very Large Array (ngVLA), CMB-S4, the LIGO-VIRGO-Kagra Collaboration and numerous other data-intensive experiments over the next decade. Each of these will provide us a different, multi-messenger window into the variable sky and the cosmos.

The rate of transients alerts from these wide-field surveys far surpasses available spectroscopic resources *already*. Less than 5% of transient events reported to the International Astronomical Union’s Transient Name Server (TNS) are followed up spectroscopically. The order-of-magnitude increase in alerts with LSST will dramatically increase the stress on available spectroscopic resources. In a cosmological context, Type Ia supernovae are especially important for measurements of the Hubble constant at a unique rung of the distance ladder. ML techniques will be essential to classify these events in real time (for active, spectroscopic follow-up) or for archival analysis.

While there has been significant progress in developing ML methods to characterize and classify events in real-time from wide-field survey data (e.g. alert broker systems such as NOIRLab’s ANTARES [61, 62], as well as algorithms such as RAPID [63], SuperNova [64], SuperRAENN [65]), much is left to be done before LSST in order to optimally understand the *diversity* of events we expect to observe. We note that the ML community has also developed many more sophisticated techniques (e.g., transformers, latent

stochastic differential equations, neural processes) which will greatly benefit our domain challenges in the near future.

There is also a major role for unsupervised learning methods in the coming decade, to identify rare events that have never been seen before (anomaly detection), and are buried within the alert stream [66]. Indeed, while only one kilonova event—GW170817—has been identified in both gravitational waves and the electromagnetic spectrum, we should expect LSST to discover many kilonovae that are beyond LIGO-Virgo-Kagra’s detection limit [67]. These unsupervised learning methods can also be used to identify events that were missed during real-time processing, and the promise of these methods has already been demonstrated by teams processing Pan-STARRS, Dark Energy Survey (DES) and Zwicky Transient Facility (ZTF) data (e.g. [68]).

Finally, the key promise of wide-field surveys is understanding *populations* of objects in the time-domain. Modeling these populations to infer fundamental physics (e.g. inferring cosmological parameters from type Ia supernovae) is a complex multi-level modeling (sometimes known as hierarchical Bayesian) problem, and these analyses have either typically made simplifying assumptions in their models to make them tractable to evaluate with traditional inference techniques such as Markov Chain Monte Carlo sampling. However, these simplifying assumptions themselves will lead to systematic biases. Building blocks of machine learning models, such as deep neural networks, can be effectively leveraged within schemes like variational inference in order to perform approximate inference on models defined in high-dimensional spaces. We expect that analysis groups will need to complement traditional techniques on high-performance computing resources with new methods that make much greater use of GPUs.

Multi-messenger Astrophysics (MMA) takes all of the requirements we have detailed above and adds additional complexity. Sources that are energetic enough to distort the fabric of space-time and cause gravitational waves are intrinsically rare, and a large subset of those events that also emit electromagnetic radiation (nearby core-collapse supernovae, neutron star mergers, and the merges of neutron stars and black holes) also evolve extremely quickly. For instance, the canonical MMA event, GW170817, will only have 1–2 detections in a survey such as LSST with its currently planned cadence [69].

Detecting and characterizing MMA events therefore requires discovery of very rare events with very sparse and heterogeneous data from multiple facilities all in real-time. ML methods, including unsupervised anomaly detection techniques [70], and hybrid architectures such as convolutional recurrent neural networks (CRNNs) [71, 72], and simulation-based inference techniques [73] will be necessary to process the smorgasbord of observations from different facilities, flag these events within hours, and automatically trigger follow-up studies. Beyond the ML techniques highlighted in this white paper, MMA will require significant investment in cross-survey cyberinfrastructure to help the community store, process and share a mix of public and private data in order to understand these enigmatic events.

3 Computing and Data

3.1 Data Processing

During and after data acquisition of modern, large-scale cosmological surveys, there are multiple opportunities to improve the speed and accuracy of data processing. We will approach these items sequentially in the nominal processing steps.

First, we envision new modalities for smart triggering—or fast selection and analysis of data of critical elements upon observation. This is important especially for the detection of transient or time-varying data that requires timely follow-up. Currently, this is typically handled highly successfully with difference-imaging (e.g., Antares in LSST). However, the predictive capabilities of this method are limited in terms of the accuracy and prediction of early detection, which constrains our understanding of which types of objects are being selected. Being able to predict the class of objects with fewer observations greatly increases the probability of capturing objects before they fade away. In general, earlier understanding of object status and classification will allow us to choose the most valuable objects for further observation. For objects that are not transient, but uniquely important for other science cases, fast and accurate identification will permit for greater coordination with other observational resources. Future solutions include embedding deep learning algorithms within specialized computational units like TPUs and FPGAs, which are designed for hyperfast data processing. This is the model under current usage within particle physics programs at the Large Hadron Collider. Preliminary work has taken place in astronomy to demonstrate this capability [74]. We envision a possible future where these fast processors are used to pre-process raw data for fast detection: they may be part of an array at the facility where the long-term data processing is occurring, or they may be located at the observational facility. In the latter case, this would bypass data transfer times, as well as facilitate adaptive action for immediate updates of the observing schedule from observation to observation.

Next, consider the cleaning and calibration of data. Multiple types of noise artifacts and confounding patterns can degrade image quality: sensor artifacts, particles (e.g., cosmic rays), bright foreground stars, point spread function variation, galaxy blending, etc. All of these need to be modeled and removed from the data before science can be performed. The first obvious usage for machine learning is the removal of basic artifacts like cosmic rays, which has already been demonstrated [75]. Foreground stars can be detected with tools like Source Extractor, which uses machine learning for e.g., source classification [76]. Patterns in sensors (like the tree rings or brighter-fatter pixels discovered in the Dark Energy Survey data) are more difficult to discover and typically require a great deal of data investigation: replacing this task with algorithmic tools that can discover these patterns has the potential to save a great deal of commissioning and science verification time. Potential tools for this include detailed sensor simulators that can generate patterns in the sensors as propagated to data, as well as inference loops that can test many flexible data models and propose sensor patterns as systematics. Object blending is already a significant challenge for surveys like LSST; the fainter and deeper our imaging becomes, the worse this problem will be. This is currently a feature case study and much progress has been made [4], and we encourage continued work on this challenge.

3.2 High-Performance Computing

We anticipate high-performance computing (HPC) to play a crucial role in our capability to use ML for cosmology this decade. The massive amount of data from cosmological surveys such as the Dark Energy Spectroscopic Instrument (DESI), *eROSITA*, *Euclid*, *Roman Space Telescope*, the *Vera C. Rubin Observatory*, Simons Observatory, and the Square Kilometre Array (SKA) will require dedicated infrastructure to host the data. Likewise, the complexity of the models used to interpret these surveys demands significant computational and storage capabilities.

Consider the example of training an ML model to interpret a galaxy survey by comparison to cosmological simulations. Every stage of this process is an HPC-intensive task: (1) simulations must be produced at large HPC facilities, using $\mathcal{O}(100\text{K}-1\text{M})$ GPU-hours (e.g. AbacusSummit [77]; Uchuu [78]); or $\mathcal{O}(10\text{M}-100\text{M})$ CPU-hours [79–82] (2) outputs are stored on large, high-throughput file systems, totalling $\mathcal{O}(1 \text{ PB})$; (3) the analysis (generation of data products) is executed in cluster environments (often using CPUs) requiring a small percent of the cost incurred in running the simulations; (4) ML models are trained on these data products (using GPUs or other accelerators). The next decade will only see these requirements increase, as upcoming surveys explore both broader sky area and greater depth, demanding simulations with both larger volume and finer resolution. This will increase the HPC expense of every step of this process.

However, ML does not just place demands on HPC, but also offers opportunities. Traditional simulations, or parts of them, can be accelerated with ML [83–85]. High-resolution simulations may be emulated based on low-resolution ones [86–88]; expensive physical calculations replaced with a machine learning interpolation thereof [82]. This may increase the slope of the relation between computational resources and the size and resolution of the simulations. However, this will pose significant challenges to the storage and manipulation of this data, that we expect to be of the order of tens to hundreds of petabytes.

In order to maximize the usage and utilization of simulations and observations, large collections of data must be placed in a publicly accessible location in national or international HPC centers where users can access and work directly with the data without creating copies of it elsewhere. At the same time, subsets of this data should be easy to download, as interesting investigations can often be done on just a few MB of data products.

Cloud computing may be a compelling option for some HPC applications, with its massive storage and distributed compute resources. Cloud computing also offers an opportunity to present users with a standard software stack, promising a solution to the byzantine task of setting up an ML-ready software environment. However, applications like simulations or some kinds of ML training that need efficient communication between many worker nodes may not run well in cloud environments. The opportunities for HPC-in-cloud will surely expand in the next decade, though, driven by ML applications, and efforts to leverage it for scientific workflows should be supported.

Ensuring reproducible research in an HPC context will require new ways for users to interact with HPC resources. Reproducibility efforts in astrophysics are now widespread, and span a spectrum from adopting software-industry testing and continuous integration practices, to fully reproducible workflows that allow readers to reproduce a paper’s results

(e.g. generate all values and figures) at the click of a button. These pipelines tend to be cloud-hosted and use small datasets (few GB) and only a few CPU-hours—far below ML requirements. But the transparency and pedagogical value of such workflows cannot be ignored by the cosmology ML community. HPC centers must offer entry points to their resources that allow automation of these workflows, including access to large datasets and compute resources (especially GPUs).

The increasing demanding computational needs in astrophysics may trigger synergistic collaborations between scientists and the industry. For instance, the methods used in the industry to compress certain classes of data (e.g. videos in YouTube) will reduce the storage needs in astrophysics, enabling deeper and more complete analyses of the data.

In summary, in the next decade HPC must support multiple workflows to maximize science at the intersection of ML and cosmology. The computational infrastructure must enable flagship efforts, such as hosting large observational data sets, running massive cosmological simulations, and training ML models to map between the two. HPC centers must also make such data (simulation and observations) publicly available, such that users can bring their computations to the data or download narrow subsets of the data to support local workflows. The software side will require dedicated, expert effort in the areas of cloud hosting, containerization, and software environments to bridge the gap between the code scientists write and the potential performance HPC hardware promises.

4 Simulations

While humans can easily identify patterns in low-dimensional data, machine learning algorithms can also perform this task in high-dimensional spaces. Thus, simulations can be used as a laboratory to identify unrecognized patterns that can help our understanding of the underlying mechanisms behind the physical process being studied. For instance, it has been shown that unknown relations between galaxy properties and parameters describing the composition of the Universe can be easily identified by employing machine learning techniques on top of state-of-the-art hydrodynamic simulations [89]. We believe that machine learning can trigger a revolution in the large variety of areas of cosmology and galaxy formation that deal with high-dimensional data.

An important question in cosmology is: where does most of the information reside? While for Gaussian density fields (e.g. our Universe on large scales or at early times) all information can be extracted using the power spectrum, for highly non-Gaussian density fields the optimal estimator is unknown. Many different works have shown that there is a wealth of cosmological information that is located on small scales and that cannot be retrieved by using the power spectrum. This motivates the use of these scales to tighten cosmological constraints. Unfortunately, these scales not only are non-perturbative (i.e. they require numerical simulations), but they may also be affected by uncertainties in astrophysical phenomena such as supernova and active galactic nuclei (AGN) feedback. Thus, in a conservative scenario one would like to marginalize over these *baryonic* effects.

In an ideal scenario, it would be desirable to constrain both cosmology and galaxy formation parameters with the highest accuracy. This task could be carried out using machine learning methods, that would require being trained using state-of-the-art cosmological hy-

drodynamic simulations. With those simulations on hand, one could train neural networks to extract the maximum amount of cosmological and astrophysical information from multi-wavelength observations [90]. This would represent the theoretical formalism needed to extract every single bit of information from cosmological surveys.

Achieving this goal requires running large sets of state-of-the-art hydrodynamic simulations and using machine learning to increase the resolution of the simulations. Alternatively, machine learning could be used to accelerate the simulations themselves. Developing and making publicly available standard datasets for cosmological tasks [91] will contribute to the development of the machine learning techniques needed to accomplish these tasks.

However, this ambitious approach is subject to some important caveats. First, it heavily relies on the outcome of numerical simulations, that may or may not, overlap with reality in high dimensions or in low-dimensional projections. For instance, can we trust the outcome of state-of-the-art hydrodynamic simulations when considering the space formed by a large number of galaxy properties? Second, a proper quantification of the statistical, and more important, systematics errors [90] are needed to determine the confidence level of potential discoveries associated to this task (e.g. a detection of the sum of the neutrino masses). Addressing these difficult questions may trigger deep developments not only in the field of hydrodynamic numerical simulations (by improving the simulations to match observations), but also the development of machine learning techniques that may be protected against systematics effects such as numerical artifacts and inaccuracies from the simulations.

With increase in size and complexity of both simulated and observational data, the need for development of more robust machine learning methods, capable of working with multiple datasets at the same time, will increase. This will allow us to combine the knowledge from hydrodynamic simulations and astrophysical surveys, but also learn from the similarities and differences between them. These differences can come from approximations, numerical artifacts, computational constraints, or even unknown physics, not included in the simulations. On the other hand, observational effects, detector errors and even compression and decompression of the observed data can introduce additional differences. These differences can cause models trained on simulations to suffer from very large decrease in performance, when applied to real data.

It has been shown in Refs. [92, 93] that domain adaptation methods offer great promise for astrophysics and cosmology. Domain adaptation methods [94], implemented during model training, help extract only features present in all datasets. This helps align data distributions and build a model that works on all datasets at the same time. Future implementation and development of these methods will be important for building more robust machine learning methods that will not be affected by differences between the datasets, or noise and other perturbations. Furthermore, because domain adaptation methods do not require datasets to be labeled, they can be applied on new observational data. This will be crucial for building automated systems initially trained on simulations, but capable of working in real-time during cosmological survey observations.

5 Survey Operations and Instrument Design

5.1 Survey Operations

Cosmological surveys such as the LSST will repeatedly image the sky, returning to the same part of the sky \sim 1000 times over a period of ten years. The cadence of these observations, the depth of each exposure, and the filters through which the data will be taken will all influence the scientific returns of the survey. For experiments such as the *Vera C. Rubin Observatory*, the need for almost real-time classification of transient events exacerbates the challenges of how to schedule the observations. With multiple scientific objectives and observing conditions that have both short- and long-term coherence timescales (e.g. atmospheric seeing and weather) traditional hand-tuned scheduling strategies or policies are not feasible. Moreover, coordinating observations amongst many telescopes—e.g., for spectroscopic follow-up or multi-wavelength/-messenger measurements—typically requires proposals to time-allocation committees with months of lag time. Finally, how do we optimize the balance of execution toward a specific scientific goal and the exploration for new phenomena?

Current methodologies, like those used in the Dark Energy Survey, rely on simulating many years of high-level metrics (e.g., effective seeing and depth) for a single survey run; humans then review statistics of these metrics, update parameters, and re-simulate until a relative optimum can be found. Adaptive approaches such as feature-based scheduling [95] that are cast as a formal optimization problem [96] in the context of reinforcement learning have shown that they can optimize multiple competing science objectives, and outperform current telescope scheduling approaches while retaining the capacity to recover from unscheduled events (e.g. instrument failures). Unsupervised methods leveraging graph neural networks have been shown in simulations to optimize for scientific objectives without the need for prescribing any observational policy models [97]. The above references are a few of the rare examples of work that have gone into AI-related algorithmic resource allocation for telescope observations.

Pursuing automated-scheduling is likely to lower costs and increase efficiency of telescope observations. Consider the time savings for individual-proposal follow-up campaigns, facility queued observations, and even multi-instrument/-site coordination. Moreover, it has the potential to improve science by allowing us to get the best data for a given goal: the telescope can focus on fields and objects (of class and quality) that will serve the prescribed scientific goal (including exploration).

There are several key challenges for growing and making the most of automated resource allocation in telescope observations. For supervised algorithms that are trained on simulated data or previous observations, there will be a bias when applied to the space of unobserved data: domain adaptation is an avenue to pursue here. Related, interpreting the decisions and policies that are generated by the scheduling algorithms, as well as the downstream economic and logistical consequences, will be essential for verification, human understanding, and safety. To grow this effort, we envision administering data benchmarks and challenges, performing cost-benefit analyses to understand economics, and scaled testing of algorithms on small and large telescopes, as well as federations of telescopes.

5.2 Experiment and Instrument Design

Observational facilities and their instruments are designed to meet scientific requirements. Developing the design specifications that match these requirements is a time- and thus human-intensive process. Due to the evolving demands of scientific discovery—e.g., increased precision and sensitivity of observations—the complexity of instruments is poised to grow significantly; and this could be a significant bottleneck in terms of time and cost for future experiments.

In the traditional design process, somewhat simplistically, humans first generate hypotheses (e.g., dark energy and dark matter affect the expansion of spacetime and thus the distribution of matter), and then they establish quantitative scientific goals (e.g., errors on the dark matter density) to test those hypotheses, generate the summary metrics that come from data sets (e.g., galaxy two-point correlation function), identify the scope and kinds of data that are required to achieve these metrics (e.g., galaxy survey), and then establish the instrument design and observing specifications to achieve these data (e.g., optical telescope and sensor design). The requirement flow-down and design processes are not always this linear from step to step: there are often feedback loops that are shorter than the full loop; an experiment typically has multiple competing science goals, and thus metrics. Each of the components in the requirement flowdown are typically developed independently within different software frameworks—whatever is required for that component. Then, inputs and outputs from each are communicated to their neighboring components, and this feedback process continues until design specifications are achieved that meet the scientific requirements.

Rarely are multiple, let alone all, of the components linked together within a single software framework. For example, in many circumstances, people working on one experimental component communicate inputs and outputs via email. Also, many instrument simulators used for design run primarily or only as standalone programs that require manual operation: ZeMax is used for optics [98], and Comsol for electromagnetic sensors [99]. These all present barriers to fast global optimization of instrument and experiment design.

Connecting these components seamlessly within an algorithmic framework has the potential to reduce bottlenecks and permit fast co-optimization of all the experiment and instrument parameters in service of a set of scientific goals. We envision the conversion of requirement flowdowns into models and numerical data. There are just a few known areas where some of this concept have been implemented. For example, the optimization of adaptive optics systems has subject of deep learning applications (see [100] for a recent review). Many works are exploring the design of optical element configurations with deep neural networks and similarly flexible algorithms [101–105]. Co-design of instruments and observational strategies are being performed to optimize black hole observations [106]. Finally, there is also a burgeoning community of scientists who are approaching design of instruments with machine learning [107] or with the outputs of machine learning-pipelines in mind [108].

The only full simulation of a cosmic experiment that has so far been developed is the SPectrOscopic Ken Simulation (SPOKES) [109], which was designed to replicate the DE-Spec prototype experiment: it includes all major components of a spectroscopic cosmic experiment from data acquisition to analysis to dark energy parameter constraints. Au-

tomating hypothesis generation has not yet been implemented in SPOKES, and there has been limited work in this area for physical experiments – primarily related to the generation of symbolic models [110, 111]. When a field of study is faced with a large landscape of possibilities for theoretical development, automating hypothesis generation has the potential to accelerate the search for new ideas to pursue; this may be of significant import for studying dark matter and dark energy.

There are many potential avenues for the continued development of this area. For example, a project to design a simple experiment (including the instruments) that is optimized toward a scientific goal can provide a conceptual launching-off point: imagine performing the automated design of a telescope system component-by-component, and then doing that all in tandem. We recommend significant investment in the following: 1) building full digital twins of experiments (including instruments); 2) exploring the capacity of machine learning for optimizing the design of individual instruments and all instruments within an experiment; and 3) developing the capacity for automating hypotheses, which is among the tightest of bottlenecks in scientific discovery. Also, some cost-benefit analysis should be performed to study the utility and potential economic gains of automating instrument design.

6 Machine Learning Architectures

Until recently, machine learning methods applied in cosmology were borrowed wholesale from the computer science literature. The specific needs of the cosmology community combined with the highly structured nature of cosmological datasets necessitates theoretical and methodological developments on the machine learning side specifically tailored for cosmology applications. Here, we outline ongoing as well as necessary efforts in this direction.

The modern deep learning revolution arguably kicked off in 2012 with deep convolutional neural networks like AlexNet showing record-breaking performance on benchmark vision tasks. Since then, convolutional neural networks have found widespread application within cosmology due to the fact that a lot of cosmological data is naturally represented in the image domain. Despite their success, traditional convolutional neural networks can only process rectangular images, restricting their use in the cosmological context. The need to extend desired implicit biases like translation and rotation symmetry to a more diverse set of domains—such as the celestial sphere—has spurred methodological developments in neural network architectures specifically for cosmological applications. Examples of these developments are custom spherical convolutional networks [112], efficient translation- and rotation-invariant normalizing flows [113], and fixed kernel-based convolutional architectures [114–116]. The needs addressed by such developments are often distinct from those in mainstream computer vision research, which is predominantly focused on application to natural images.

As outlined above, of particular importance for cosmology surveys is the development of methodologies for the analysis of time series data (e.g. classification of light curves from Type 1a supernovae). Recurrent Neural Networks [117], Long Short-Term Memory Networks [118], and Transformers [119] have made significant advances for classifica-

tion problems, but sparsely sampled light curves with noisy observations present many challenges for standard neural network architectures common in applications outside of astrophysics. Physics-inspired and physics-constrained neural networks [120] have shown promise in reducing the dimensionality of the underlying latent space of a network with an associated reduction in the size of data sets needed to train the network. Most of these applications, however, simply adopt existing network architectures without extending them to the specific requirements of cosmology. For example, many network architectures are sequential and recursive, which does not easily allow parallel computation (to account for the size of cosmological data sets) or they assume a Markov process, which means that they cannot easily learn long-range dependencies.

A specific effort to bridge methodological developments in the computer science literature with needs within cosmology is therefore necessary. As a concrete example, spherical convolutional neural networks have seen widespread recent adoption for various cosmological applications (e.g., Refs. [121–123]) due to the existence of easy-to-use implementations tailored to common pixelizations schemes like HEALPix which are widely used in cosmology [112, 124]. This showcases the need for continuous integration and cross-talk between the communities.

The reliability of machine learning models for scientific discovery can be interrogated through the development and application of interpretable and explainable ML systems. Although a mature subfield, explainable AI (XAI) has made limited inroads within cosmology [125, 126]. This can at least in part be attributed to the lack of methods that can be robustly applied in typical scenarios of interest; we therefore see a critical need for methodological development in this direction. Post-hoc application of symbolic regression methods to trained neural networks [53, 127–129] towards model interpretability and model discovery has recently demonstrated promise within cosmology, and we expect this direction to continue to flourish in the near future.

7 Uncertainty Quantification and Bias

A major historic barrier to using machine learning models such as deep neural networks within cosmology has been the lack of a robust and validatable statistical description of their output. With increasing adoption of these methods within cosmology, and more broadly in the sciences, the need for such descriptions have become acute. Here, we summarize recent developments in uncertainty and bias quantification, and outline necessary developments for the widespread adoption of machine learning models in the cosmological context.

Perhaps the most common task in cosmological data analysis is parameter inference—i.e., describing what a given set of observations tells us about a set of physically-meaningful parameters in a model (for example, the six Λ CDM parameters). In the supervised learning setting, the simplest and historically most common way to frame this is as a regression task—e.g., teaching a neural network to minimize the mean square error (MSE) between its output and a set of target parameters. With the expressiveness of the neural network and number of simulations under control, it can be shown the minimizing the MSE loss provides the maximum-likelihood estimate (MLE) [130]. Although statistically well-defined,

the MLE and, more generally, any point estimate, is of limited scientific value when the goal is to quantify if, and how well, a given point in the parameter space is compatible with the data. In cosmology, this fact is further exacerbated since physical degeneracies between parameters are ubiquitous, and there is often the need to compare and/or combine results stemming from different experiments and observations.

Several methods for uncertainty quantification have been developed in the machine learning literature, e.g., Monte Carlo dropout [131], Bayesian neural networks [132], and deep ensembles [133]. Direct application of these methods within cosmology and their statistical interpretation is often stymied due to a difference in language in how the two fields talk about uncertainty. In particular, *aleatoric* and *epistemic* uncertainties (sometimes called *data* and *model* uncertainties, respectively) are often distinguished in the machine learning literature. Although these terms provide useful abstractions, their use in physics can be confusing since they do not neatly map onto the more familiar concept of statistical and systematic uncertainty.

The difference in perspectives can be at least partially attributed to what kind of uncertainty is most important in either field. In machine learning, *model* usually refers to the trained model (e.g., a neural network), since the data (e.g., natural language or images) is often assumed to be fixed with little insight into its generating process. On the other hand, in physics we often have substantial insight into the data-generating processes, with *model* encompassing the underlying physics processes, including detector effects that led to the generation of the dataset. In typical machine learning applications, the aforementioned methods are often used to account for the uncertainty in the machine learning model. When adapted for physics applications, a subset can also account for uncertainty associated with the data-generating process. These methods can produce biased or mis-calibrated uncertainty estimates, even when applied to simple physical systems [134].

Probabilistic machine learning methods have recently demonstrated significant promise for scientific applications, and particularly within cosmology. Methods under the umbrella of *simulation-based inference* (SBI) aim to estimate full Bayesian posteriors, as well as likelihoods or likelihood ratios (and associated confidence regions) in the setting where the likelihood $p(x | \theta)$ of the data-generating process associated with data x and parameters θ is unavailable or intractable, but samples $x, \theta \sim p(x | \theta)$ can be simulated from a mechanistic forward model. See Ref. [135] for a recent review. These methods are related in lineage to Approximate Bayesian Computation (ABC), which has been extensively used within cosmology in the last few decades. A typical application involves using a domain-specific neural network, or a projected set of summary statistics, as an input to a parametrized classifier or a conditioning context for a density estimator in order to learn an estimator for the quantity of interest (e.g., posterior or likelihood). Indeed, cosmology-inspired research has motivated the development of efficient SBI methods and algorithms [136–142].

Training parametrized classifiers and density estimators is a challenging task in machine learning, and often requires a large number of simulations and/or a post-hoc calibration procedure in order to produce satisfactory results. It was recently demonstrated that typically-used SBI algorithms tend to produce overly confident posterior estimates [143]—an unacceptable outcome in cosmological applications. It is therefore imperative to perform diagnostic coverage tests (as done in, e.g., Refs. [143–145]) to ensure that a trained machine learning-based estimator produces well-calibrated and statistically-consistent re-

sults, as well as incentivize continued development of simulation-based inference methods towards robust scientific discovery.

8 Education and Outreach

Cosmology, and in fact science at large, is at the cusp of a data-driven evolution driven by the ever increasing open-source software and massive public datasets (both real and simulated), as mentioned above. With this revolution, we have the opportunity to create a more equitable scientific community and greatly increase the reach of our scientific impacts (i.e., create an “inclusion revolution”, [146]). Here we touch upon educational needs for cross-disciplinary scientists and the requirements for effective outreach programs to the broader public.

Cutting-edge advances in data-driven cosmology require increasingly complex techniques. Cosmologists now regularly publish in ML venues such as NeurIPS and ICML (see, e.g., the recurrent Machine Learning and the Physical Sciences Workshop at NeurIPS). Competitive progress therefore requires mastery of physics (theoretical and/or observational), statistics and data science—a tall order for the upcoming generation of scientists. Massive Open Online Courses (MOOCs) and educational-focused journals (e.g., distill.pub) have addressed this need in varied educational backgrounds by providing free, online supplementary courses to a more traditional physics background. However, few resources exist *by* physicists and *for* physicists that detail scientifically-informed, data-driven methodologies most relevant to our field. Interdisciplinary AI centers are increasingly common, as well as open seminar series which are often broadcast online. We encourage these new foci to continue to facilitate open discussion. We additionally encourage opportunities for collaborative development of new curricula for physics-informed ML. Finally, in the aims of diversifying the field, we encourage funding agencies to directly track progress on inclusion of funded programs by requiring documentation of diversity, inclusion and equity efforts [147].

Looking beyond the academic community, cosmologists are additionally at the forefront of new opportunities to engage with the broader public, especially those of historically marginalized groups. On the broadest scale, computational capabilities of handheld electronics (e.g., phones, tablets) allows for unprecedented interactivity with high-quality simulations and graphics. Cosmologists can also capitalize on the growing the community of online educational tools and communities for ML by providing datasets to these communities. One potential avenue for this is creating public competitions which is free for participants. We highlight two successful competitions: Galaxy Zoo [148] (in which competitors are tasked with classifying the morphologies of galaxies from images) and PLAsTiCC [149] (in which competitors must classify transients simulated in an LSST-like datastream). Both of these were held on the website Kaggle and garnered hundreds of submissions. However, we do caution that these communities and competitions likely do not reach out to the most marginalized communities, and indeed may contribute to widening the gap between those with and without dedicated educational resources (as is well known for a need for dedicated education and outreach expertise [150]).

9 Workforce Development

It is undeniable that the rapid progress in AI/ML has been guided by the commercial gains of these techniques, especially in various technology sectors. At the same time, more than 60% of physics college graduates end up in the private sector; within this fraction, over 50% of graduates go into engineering or computational sciences [151]. Both academic and industry settings require knowledge of algorithms, high-performance computing, software carpentry, statistics and machine learning. Integrated computational education within physics curricula is increasingly relevant for preparing a tech savvy workforce. While students can supplement their core curricula with computer science courses, focused courses which teach these fundamental skillsets are becoming increasingly necessary [152].

Industry is also taking advantage of the open-access, large datasets from cosmology and increasingly creating opportunities to collaborate with academics on these problems (e.g., the Pitt-Google LSST Broker being developed for the community). We encourage continued collaboration with industry, through visiting scholar programs or more junior funded research opportunities.

It is increasingly important that a broad and diverse range of people be involved in these new technologies in order to bring critical perspectives to bear. If innovation is to flourish, the goal must be to nurture critical and questioning perspectives, shaped by a wide breath of experience, particularly experiences different from the traditional norms currently represented in physics, math and data science. While such a goal is beneficial to the study of these disciplines, it is crucial for the development of applications and technologies that will exploit algorithms developed by any future workforce. As the physics community engages in the education of the next generation to design, develop and build the future tools, a broad perspective must be taken to train and engage students in the ethics and efficacy of these technologies. Recruiting and retaining currently minoritized groups is an important piece of expanding the innovation quotient of the workforce.

The lack of input from groups who are often minoritized in the fields of physics and data science can lead to algorithms and AI technologies that reflect the biases and prejudice (conscious or unconscious) of a more narrowly experienced community. Algorithms that mimic racial and gender discrimination are rife within our society. Without opportunities to broaden the range of those who design and build these technologies, we will be unable to advance justice in the scientific culture.

10 Recommendations and Vision

Machine learning is set to play an increasingly important role in many facets of cosmology in the next decade. We conclude here by outlining a set of recommendations aimed at realizing our vision for the next generation in the development of artificial intelligence applied to cosmology. We target these recommendations to a broad community of researchers in the high-energy physics and astrophysics communities.

- Incentivize the use of machine learning techniques when there is a specific advantage towards achieving the stated scientific goal over more thoroughly-vetted traditional

statistical (e.g., providing the ability to extract more useful information from the data or to emulate expensive simulations).

- For parameter inference, eschew point estimation in favor of principled uncertainty quantification, e.g. through simulation-based inference techniques. Assess the statistical soundness of inference pipelines through coverage tests ensuring that the derived quantities of interest (e.g., parameter posteriors) are well-calibrated.
- Encourage the production and public dissemination of general-purpose cosmological simulations geared towards machine learning applications. Enable this through the availability of centralized and decentralized institutional resources (e.g. HPC clusters as well as cloud computing and storage facilities).
- Incentivize modular and reproducible workflows in cosmological analyses that can be built upon, independently validated, and used for pedagogical purposes.
- Invest in cross-survey cyberinfrastructure in order to fully exploit synergies between different cosmological surveys for real-time as well as archival multi-messenger analysis.
- Recognizing that our models for complex processes in the Universe (reflected in, e.g., simulations) are often not perfect, understand and account for systematic differences between the training and the data distributions through, e.g., principled uncertainty propagation, the use of nuisance parameters defining a flexible family of training distributions, or domain adaptation techniques. Develop and use methods that are robust to systematic model mis-specification in order to minimize the risk of biases.
- Adopt and develop explainable artificial intelligence techniques with the goal of model interpretability and discovery in cosmology.
- To the extent possible, minimize data augmentation by leveraging machine learning architectures that encode the expected physical symmetries of the problem and the structure of the data domain. Given the diversity of cosmological datasets, develop efficient architectures suited to processing common cosmological data modalities.
- Nurture, through education and workforce development, career paths for practitioners with expertise overlap in machine learning methods as well as cosmological applications. Develop curriculum-based learning at the intersection of cosmology and data science aimed at PhD-level and early-career researchers.
- Develop community-driven data benchmarks and challenges with specific scientific goals in mind; these should be fostered by funding agencies to promote participation.
- Understand and consider our responsibility towards society at large when developing techniques based on artificial intelligence. Critically assess the broader societal (e.g., potential for harm) and environmental (e.g., carbon footprint) impact of research performed at the intersection of artificial intelligence and cosmology.

Acknowledgments

CD is partially supported by NSF grant AST-1813694. SM is partially supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics of U.S. Department of Energy under grant Contract Number DE-SC0012567. VAV is partially supported by the NSF grant AST-2108676. This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. This work is supported by the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions).

References

- [1] R. Joseph, P. Melchior, and F. Moolekamp, “Joint survey processing: combined resampling and convolution for galaxy modelling and deblending,” *arXiv e-prints* (July, 2021) arXiv:2107.06984, arXiv:2107.06984 [astro-ph.IM].
- [2] B. J. Brewer, D. Foreman-Mackey, and D. W. Hogg, “Probabilistic Catalogs for Crowded Stellar Fields,” *AJ* **146** no. 1, (July, 2013) 7, arXiv:1211.5805 [astro-ph.IM].
- [3] T. Daylan, S. K. N. Portillo, and D. P. Finkbeiner, “Inference of Unresolved Point Sources At High Galactic Latitudes Using Probabilistic Catalogs,” *Astrophys. J.* **839** no. 1, (2017) 4, arXiv:1607.04637 [astro-ph.IM].
- [4] R. Liu, J. D. McAuliffe, and J. Regier, “Variational Inference for Deblending Crowded Starfields,” *arXiv e-prints* (Feb., 2021) arXiv:2102.02409, arXiv:2102.02409 [astro-ph.IM].
- [5] R. B. Metcalf *et al.*, “The Strong Gravitational Lens Finding Challenge,” *Astron. Astrophys.* **625** (2019) A119, arXiv:1802.03609 [astro-ph.GA].
- [6] C. Jacobs, T. Collett, K. Glazebrook, C. McCarthy, A. K. Qin, T. M. C. Abbott, F. B. Abdalla, J. Annis, S. Avila, K. Bechtol, E. Bertin, D. Brooks, E. Buckley-Geer, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, L. N. da Costa, C. Davis, J. De Vicente, S. Desai, H. T. Diehl, P. Doel, T. F. Eifler, B. Flaugher, J. Frieman, J. García-Bellido, E. Gaztanaga, D. W. Gerdes, D. A. Goldstein, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, W. G. Hartley, D. L. Hollowood, K. Honscheid, B. Hoyle, D. J. James, K. Kuehn, N. Kuropatkin, O. Lahav, T. S. Li, M. Lima, H. Lin, M. A. G. Maia, P. Martini, C. J. Miller, R. Miquel, B. Nord, A. A. Plazas, E. Sanchez, V. Scarpine, M. Schubnell, S. Serrano, I. Sevilla-Noarbe, M. Smith, M. Soares-Santos, F. Sobreira, E. Suchyta, M. E. C. Swanson, G. Tarle, V. Vikram, A. R. Walker, Y. Zhang, J. Zuntz, and DES Collaboration, “Finding high-redshift strong lenses in DES using convolutional neural networks,” *MNRAS* **484** no. 4, (Apr., 2019) 5330–5349, arXiv:1811.03786 [astro-ph.GA].

- [7] **DES** Collaboration, J. H. O'Donnell *et al.*, “The DES Bright Arcs Survey: Candidate Strongly Lensed Galaxy Systems from the Dark Energy Survey 5,000 Sq. Deg. Footprint,” arXiv:2110.02418 [astro-ph.CO].
- [8] X. Huang, C. Storfer, A. Gu, V. Ravi, A. Pilon, W. Sheu, R. Venguswamy, S. Banka, A. Dey, M. Landriau, D. Lang, A. Meisner, J. Moustakas, A. D. Myers, R. Sajith, E. F. Schlaflly, and D. J. Schlegel, “Discovering New Strong Gravitational Lenses in the DESI Legacy Imaging Surveys,” *ApJ* **909** no. 1, (Mar., 2021) 27, arXiv:2005.04730 [astro-ph.IM].
- [9] G. Stein, J. Blaum, P. Harrington, T. Medan, and Z. Lukic, “Mining for strong gravitational lenses with self-supervised learning,” arXiv:2110.00023 [astro-ph.IM].
- [10] Y. D. Hezaveh, L. Perreault Levasseur, and P. J. Marshall, “Fast Automated Analysis of Strong Gravitational Lenses with Convolutional Neural Networks,” *Nature* **548** (2017) 555–557, arXiv:1708.08842 [astro-ph.IM].
- [11] L. Perreault Levasseur, Y. D. Hezaveh, and R. H. Wechsler, “Uncertainties in Parameters Estimated with Neural Networks: Application to Strong Gravitational Lensing,” *Astrophys. J. Lett.* **850** no. 1, (2017) L7, arXiv:1708.08843 [astro-ph.CO].
- [12] W. R. Morningstar, Y. D. Hezaveh, L. Perreault Levasseur, R. D. Blandford, P. J. Marshall, P. Putzky, and R. H. Wechsler, “Analyzing interferometric observations of strong gravitational lenses with recurrent and convolutional neural networks,” arXiv:1808.00011 [astro-ph.IM].
- [13] **LSST Dark Energy Science** Collaboration, S. Wagner-Carena, J. W. Park, S. Birrer, P. J. Marshall, A. Roodman, and R. H. Wechsler, “Hierarchical Inference with Bayesian Neural Networks: An Application to Strong Gravitational Lensing,” *Astrophys. J.* **909** no. 2, (2021) 187, arXiv:2010.13787 [astro-ph.CO].
- [14] M. Chianese, A. Coogan, P. Hofma, S. Otten, and C. Weniger, “Differentiable Strong Lensing: Uniting Gravity and Neural Nets through Differentiable Probabilistic Programming,” *Mon. Not. Roy. Astron. Soc.* **496** no. 1, (2020) 381–393, arXiv:1910.06157 [astro-ph.CO].
- [15] K. Karchev, A. Coogan, and C. Weniger, “Strong-lensing source reconstruction with variationally optimised Gaussian processes,” arXiv:2105.09465 [astro-ph.IM].
- [16] W. R. Morningstar, L. Perreault Levasseur, Y. D. Hezaveh, R. Blandford, P. Marshall, P. Putzky, T. D. Rueter, R. Wechsler, and M. Welling, “Data-Driven Reconstruction of Gravitationally Lensed Galaxies using Recurrent Inference Machines,” arXiv:1901.01359 [astro-ph.IM].
- [17] **LSST Dark Energy Science** Collaboration, J. W. Park, S. Wagner-Carena, S. Birrer, P. J. Marshall, J. Y.-Y. Lin, and A. Roodman, “Large-Scale Gravitational Lens

Modeling with Bayesian Neural Networks for Accurate and Precise Inference of the Hubble Constant,” *Astrophys. J.* **910** no. 1, (2021) 39, arXiv:2012.00042 [astro-ph.IM].

- [18] T. Daylan, F.-Y. Cyr-Racine, A. Diaz Rivero, C. Dvorkin, and D. P. Finkbeiner, “Probing the small-scale structure in strongly lensed systems via transdimensional inference,” *Astrophys. J.* **854** no. 2, (2018) 141, arXiv:1706.06111 [astro-ph.CO].
- [19] S. Alexander, S. Gleyzer, E. McDonough, M. W. Toomey, and E. Usai, “Deep Learning the Morphology of Dark Matter Substructure,” *Astrophys. J.* **893** (2020) 15, arXiv:1909.07346 [astro-ph.CO].
- [20] A. Diaz Rivero and C. Dvorkin, “Direct Detection of Dark Matter Substructure in Strong Lens Images with Convolutional Neural Networks,” *Phys. Rev. D* **101** no. 2, (2020) 023515, arXiv:1910.00015 [astro-ph.CO].
- [21] S. Alexander, S. Gleyzer, H. Parul, P. Reddy, M. W. Toomey, E. Usai, and R. Von Klar, “Decoding Dark Matter Substructure without Supervision,” arXiv:2008.12731 [astro-ph.CO].
- [22] B. Ostdiek, A. Diaz Rivero, and C. Dvorkin, “Extracting the Subhalo Mass Function from Strong Lens Images with Image Segmentation,” *Astrophys. J.* **927** no. 1, (2022) 83, arXiv:2009.06639 [astro-ph.CO].
- [23] B. Ostdiek, A. Diaz Rivero, and C. Dvorkin, “Image segmentation for analyzing galaxy-galaxy strong lensing systems,” *Astron. Astrophys.* **657** (2022) L14, arXiv:2009.06663 [astro-ph.CO].
- [24] A. Coogan, K. Karchev, and C. Weniger, “Targeted Likelihood-Free Inference of Dark Matter Substructure in Strongly-Lensed Galaxies,” in *34th Conference on Neural Information Processing Systems*. 10, 2020. arXiv:2010.07032 [astro-ph.CO].
- [25] Q. He, A. Robertson, J. Nightingale, S. Cole, C. S. Frenk, R. Massey, A. Amvrosiadis, R. Li, X. Cao, and A. Etherington, “A forward-modelling method to infer the dark matter particle mass from strong gravitational lenses,” *Mon. Not. Roy. Astron. Soc.* **511** no. 2, (2022) 3046–3062, arXiv:2010.13221 [astro-ph.CO].
- [26] R. Legin, Y. Hezaveh, L. P. Levasseur, and B. Wandelt, “Simulation-Based Inference of Strong Gravitational Lensing Parameters,” arXiv:2112.05278 [astro-ph.CO].
- [27] S. Alexander, S. Gleyzer, P. Reddy, M. Tidball, and M. W. Toomey, “Domain Adaptation for Simulation-Based Dark Matter Searches Using Strong Gravitational Lensing,” arXiv:2112.12121 [astro-ph.CO].

- [28] S. Wagner-Carena, J. Aalbers, S. Birrer, E. O. Nadler, E. Darragh-Ford, P. J. Marshall, and R. H. Wechsler, “From Images to Dark Matter: End-To-End Inference of Substructure From Hundreds of Strong Gravitational Lenses,” [arXiv:2203.00690](https://arxiv.org/abs/2203.00690) [astro-ph.CO].
- [29] J. Brehmer, S. Mishra-Sharma, J. Hermans, G. Louppe, and K. Cranmer, “Mining for Dark Matter Substructure: Inferring subhalo population properties from strong lenses with machine learning,” *Astrophys. J.* **886** no. 1, (2019) 49, [arXiv:1909.02005](https://arxiv.org/abs/1909.02005) [astro-ph.CO].
- [30] N. Jeffrey, J. Alsing, and F. Lanusse, “Likelihood-free inference with neural compression of DES SV weak lensing map statistics,” *Mon. Not. Roy. Astron. Soc.* **501** no. 1, (2021) 954–969, [arXiv:2009.08459](https://arxiv.org/abs/2009.08459) [astro-ph.CO].
- [31] A. Gupta, J. M. Z. Matilla, D. Hsu, and Z. Haiman, “Non-Gaussian information from weak lensing data via deep learning,” *Phys. Rev. D* **97** no. 10, (2018) 103515, [arXiv:1802.01212](https://arxiv.org/abs/1802.01212) [astro-ph.CO].
- [32] J. Fluri, T. Kacprzak, A. Refregier, A. Amara, A. Lucchi, and T. Hofmann, “Cosmological constraints from noisy convergence maps through deep learning,” *Phys. Rev. D* **98** no. 12, (2018) 123518, [arXiv:1807.08732](https://arxiv.org/abs/1807.08732) [astro-ph.CO].
- [33] D. Ribli, B. A. Pataki, J. M. Zorrilla Matilla, D. Hsu, Z. Haiman, and I. Csabai, “Weak lensing cosmology with convolutional neural networks on noisy data,” *Mon. Not. Roy. Astron. Soc.* **490** no. 2, (2019) 1843–1860, [arXiv:1902.03663](https://arxiv.org/abs/1902.03663) [astro-ph.CO].
- [34] J. Fluri, T. Kacprzak, A. Lucchi, A. Refregier, A. Amara, T. Hofmann, and A. Schneider, “Cosmological constraints with deep learning from KiDS-450 weak lensing maps,” *Phys. Rev. D* **100** no. 6, (2019) 063514, [arXiv:1906.03156](https://arxiv.org/abs/1906.03156) [astro-ph.CO].
- [35] A. Mootoovaloo, A. F. Heavens, A. H. Jaffe, and F. Leclercq, “Parameter Inference for Weak Lensing using Gaussian Processes and MOPED,” *Mon. Not. Roy. Astron. Soc.* **497** no. 2, (2020) 2213–2226, [arXiv:2005.06551](https://arxiv.org/abs/2005.06551) [astro-ph.CO].
- [36] T. Lu, Z. Haiman, and J. M. Z. Matilla, “Simultaneously constraining cosmology and baryonic physics via deep learning from weak lensing,” *Mon. Not. Roy. Astron. Soc.* **511** no. 1, (2022) 1518–1528, [arXiv:2109.11060](https://arxiv.org/abs/2109.11060) [astro-ph.CO].
- [37] M. Tewes, T. Kuntzer, R. Nakajima, F. Courbin, H. Hildebrandt, and T. Schrabback, “Weak-lensing shear measurement with machine learning: teaching artificial neural networks about feature noise,” *Astron. Astrophys.* **621** (2019) A36, [arXiv:1807.02120](https://arxiv.org/abs/1807.02120) [astro-ph.CO].

- [38] O. M. Springer, E. O. Ofek, Y. Weiss, and J. Merten, “Weak lensing shear estimation beyond the shape-noise limit: a machine learning approach,” *Mon. Not. Roy. Astron. Soc.* **491** no. 4, (2020) 5301–5316, arXiv:1808.07491 [astro-ph.CO].
- [39] N. Jeffrey, F. Lanusse, O. Lahav, and J.-L. Starck, “Deep learning dark matter map reconstructions from DES SV weak lensing data,” *Mon. Not. Roy. Astron. Soc.* **492** no. 4, (2020) 5023–5029, arXiv:1908.00543 [astro-ph.CO].
- [40] B. Remy, F. Lanusse, N. Jeffrey, J. Liu, J.-L. Starck, K. Osato, and T. Schrabback, “Probabilistic Mass Mapping with Neural Score Estimation,” arXiv:2201.05561 [astro-ph.CO].
- [41] J. a. Caldeira, W. L. K. Wu, B. Nord, C. Avestruz, S. Trivedi, and K. T. Story, “DeepCMB: Lensing Reconstruction of the Cosmic Microwave Background with Deep Neural Networks,” *Astron. Comput.* **28** (2019) 100307, arXiv:1810.01483 [astro-ph.CO].
- [42] K. Aylor, M. Haq, L. Knox, Y. Hezaveh, and L. Perreault-Levasseur, “Cleaning our own dust: simulating and separating galactic dust foregrounds with neural networks,” *MNRAS* **500** no. 3, (Jan., 2021) 3889–3897, arXiv:1909.06467 [astro-ph.IM].
- [43] M. A. Petroff, G. E. Addison, C. L. Bennett, and J. L. Weiland, “Full-sky Cosmic Microwave Background Foreground Cleaning Using Machine Learning,” *Astrophys. J.* **903** no. 2, (2020) 104, arXiv:2004.11507 [astro-ph.CO].
- [44] N. Jeffrey, F. Boulanger, B. D. Wandelt, B. R.-S. Blancard, E. Ally, and F. Levrier, “Single frequency CMB B-mode inference with realistic foregrounds from a single training image,” *Mon. Not. Roy. Astron. Soc.* **510** no. 1, (2022) L1–L6, arXiv:2111.01138 [astro-ph.CO].
- [45] E. Guzman and J. Meyers, “Reconstructing cosmic polarization rotation with ResUNet-CMB,” *JCAP* **01** no. 01, (2022) 030, arXiv:2109.09715 [astro-ph.CO].
- [46] **Simons Observatory** Collaboration, P. Ade *et al.*, “The Simons Observatory: Science goals and forecasts,” *JCAP* **02** (2019) 056, arXiv:1808.07445 [astro-ph.CO].
- [47] **CMB-S4** Collaboration, K. N. Abazajian *et al.*, “CMB-S4 Science Book, First Edition,” arXiv:1610.02743 [astro-ph.CO].
- [48] N. Sehgal *et al.*, “CMB-HD: An Ultra-Deep, High-Resolution Millimeter-Wave Survey Over Half the Sky,” arXiv:1906.10134 [astro-ph.CO].
- [49] Z. Lin, N. Huang, C. Avestruz, W. L. K. Wu, S. Trivedi, J. Caldeira, and B. Nord, “DeepSZ: identification of Sunyaev–Zel’dovich galaxy clusters using deep

learning,” *Mon. Not. Roy. Astron. Soc.* **507** no. 3, (2021) 4149–4164, arXiv:2102.13123 [astro-ph.CO].

[50] D. de Andres, W. Cui, F. Ruppin, M. De Petris, G. Yepes, I. Lahouli, G. Aversano, R. Dupuis, and M. Jarraya, “Mass Estimation of Planck Galaxy Clusters using Deep Learning,” *EPJ Web Conf.* **257** (2022) 00013, arXiv:2111.01933 [astro-ph.CO].

[51] N. Gupta and C. L. Reichardt, “Mass Estimation of Galaxy Clusters with Deep Learning. I. Sunyaev–Zel’dovich Effect,” *Astrophys. J.* **900** no. 2, (2020) 110, arXiv:2003.06135 [astro-ph.CO].

[52] N. Gupta and C. L. Reichardt, “Mass Estimation of Galaxy Clusters with Deep Learning II. Cosmic Microwave Background Cluster Lensing,” *Astrophys. J.* **923** no. 1, (2021) 96, arXiv:2005.13985 [astro-ph.CO].

[53] D. Wadekar, F. Villaescusa-Navarro, S. Ho, and L. Perreault-Levasseur, “Modeling assembly bias with machine learning and symbolic regression,” arXiv:2012.00111 [astro-ph.CO].

[54] S. Hassan, S. Andrianomena, and C. Doughty, “Constraining the astrophysics and cosmology from 21 cm tomography using deep learning with the SKA,” *Mon. Not. Roy. Astron. Soc.* **494** no. 4, (2020) 5761–5774, arXiv:1907.07787 [astro-ph.CO].

[55] S. Neutsch, C. Heneka, and M. Brüggen, “Inferring Astrophysics and Dark Matter Properties from 21cm Tomography using Deep Learning,” arXiv:2201.07587 [astro-ph.CO].

[56] S. Sikder, R. Barkana, I. Reis, and A. Fialkov, “Machine Learning to Decipher the Astrophysical Processes at Cosmic Dawn,” arXiv:2201.08205 [astro-ph.CO].

[57] H. J. Hortúa, L. Malago, and R. Volpi, “Constraining the Reionization History using Bayesian Normalizing Flows,” *Mach. Learn. Sci. Tech.* **1** (2020) 035014, arXiv:2005.07694 [astro-ph.CO].

[58] T. L. Makinen, L. Lancaster, F. Villaescusa-Navarro, P. Melchior, S. Ho, L. Perreault-Levasseur, and D. N. Spergel, “deep21: a deep learning method for 21 cm foreground removal,” *JCAP* **04** (2021) 081, arXiv:2010.15843 [astro-ph.CO].

[59] P. Villanueva-Domingo and F. Villaescusa-Navarro, “Removing Astrophysics in 21 cm maps with Neural Networks,” *Astrophys. J.* **907** no. 1, (2021) 44, arXiv:2006.14305 [astro-ph.CO].

[60] A. Ghosh *et al.*, “Foreground modelling via Gaussian process regression: an application to HERA data,” *Mon. Not. Roy. Astron. Soc.* **495** no. 3, (2020) 2813–2826, arXiv:2004.06041 [astro-ph.CO].

- [61] ANTARES Collaboration, G. Narayan *et al.*, “Machine Learning-based Brokers for Real-time Classification of the LSST Alert Stream,” *Astrophys. J. Suppl.* **236** no. 1, (2018) 9, arXiv:1801.07323 [astro-ph.IM].
- [62] T. Matheson, C. Stubens, N. Wolf, C.-H. Lee, G. Narayan, A. Saha, A. Scott, M. Soraisam, A. S. Bolton, B. Hauger, D. R. Silva, J. Kececioglu, C. Scheidegger, R. Snodgrass, P. D. Aleo, E. Evans-Jacquez, N. Singh, Z. Wang, S. Yang, and Z. Zhao, “The ANTARES Astronomical Time-domain Event Broker,” *AJ* **161** no. 3, (Mar., 2021) 107, arXiv:2011.12385 [astro-ph.IM].
- [63] D. Muthukrishna, G. Narayan, K. S. Mandel, R. Biswas, and R. Hložek, “RAPID: Early Classification of Explosive Transients using Deep Learning,” *Publ. Astron. Soc. Pac.* **131** no. 1005, (2019) 118002, arXiv:1904.00014 [astro-ph.IM].
- [64] A. Möller and T. de Boissière, “SuperNNova: an open-source framework for Bayesian, neural network-based supernova classification,” *MNRAS* **491** no. 3, (Jan., 2020) 4277–4293, arXiv:1901.06384 [astro-ph.IM].
- [65] V. A. Villar *et al.*, “SuperRAENN: A Semisupervised Supernova Photometric Classification Pipeline Trained on Pan-STARRS1 Medium-Deep Survey Supernovae,” *Astrophys. J.* **905** no. 2, (2020) 94, arXiv:2008.04921 [astro-ph.HE].
- [66] X. Li, F. Ragosta, W. I. Clarkson, and F. B. Bianco, “Preparing to discover the unknown with rubin lsst: Time domain,” *The Astrophysical Journal Supplement Series* **258** no. 1, (2021) 2.
- [67] I. Andreoni, R. Margutti, O. S. Salafia, B. Parazin, V. A. Villar, M. W. Coughlin, P. Yoachim, K. Mortensen, D. Brethauer, S. Smartt, *et al.*, “Target of opportunity observations of gravitational wave events with vera c. rubin observatory,” *arXiv preprint arXiv:2111.01945* (2021) .
- [68] K. L. Malanchev, M. V. Pruzhinskaya, V. S. Korolev, P. D. Aleo, M. V. Kornilov, E. E. O. Ishida, V. V. Krushinsky, F. Mondon, S. Sreejith, A. A. Volnova, A. A. Belinski, A. V. Dodin, A. M. Tatarnikov, S. G. Zheltoukhov, and (The SNAD Team), “Anomaly detection in the Zwicky Transient Facility DR3,” *MNRAS* **502** no. 4, (Apr., 2021) 5147–5175, arXiv:2012.01419 [astro-ph.IM].
- [69] P. Cowperthwaite, V. Villar, D. Scolnic, and E. Berger, “Lsst target-of-opportunity observations of gravitational-wave events: Essential and efficient,” *The Astrophysical Journal* **874** no. 1, (2019) 88.
- [70] F. Morawski, M. Bejger, E. Cuoco, and L. Petre, “Anomaly detection in gravitational waves data using convolutional autoencoders,” *Machine Learning: Science and Technology* **2** no. 4, (2021) 045014.

- [71] T. D. Gebhard, N. Kilbertus, I. Harry, and B. Schölkopf, “Convolutional neural networks: A magic bullet for gravitational-wave detection?,” *Physical Review D* **100** no. 6, (2019) 063015.
- [72] P. G. Krastev, K. Gill, V. A. Villar, and E. Berger, “Detection and parameter estimation of gravitational waves from binary neutron-star mergers in real ligo data using deep learning,” *Physics Letters B* **815** (2021) 136161.
- [73] M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf, “Real-time gravitational wave science with neural posterior estimation,” *Physical review letters* **127** no. 24, (2021) 241103.
- [74] C. Christ, B. Nord, K. Gozman, and K. Ottenbreit, “Fast identification of strong gravitational lenses with deep learning on FPGAs and other heterogeneous computing devices,” in *American Astronomical Society Meeting Abstracts #235*, vol. 235 of *American Astronomical Society Meeting Abstracts*, p. 303.03. Jan., 2020.
- [75] K. Zhang and J. S. Bloom, “deepCR: Cosmic Ray Rejection with Deep Learning,” *ApJ* **889** no. 1, (Jan., 2020) 24, arXiv:1907.09500 [astro-ph.IM].
- [76] E. Bertin and S. Arnouts, “SExtractor: Software for source extraction.,” *A&AS* **117** (June, 1996) 393–404.
- [77] N. A. Maksimova, L. H. Garrison, D. J. Eisenstein, B. Hadzhiyska, S. Bose, and T. P. Satterthwaite, “ABACUSSUMMIT: a massive set of high-accuracy, high-resolution N-body simulations,” *MNRAS* **508** no. 3, (Dec., 2021) 4017–4037, arXiv:2110.11398 [astro-ph.CO].
- [78] T. Ishiyama, F. Prada, A. A. Klypin, M. Sinha, R. B. Metcalf, E. Jullo, B. Altieri, S. A. Cora, D. Croton, S. de la Torre, D. E. Millán-Calero, T. Oogi, J. Ruedas, and C. A. Vega-Martínez, “The Uchuu simulations: Data Release 1 and dark matter halo concentrations,” *MNRAS* **506** no. 3, (Sept., 2021) 4210–4231, arXiv:2007.14720 [astro-ph.CO].
- [79] F. Villaescusa-Navarro, C. Hahn, E. Massara, A. Banerjee, A. M. Delgado, D. K. Ramanah, T. Charnock, E. Giusarma, Y. Li, E. Ally, A. Brochard, C. Uhlemann, C.-T. Chiang, S. He, A. Pisani, A. Obuljen, Y. Feng, E. Castorina, G. Contardo, C. D. Kreisch, A. Nicola, J. Alsing, R. Scoccimarro, L. Verde, M. Viel, S. Ho, S. Mallat, B. Wandelt, and D. N. Spergel, “The Quijote Simulations,” *ApJS* **250** no. 1, (Sept., 2020) 2, arXiv:1909.05273 [astro-ph.CO].
- [80] D. Nelson, V. Springel, A. Pillepich, V. Rodriguez-Gomez, P. Torrey, S. Genel, M. Vogelsberger, R. Pakmor, F. Marinacci, R. Weinberger, L. Kelley, M. Lovell, B. Diemer, and L. Hernquist, “The IllustrisTNG simulations: public data release,” *Computational Astrophysics and Cosmology* **6** no. 1, (May, 2019) 2, arXiv:1812.05609 [astro-ph.GA].

- [81] Y. Ni, T. Di Matteo, S. Bird, R. Croft, Y. Feng, N. Chen, M. Tremmel, C. DeGraf, and Y. Li, “The ASTRID simulation: the evolution of Supermassive Black Holes,” *arXiv e-prints* (Oct., 2021) arXiv:2110.14154, arXiv:2110.14154 [astro-ph.GA].
- [82] F. Villaescusa-Navarro, D. Anglés-Alcázar, S. Genel, D. N. Spergel, R. S. Somerville, R. Dave, A. Pillepich, L. Hernquist, D. Nelson, P. Torrey, D. Narayanan, Y. Li, O. Philcox, V. La Torre, A. Maria Delgado, S. Ho, S. Hassan, B. Burkhardt, D. Wadekar, N. Battaglia, G. Contardo, and G. L. Bryan, “The CAMELS Project: Cosmology and Astrophysics with Machine-learning Simulations,” *ApJ* **915** no. 1, (July, 2021) 71, arXiv:2010.00619 [astro-ph.CO].
- [83] S. He, Y. Li, Y. Feng, S. Ho, S. Ravanbakhsh, W. Chen, and B. Póczos, “Learning to predict the cosmological structure formation,” *Proceedings of the National Academy of Science* **116** no. 28, (July, 2019) 13825–13832, arXiv:1811.06533 [astro-ph.CO].
- [84] R. Alves de Oliveira, Y. Li, F. Villaescusa-Navarro, S. Ho, and D. N. Spergel, “Fast and Accurate Non-Linear Predictions of Universes with Deep Learning,” *arXiv e-prints* (Nov., 2020) arXiv:2012.00240, arXiv:2012.00240 [astro-ph.CO].
- [85] N. Kaushal, F. Villaescusa-Navarro, E. Giusarma, Y. Li, C. Hawry, and M. Reyes, “NECOLA: Towards a Universal Field-level Cosmological Emulator,” *arXiv e-prints* (Nov., 2021) arXiv:2111.02441, arXiv:2111.02441 [astro-ph.CO].
- [86] D. Kodi Ramanah, T. Charnock, F. Villaescusa-Navarro, and B. D. Wandelt, “Super-resolution emulator of cosmological simulations using deep physical models,” *MNRAS* **495** no. 4, (July, 2020) 4227–4236, arXiv:2001.05519 [astro-ph.CO].
- [87] Y. Li, Y. Ni, R. A. C. Croft, T. Di Matteo, S. Bird, and Y. Feng, “AI-assisted superresolution cosmological simulations,” *Proceedings of the National Academy of Science* **118** no. 19, (May, 2021) 2022038118, arXiv:2010.06608 [astro-ph.CO].
- [88] Y. Ni, Y. Li, P. Lachance, R. A. C. Croft, T. Di Matteo, S. Bird, and Y. Feng, “AI-assisted superresolution cosmological simulations - II. Halo substructures, velocities, and higher order statistics,” *MNRAS* **507** no. 1, (Oct., 2021) 1021–1033, arXiv:2105.01016 [astro-ph.CO].
- [89] F. Villaescusa-Navarro, J. Ding, S. Genel, S. Tonnesen, V. La Torre, D. N. Spergel, R. Teyssier, Y. Li, C. Heneka, P. Lemos, D. Anglés-Alcázar, D. Nagai, and M. Vogelsberger, “Cosmology with one galaxy?,” *arXiv e-prints* (Jan., 2022) arXiv:2201.02202, arXiv:2201.02202 [astro-ph.CO].
- [90] F. Villaescusa-Navarro, D. Anglés-Alcázar, S. Genel, D. N. Spergel, Y. Li, B. Wandelt, A. Nicola, L. Thiele, S. Hassan, J. M. Zorrilla Matilla, D. Narayanan, R. Dave, and M. Vogelsberger, “Multifield Cosmology with Artificial Intelligence,” *arXiv e-prints* (Sept., 2021) arXiv:2109.09747, arXiv:2109.09747 [astro-ph.CO].

[91] F. Villaescusa-Navarro, S. Genel, D. Angles-Alcazar, L. Thiele, R. Dave, D. Narayanan, A. Nicola, Y. Li, P. Villanueva-Domingo, B. Wandelt, D. N. Spergel, R. S. Somerville, J. M. Zorrilla Matilla, F. G. Mohammad, S. Hassan, H. Shao, D. Wadekar, M. Eickenberg, K. W. K. Wong, G. Contardo, Y. Jo, E. Moser, E. T. Lau, L. F. Machado Poletti Valle, L. A. Perez, D. Nagai, N. Battaglia, and M. Vogelsberger, “The CAMELS Multifield Dataset: Learning the Universe’s Fundamental Parameters with Artificial Intelligence,” *arXiv e-prints* (Sept., 2021) arXiv:2109.10915, arXiv:2109.10915 [cs.LG].

[92] A. Ćiprijanović, D. Kafkes, K. Downey, S. Jenkins, G. N. Perdue, S. Madireddy, T. Johnston, G. F. Snyder, and B. Nord, “DeepMerge - II. Building robust deep learning algorithms for merging galaxy identification across domains,” *MNRAS* **506** no. 1, (Sept., 2021) 677–691, arXiv:2103.01373 [astro-ph.IM].

[93] A. Ćiprijanović, D. Kafkes, G. Snyder, F. J. Sánchez, G. N. Perdue, K. Pedro, B. Nord, S. Madireddy, and S. M. Wild, “DeepAdversaries: Examining the Robustness of Deep Learning Models for Galaxy Morphology Classification,” *arXiv e-prints* (Dec., 2021) arXiv:2112.14299, arXiv:2112.14299 [cs.LG].

[94] G. Csurka, “A comprehensive survey on domain adaptation for visual applications,” in *Domain Adaptation in Computer Vision Applications*, pp. 1–35. Springer International Publishing, Cham, 2017.

[95] E. Naghib, P. Yoachim, R. J. Vanderbei, A. J. Connolly, and R. L. Jones, “A Framework for Telescope Schedulers: With Applications to the Large Synoptic Survey Telescope,” *AJ* **157** no. 4, (Apr., 2019) 151, arXiv:1810.04815 [astro-ph.IM].

[96] S. Lampoudi, E. Saunders, and J. Eastman, “An Integer Linear Programming Solution to the Telescope Network Scheduling Problem,” *arXiv e-prints* (Mar., 2015) arXiv:1503.07170, arXiv:1503.07170 [astro-ph.IM].

[97] M. Cranmer, P. Melchior, and B. Nord, “Unsupervised Resource Allocation with Graph Neural Networks,” in *34th Conference on Neural Information Processing Systems*. 6, 2021. arXiv:2106.09761 [cs.LG].

[98] ZeMax: An Ansys Company, “Zemax.” <https://www.zemax.com/>.

[99] COMSOL, Inc., “Comsol.” <https://www.comsol.com/>.

[100] Y. Guo, L. Zhong, L. Min, J. Wang, Y. Wu, K. Chen, K. Wei, and C. Rao, “Adaptive optics based on machine learning: a review,” *Opto-Electronic Advances* (2022) 200082–1.

[101] B. Cohen, B. Nord, and t. . A. in prep.

[102] G. Côté, J.-F. Lalonde, and S. Thibault, “Extrapolating from lens design databases using deep learning,” *Opt. Express* **27** no. 20, (Sep, 2019) 28279–28292. <http://www.opticsexpress.org/abstract.cfm?URI=oe-27-20-28279>.

[103] G. Côté, J.-F. Lalonde, and S. Thibault, “Introducing a dynamic deep neural network to infer lens design starting points,” in *Current Developments in Lens Design and Optical Engineering XX*, R. B. Johnson, V. N. Mahajan, and S. Thibault, eds., vol. 11104, pp. 8 – 14, International Society for Optics and Photonics. SPIE, 2019. <https://doi.org/10.1117/12.2528866>.

[104] K. Höschel and V. Lakshminarayanan, “Genetic algorithms for lens design: a review,” *Journal of Optics* **48** (12, 2018) .

[105] T. Yang, G.-F. Jin, and J. Zhu, “Automated design of freeform imaging systems,” *Light: Science & Applications* **6** no. 10, (Oct, 2017) e17081–e17081. <https://doi.org/10.1038/lsa.2017.81>.

[106] H. Sun, A. V. Dalca, and K. L. Bouman, “Learning a Probabilistic Strategy for Computational Imaging Sensor Selection,” *arXiv e-prints* (Mar., 2020) arXiv:2003.10424, arXiv:2003.10424 [eess.IV].

[107] **MODE** Collaboration, A. G. Baydin *et al.*, “Toward Machine Learning Optimization of Experimental Design,” *Nucl. Phys. News* **31** no. 1, (2021) 25–28.

[108] K. Bundy, K. Westfall, N. MacDonald, R. Kupke, M. Savage, C. Poppett, A. Alabi, G. Becker, J. Burchett, P. Capak, *et al.*, “Fobos: A next-generation spectroscopic facility at the w. m. keck observatory,” *arXiv preprint arXiv:1907.07195* (2019) .

[109] B. D. Nord *et al.*, “SPOKES: an End-to-End Simulation Facility for Spectroscopic Cosmological Surveys,” *Astron. Comput.* **15** (2016) 1–15, arXiv:1602.01480 [astro-ph.IM].

[110] M. Cranmer, A. Sanchez-Gonzalez, P. Battaglia, R. Xu, K. Cranmer, D. Spergel, and S. Ho, “Discovering Symbolic Models from Deep Learning with Inductive Biases,” *arXiv e-prints* (June, 2020) arXiv:2006.11287, arXiv:2006.11287 [cs.LG].

[111] S.-M. Udrescu and M. Tegmark, “Ai feynman: A physics-inspired method for symbolic regression,” *Science Advances* **6** no. 16, (2020) , <https://advances.sciencemag.org/content/6/16/eaay2631.full.pdf>. <https://advances.sciencemag.org/content/6/16/eaay2631>.

[112] N. Perraudeau, M. Defferrard, T. Kacprzak, and R. Sgier, “DeepSphere: Efficient spherical convolutional neural network with HEALPix sampling for cosmological applications,” *Astronomy and Computing* **27** (Apr., 2019) 130, arXiv:1810.12186 [astro-ph.CO].

[113] B. Dai and U. Seljak, “Translation and Rotation Equivariant Normalizing Flow (TRENF) for Optimal Cosmological Analysis,” *arXiv e-prints* (Feb., 2022) arXiv:2202.05282, arXiv:2202.05282 [astro-ph.CO].

[114] A. K. Saydjari and D. P. Finkbeiner, “Equivariant Wavelets: Fast Rotation and Translation Invariant Wavelet Scattering Transforms,” *arXiv e-prints* (Apr., 2021) arXiv:2104.11244, arXiv:2104.11244 [cs.CV].

[115] S. Cheng, Y.-S. Ting, B. Ménard, and J. Bruna, “A new approach to observational cosmology using the scattering transform,” *MNRAS* **499** no. 4, (Dec., 2020) 5902–5914, arXiv:2006.08561 [astro-ph.CO].

[116] G. Valogiannis and C. Dvorkin, “Towards an Optimal Estimation of Cosmological Parameters with the Wavelet Scattering Transform,” arXiv:2108.07821 [astro-ph.CO].

[117] V. A. Villar, G. Hosseinzadeh, E. Berger, M. Ntampaka, D. O. Jones, P. Challis, R. Chornock, M. R. Drout, R. J. Foley, R. P. Kirshner, R. Lunnan, R. Margutti, D. Milisavljevic, N. Sanders, Y.-C. Pan, A. Rest, D. M. Scolnic, E. Magnier, N. Metcalfe, R. Wainscoat, and C. Waters, “SuperRAENN: A Semisupervised Supernova Photometric Classification Pipeline Trained on Pan-STARRS1 Medium-Deep Survey Supernovae,” *ApJ* **905** no. 2, (Dec., 2020) 94, arXiv:2008.04921 [astro-ph.HE].

[118] R. Morgan, B. Nord, K. Bechtol, S. J. González, E. Buckley-Geer, A. Möller, J. W. Park, A. G. Kim, S. Birrer, M. Aguena, J. Annis, S. Bocquet, D. Brooks, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, R. Cawthon, L. N. da Costa, T. M. Davis, J. De Vicente, P. Doel, I. Ferrero, D. Friedel, J. García-Bellido, M. Gatti, E. Gaztanaga, G. Giannini, D. Gruen, R. A. Gruendl, G. Gutierrez, D. L. Hollowood, K. Honscheid, D. J. James, K. Kuehn, N. Kuropatkin, M. A. G. Maia, R. Miquel, A. Palmese, F. Paz-Chinchón, M. E. S. Pereira, A. Pieres, A. A. Plazas Malagón, K. Reil, A. Roodman, E. Sanchez, M. Smith, E. Suchyta, M. E. C. Swanson, G. Tarle, and C. To, “DeepZipper: A Novel Deep Learning Architecture for Lensed Supernovae Identification,” *arXiv e-prints* (Dec., 2021) arXiv:2112.01541, arXiv:2112.01541 [astro-ph.CO].

[119] J. Allam, Tarek and J. D. McEwen, “Paying Attention to Astronomical Transients: Photometric Classification with the Time-Series Transformer,” *arXiv e-prints* (May, 2021) arXiv:2105.06178, arXiv:2105.06178 [astro-ph.IM].

[120] K. Boone, “ParSNIP: Generative Models of Transient Light Curves with Physics-enabled Deep Learning,” *AJ* **162** no. 6, (Dec., 2021) 275, arXiv:2109.13999 [astro-ph.IM].

[121] J. Fluri, T. Kacprzak, A. Refregier, A. Lucchi, and T. Hofmann, “Cosmological parameter estimation and inference using deep summaries,” *Phys. Rev. D* **104** no. 12, (Dec., 2021) 123526, arXiv:2107.09002 [astro-ph.CO].

[122] S. Mishra-Sharma, “Inferring dark matter substructure with astrometric lensing beyond the power spectrum,” *Machine Learning: Science and Technology* **3** no. 1, (Mar., 2022) 01LT03, arXiv:2110.01620 [astro-ph.CO].

- [123] J. Fluri, T. Kacprzak, A. Lucchi, A. Schneider, A. Refregier, and T. Hofmann, “A Full w CDM Analysis of KiDS-1000 Weak Lensing Maps using Deep Learning,” *arXiv e-prints* (Jan., 2022) arXiv:2201.07771, arXiv:2201.07771 [astro-ph.CO].
- [124] N. Krachmalnicoff and M. Tomasi, “Convolutional neural networks on the HEALPix sphere: a pixel-based algorithm and its application to CMB data analysis,” *A&A* **628** (Aug., 2019) A129, arXiv:1902.04083 [astro-ph.IM].
- [125] W. Ha, C. Singh, F. Lanusse, S. Upadhyayula, and B. Yu, “Adaptive wavelet distillation from neural networks through interpretations,” *arXiv e-prints* (July, 2021) arXiv:2107.09145, arXiv:2107.09145 [stat.ML].
- [126] C. Singh, W. Ha, F. Lanusse, V. Boehm, J. Liu, and B. Yu, “Transformation Importance with Applications to Cosmology,” *arXiv e-prints* (Mar., 2020) arXiv:2003.01926, arXiv:2003.01926 [stat.ML].
- [127] P. Lemos, N. Jeffrey, M. Cranmer, S. Ho, and P. Battaglia, “Rediscovering orbital mechanics with machine learning,” *arXiv e-prints* (Feb., 2022) arXiv:2202.02306, arXiv:2202.02306 [astro-ph.EP].
- [128] D. Wadekar, L. Thiele, F. Villaescusa-Navarro, J. C. Hill, M. Cranmer, D. N. Spergel, N. Battaglia, D. Anglés-Alcázar, L. Hernquist, and S. Ho, “Augmenting astrophysical scaling relations with machine learning : application to reducing the SZ flux-mass scatter,” arXiv:2201.01305 [astro-ph.CO].
- [129] A. M. Delgado, D. Wadekar, B. Hadzhiyska, S. Bose, L. Hernquist, and S. Ho, “Modeling the galaxy-halo connection with machine learning,” arXiv:2111.02422 [astro-ph.CO].
- [130] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [131] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML’16, p. 1050–1059. JMLR.org, 2016.
- [132] A. Graves, “Practical variational inference for neural networks,” in *Proceedings of the 24th International Conference on Neural Information Processing Systems*, NIPS’11, p. 2348–2356. Curran Associates Inc., Red Hook, NY, USA, 2011.
- [133] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS’17, p. 6405–6416. Curran Associates Inc., Red Hook, NY, USA, 2017.
- [134] J. a. Caldeira and B. Nord, “Deeply Uncertain: Comparing Methods of Uncertainty Quantification in Deep Learning Algorithms,” *Sci. Technol.* **2** (2021) 015002, arXiv:2004.10710 [cs.LG].

- [135] K. Cranmer, J. Brehmer, and G. Louppe, “The frontier of simulation-based inference,” *Proc. Nat. Acad. Sci.* **117** no. 48, (2020) 30055–30062, arXiv:1911.01429 [stat.ML].
- [136] J. Alsing, B. Wandelt, and S. Feeney, “Massive optimal data compression and density estimation for scalable, likelihood-free inference in cosmology,” *Mon. Not. Roy. Astron. Soc.* **477** no. 3, (2018) 2874–2885, arXiv:1801.01497 [astro-ph.CO].
- [137] T. Charnock, G. Lavaux, and B. D. Wandelt, “Automatic physical inference with information maximizing neural networks,” *Phys. Rev. D* **97** no. 8, (Apr., 2018) 083004, arXiv:1802.03537 [astro-ph.IM].
- [138] J. Alsing and B. Wandelt, “Nuisance hardened data compression for fast likelihood-free inference,” *Mon. Not. Roy. Astron. Soc.* **488** no. 4, (2019) 5093–5103, arXiv:1903.01473 [astro-ph.CO].
- [139] J. Alsing, T. Charnock, S. Feeney, and B. Wandelt, “Fast likelihood-free cosmology with neural density estimators and active learning,” *Mon. Not. Roy. Astron. Soc.* **488** no. 3, (2019) 4440–4458, arXiv:1903.00007 [astro-ph.CO].
- [140] A. Diaz Rivero and C. Dvorkin, “Flow-Based Likelihoods for Non-Gaussian Inference,” *Phys. Rev. D* **102** no. 10, (2020) 103507, arXiv:2007.05535 [astro-ph.CO].
- [141] N. Jeffrey and B. D. Wandelt, “Solving high-dimensional parameter inference: marginal posterior densities & Moment Networks,” in *34th Conference on Neural Information Processing Systems*. 11, 2020. arXiv:2011.05991 [stat.ML].
- [142] T. L. Makinen, T. Charnock, J. Alsing, and B. D. Wandelt, “Lossless, scalable implicit likelihood inference for cosmological fields,” *JCAP* **11** no. 11, (2021) 049, arXiv:2107.07405 [astro-ph.CO].
- [143] J. Hermans, A. Delaunoy, F. Rozet, A. Wehenkel, and G. Louppe, “Averting A Crisis In Simulation-Based Inference,” *arXiv e-prints* (Oct., 2021) arXiv:2110.06581, arXiv:2110.06581 [stat.ML].
- [144] J. Hermans, N. Banik, C. Weniger, G. Bertone, and G. Louppe, “Towards constraining warm dark matter with stellar streams through neural simulation-based inference,” *Mon. Not. Roy. Astron. Soc.* **507** no. 2, (2021) 1999–2011, arXiv:2011.14923 [astro-ph.GA].
- [145] S. Mishra-Sharma, “Inferring dark matter substructure with astrometric lensing beyond the power spectrum,” *Mach. Learn. Sci. Tech.* **3** no. 1, (2022) 01LT03, arXiv:2110.01620 [astro-ph.CO].

- [146] D. Norman, “Can Big Data Lead an Inclusion Revolution?,” in *Advancing Astronomy for All: ASP 2018*, G. Schultz, J. Barnes, and L. Shore, eds., vol. 524 of *Astronomical Society of the Pacific Conference Series*, p. 199. Nov., 2019.
- [147] D. Norman, T. Brandt, Z. Berta-Thompson, N. Lewandowska, K. Knierman, N. Chanover, J. Whitley, A. Venkatesan, K. Coble, A. Burgasser, M. Lemoine-Busserolle, K. Cruz, C. S. Moore, K. Osei-Sarfo, and S. Kannappan, “Tying Research Funding to Progress on Inclusion,” in *Bulletin of the American Astronomical Society*, vol. 51, p. 14. Sept., 2019. arXiv:1907.13202 [astro-ph.IM].
- [148] C. J. Lintott *et al.*, “Galaxy Zoo : Morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey,” *Mon. Not. Roy. Astron. Soc.* **389** (2008) 1179–1189, arXiv:0804.4483 [astro-ph].
- [149] **LSST Dark Energy Science, LSST Transient, Variable Stars** Collaboration, R. Hložek *et al.*, “Results of the Photometric LSST Astronomical Time-series Classification Challenge (PLAsTiCC),” arXiv:2012.12392 [astro-ph.IM].
- [150] R. F. Kizilcec, A. J. Saltarelli, J. Reich, and G. L. Cohen, “Closing global achievement gaps in moocs,” *Science* **355** no. 6322, (2017) 251–252.
- [151] L. Hiscott, “Careers guide 2022: What do physics graduates do?,” 2022. <https://www.aps.org/careers/guidance/guide/index.cfm>.
- [152] D. Norman, K. Cruz, V. Desai, B. Lundgren, E. Bellm, F. Economou, A. Smith, A. Bauer, B. Nord, C. Schafer, *et al.*, “The growing importance of a tech savvy astronomy and astrophysics workforce,” *arXiv preprint arXiv:1910.08376* (2019) .