

# REFORMS: Reporting Standards for Machine Learning Based Science

Sayash Kapoor<sup>1</sup> Emily Cantrell Kenny Peng Thanh Hien Pham  
Christopher A. Bail Odd Erik Gundersen Jake M. Hofman Jessica Hullman  
Michael A. Lones Momin M. Malik Priyanka Nanayakkara Russell A. Poldrack  
Inioluwa Deborah Raji Michael Roberts Matthew J. Salganik Marta Serra-Garcia  
Brandon M. Stewart Gilles Vandewiele Arvind Narayanan

August 15, 2023

## Abstract

Machine learning (ML) methods are proliferating in scientific research. However, the adoption of these methods has been accompanied by failures of validity, reproducibility, and generalizability. These failures can hinder scientific progress, lead to false consensus around invalid claims, and undermine the credibility of ML-based science. ML methods are often applied and fail in similar ways across disciplines. Motivated by this observation, our goal is to provide clear reporting standards for ML-based science. Drawing from an extensive review of past literature, we present the REFORMS checklist (**R**eporting **S**tandards **F**or **M**achine Learning **B**ased **S**cience). It consists of 32 questions and a paired set of guidelines. REFORMS was developed based on a consensus of 19 researchers across computer science, data science, mathematics, social sciences, and biomedical sciences. REFORMS can serve as a resource for researchers when designing and implementing a study, for referees when reviewing papers, and for journals when enforcing standards for transparency and reproducibility.

## Introduction

ML methods are being widely adopted for scientific research [1–11]. Compared to older statistical methods, they offer increased predictive accuracy [1], the ability to process large amounts of data [12], and the ability to use different types of data for scientific research, such as text, images, and video [7]. However, the rapid uptake of ML methods has been accompanied by concerns of validity, reproducibility, and generalizability [13–19]. There are several reasons for concern. Performance evaluation is notoriously tricky in ML [20–23]. ML code tends to be complex and as yet lacks standardization [24, 25], leading to a lack of computational reproducibility [26]. Subtle pitfalls arise from the differences between explanatory and predictive modeling [27]. The hype and over-optimism about commercial AI may spill over into scientific research [28]. In addition, publication biases that have led to past reproducibility crises [29] are also present in ML research [30, 31]. If left unchecked, these

flaws can lead to a feedback loop of overoptimism since non-replicable findings are cited more than replicable ones [32]. There is an urgent need to systematically address errors in ML-based science rather than finding errors in individual studies after publication.

In this paper, we focus on a specific subset of ML applications: ML-based science. We use this term to refer to research that makes a scientific claim using the performance of an ML model as evidence. For example, Salganik et al. [33] use ML methods to investigate the predictability of life outcomes. This contrasts with ML methods research, which involves improving widely applicable ML methods instead of making scientific claims using ML methods. In the next section, we clarify the distinctions between ML methods research and ML-based science and outline the scope of the paper in greater detail. Box 1 summarizes this discussion.

One promising way to detect and prevent errors in scientific research is by improving reporting standards [34–36]. Clear expectations for using ML methods can allow researchers and referees to spot errors early. Despite the use of ML methods across disciplines, there are no widely applicable best practices for reporting the design, implementation, and evaluation of ML-based science. This leads to different, and often no fixed reporting standards in each field adopting ML methods. As a result, common failure modes in using ML methods recur across disciplines [37, 38].

In this paper, we introduce a set of reporting standards for ML-based science, with the goal of preventing known but common errors that occur when ML is used in scientific research. To that end, we review the literature on best practices and common errors in ML-based science. We introduce the REFORMS checklist and guidelines for reporting ML-based science, which consists of 32 items across 8 modules (Table 2). We accompany REFORMS with a detailed set of guidelines to set expectations for each item (Appendix B).

Checklists have been adopted in many scientific fields [34–36, 39], and they have been impactful in improving reporting practices [40, 41]. In 2014, the U.S. National Institutes of Health (NIH) created principles to improve reproducibility and rigor, endorsed by several journals. One item was

<sup>1</sup>Corresponding author: sayashk@princeton.edu. The latest version of the paper and the appendices are available at <https://reforms.cs.princeton.edu>

the creation of reporting standards and checklists for journals [42]. The EQUATOR network collects reporting guidelines for health research and includes over 500 checklists [43]. Several checklists have been proposed in ML methods research [31, 44, 45].

REFORMS differs from this large body of past work in crucial ways. First, past checklists for scientific research are field or method-specific. For example, the *CLAIM* checklist [34] provides best practices for reporting research on AI in medical imaging. As a result, many items in the checklist do not apply to other scientific fields adopting ML methods. Since ML methods are being rapidly adopted across fields, and research that uses ML methods suffers from similar errors, we aimed to make the REFORMS checklist field-agnostic. We selected items that broadly apply to fields that use ML methods. Second, past checklists for ML methods research focus on common errors in developing ML methods. But these errors differ from the ones that arise in scientific research (See Box 1 for the distinction between ML methods research and ML-based science). For instance, Pineau et al.’s checklist [31] does not include questions about the distribution about which the claims are made, since ML methods research often focuses on improving a model’s performance on a benchmark dataset [46]. In contrast, clearly specifying the distribution of interest is core to a scientific claim. Still, past work in both scientific research and ML methods research has helped inform our checklist.

We present findings from an extensive review of past research on best practices and common shortcomings relevant to each of the 8 checklist modules. Alongside the checklist, we release guidelines paired with each item in the checklist. Similar guidelines have been included in past efforts at establishing reporting standards. For instance, the STROBERDS statement [36] included an “explanation and elaboration” document to clarify expectations for each item in the reporting standards. Our guidelines aim to increase the usability of the REFORMS checklist and to provide pointers to best practices for ML-based science.

The REFORMS checklist can help address common failure modes that lead to errors in ML-based science. To guard against corners being cut due to time and publication pressures, the standards provide a set of clear expectations. To aid researchers new to ML-based science, the guidelines for each item identify resources and relevant past literature. While no single document would be enough to familiarize researchers with all the nuances of ML-based science, our hope is that the guidelines can be one useful pedagogical resource. Finally, there are many steps involved in successfully reporting ML-based science. It is hard to keep all of these items in mind when writing up a study. REFORMS provides all 32 items in one document to prevent omission errors.

### Box 1: What is ML-based science?

ML-based science refers to scientific research that uses machine learning (ML) models to contribute to scientific knowledge. This includes making predictions, conducting measurements, or performing other tasks that help answer scientific questions.

ML-based science could use a variety of ML techniques such as supervised learning, unsupervised learning, reinforcement learning, and deep learning. The research should be geared toward answering a scientific question of interest.

**Exclusions.** Not all ML research and applications qualify. For example, ML methods research and predictive analytics fall outside the scope. Similarly, quantitative research not using ML methods, such as explanatory modeling and simulations, do not qualify as ML-based science.

**Applicability.** Our checklist for ML-based science will be more useful for some types of research than others. For instance, research on the use of ML for predictions may benefit more than research on using ML for search tasks within vast and complex spaces.

### The scope of our claims: ML-based science

We define ML-based science as scientific research that uses the performance of a machine learning model as evidence for a scientific claim. This includes, but is not limited to, making predictions, conducting measurements, or performing other tasks that contribute to the body of scientific knowledge. This definition has two parts: first, the research should be geared toward answering a scientific question of interest. This means that other types of ML research and applications do not fit under the umbrella of ML-based science:

- **ML methods research.** Research focusing on developing and refining machine learning methods, such as a typical NeurIPS paper, does not constitute ML-based science. While such work does contribute new ML methods, the main objective is not to solve specific scientific problems of interest about generalizable populations. (Still, elements of the REFORMS checklist could be valuable to ML methods research. This is particularly the case when these newly developed methods are evaluated on benchmarks that directly influence their application in scientific contexts despite not being representative of real problems.)

- **Predictive analytics.** Many real-world applications of ML models emphasize predictive accuracy but are not conducted to gain scientific insights. For example, social media platforms use ML to predict if a user will click on ads millions of times a day. There are many differences between predictive analytics and ML-based science. In predictive analytics, the population or distribution of interest may not be clearly defined. The relationships found using ML models only need to hold within a company or organization, and need not generalize. Relative accuracy is more important than absolute accuracy numbers since the decision to deploy a model has often already been made—so the only question being answered by the modeling activity is *which* of a given set of models should be deployed in production. And verifying if the predictions later come true is often easy. This feedback is a better test for whether an application works as intended (compared to using a checklist).

Second, the research should utilize ML methods. By this, we refer to a variety of techniques including, but not limited to, supervised learning, unsupervised learning, reinforcement learning, and deep learning. This even includes foundation models [47], which are a type of ML model, although our focus in this paper is more general.<sup>1</sup> Consequently, other types of quantitative research not employing these methods do not fall under the umbrella of ML-based science.

- **Explanatory modeling.** Such research utilizes traditional statistical methods for improving our understanding of real-world phenomena rather than predicting outcomes [49]. This is an example of quantitative scientific research that does not use ML methods. While explanatory modeling suffers from many shortcomings [13], errors in this domain are subtly different compared to ML-based science, which means much of our checklist does not apply to explanatory modeling research.
- **Simulation.** Similarly, physics-based models and simulation methods are sometimes evaluated iteratively by output and by their fidelity with existing theory [50, 51], and do not involve machine learning or fitting models to data—except potentially as a tool to understand and summarize the simulation results. As a result, many items in the REFORMS checklist do not apply. Note that some algorithms used in machine learning, like Markov Chain Monte Carlo, are examples of numerical simulations; but these are used only to fit models to data, which is different from simulating an underlying phenomenon of interest.

<sup>1</sup>There are several challenges with using proprietary foundation models for scientific research. For example, they are non-deterministic and can change without adequate notice [48]. This could lead to hard-to-resolve shortcomings in computational reproducibility. Foundation models are one example of a grey area in the definition of ML-based science. In such cases, we leave it up to the authors and referees to decide if the REFORMS checklist is useful.

We call research at the intersection of ML methods and quantitative science “ML-based science”. Salganik et al.’s work on the Fragile Families Challenge [33] is an example of this category. They used machine learning to predict children’s life outcomes and answer scientific questions about the predictability of outcomes studied by sociologists. We discuss many other examples in our review.

Note that our checklist will likely be more helpful for some types of ML-based science than others. For instance, it is likely to be more useful for predictive modeling compared to research that uses machine learning methods for search tasks in vast and complex spaces, such as the search for new materials or new phases of matter [52]. In cases where verifying the result of an ML-based experiment is easy (for instance, verifying the properties of a new drug in a lab), our checklist might be less useful than such verification, though it could help ensure the validity of the experiments prior to verification.

## Methods

To develop the REFORMS checklist, we started with a focus on steps in a canonical ML pipeline, drew from previous checklists used in other domains, and went through a consensus process with all authors, involving multiple rounds of feedback and a virtual discussion. Table 1 lists the modules in our checklist and the corresponding stages in the ML pipeline. For each module, we focus on three goals: (1) establishing the scientific claim and its relationship to the ML modeling process; (2) providing an overview of the best practices and common shortcomings in building ML models correctly; and (3) enabling the verification of the results by an independent researcher. In other words, we aim to decrease the likelihood of errors of interpretation (goal 1) or execution (goal 2), and to make it easier for independent researchers to spot errors (goal 3). Box 2 outlines these goals in more detail.

To build on previous efforts at improving the reporting quality of research, we used three past checklists to ensure our coverage of important items in reporting an ML model. Pineau et al. [31] provided the checklist used alongside papers submitted to NeurIPS 2020, a prominent ML methods conference. Collins et al. [39] provided the *TRIPOD* checklist for prediction models in health research. Mongan et al. [34] introduced the *CLAIM* checklist for AI models in clinical imaging. We chose these checklists because they covered diverse modeling approaches and were applicable in different settings (ML methods research, models for individual diagnosis and prognosis, and ML for medical imaging, respectively).

### Box 3: How authors, referees, and journals can use REFORMS

Our reporting standards can help improve the quality of ML-based science in multiple ways.

- **Authors** can self-regulate by using the REFORMS to identify errors and preemptively address concerns about using ML methods in their paper [41]. This can also help increase the credibility of their paper, especially in fields that are newly adopting ML methods. We expect that REFORMS will be useful to authors throughout the study—during conceptualization, implementation, and communication of the results (See Table 1 for a list of checklist modules corresponding to these goals). The checklist can be included as part of the supplementary materials released alongside a paper. The guidelines can help authors learn how to correctly apply the REFORMS checklist in their own work and introduce them to underlying theories of evidence.
- **Referees** can use REFORMS to determine whether a study they are reviewing falls short. If they have concerns about a study, they can ask researchers to include the filled-out checklist in a revised version. For example, Roberts et al. [15] use the CLAIM checklist [34] to filter papers for a systematic review based on compliance with the checklist.
- **Journals** can require authors to submit a checklist along with their papers to set reporting standards for ML-based science. Similar checklists are in place in a number of journals [53, 54]; however, they are usually used for specific disciplines rather than for methods that are prevalent across disciplines [55]. Since ML-based science is proliferating across disciplines, REFORMS offers a method-specific (rather than discipline-specific) intervention.

Note that the REFORMS checklist is additive to field-specific norms. It is not a replacement for existing requirements within fields, such as pre-registration, ethics reviews, or using discipline-specific checklists for other parts of the research process. It might seem burdensome to ask researchers to adhere to another set of standards, but our work was born out of painful necessity: studies across fields have repeatedly found reproducibility errors in ML-based science [38], and in the absence of a systematic intervention, this is likely to worsen.

**Consensus process for developing the REFORMS checklist.** Once we had an initial set of items, authors met virtually for a discussion. One of the main outcomes of the discussion was the need for a paired set of guidelines alongside each item to clarify how these items should be reported, which we discuss in more detail below. Then, the authors collaboratively edited the checklist to choose commonly applicable items across disciplines. We paid close attention to usability: to decrease the time and cognitive load that using the checklist would entail, we removed items that were

too specific and would apply to a small subset of ML-based science. Finally, the authors independently flagged unclear items to improve the quality of the reporting standards.

A recurring theme in our conversation was making the reporting standards easy to use. To that end, we developed a set of accompanying guidelines to help researchers understand the motivation for each section and clarify what is expected for each item in REFORMS. The guidelines (Appendix B) are based on our review of past literature for items in the checklist (this review is presented next). We include references to key prior work to help onboard researchers new to using ML methods. This includes a mix of peer-reviewed scientific research that details best practices, as well as resources from the ML methods community that outline common shortcomings and ways to address them. Crucially, we do not take prescriptive stances on matters of ongoing methodological debate. Instead, we present best practices to minimize and detect known types of errors in ML-based science.

**Organization of the paper.** In the remainder of the paper, we present the REFORMS reporting standards. They comprise eight modules based on the stages of an ML-based science study (see Table 1). For each module, we motivate why items in this module are important to address in ML-based science. For each item, we include expectations about what it means to address the item sufficiently. In our review, we draw from

Stage of scientific study	Section of the checklist
Study design	Study goals (Module 1) Computational reproducibility (Module 2)
Data collection and preparation	Data quality (Module 3) Data preprocessing (Module 4)
Modeling	Modeling decisions (Module 5)
Evaluation	Data leakage (Module 6) Metrics and uncertainty quantification (Module 7)
Scope and limitations	Generalizability and limitations (Module 8)

**Table 1.** Stages of ML-based science and corresponding checklist sections.

Section	Item
<b>Study goals</b>	1a. Population or distribution about which the scientific claim is made
	1b. Motivation for choosing this population or distribution ( <i>1a.</i> )
	1c. Motivation for the use of ML methods in the study
<b>Computational reproducibility</b>	2a. Dataset used for training and evaluating the model along with link or DOI to uniquely identify the dataset
	2b. Code used to train and evaluate the model and produce the results reported in the paper along with link or DOI to uniquely identify the version of the code used
	2c. Description of the computing infrastructure used
	2d. README file which contains instructions for generating the results using the provided dataset and code
	2e. Reproduction script to produce all results reported in the paper
<b>Data quality</b>	3a. Source(s) of data, separately for the training and evaluation datasets (if applicable), along with the time when the dataset(s) are collected, the source and process of ground-truth annotations, and other data documentation
	3b. Distribution or set from which the dataset is sampled (i.e., the sampling frame)
	3c. Justification for why the dataset is useful for the modeling task at hand
	3d. The outcome variable of the model, along with descriptive statistics (split by class for a categorical outcome variable) and its definition.
	3e. Sample size and outcome frequencies
	3f. Percentage of missing data, split by class for a categorical outcome variable
	3g. Justification for why the distribution or set from which the dataset is drawn ( <i>3b.</i> ) is representative of the one about which the scientific claim is being made ( <i>1a.</i> )
<b>Data preprocessing</b>	4a. Identification of whether any samples are excluded with a rationale for why they are excluded
	4b. How impossible or corrupt samples are dealt with
	4c. All transformations of the dataset from its raw form ( <i>3a.</i> ) to the form used in the model, for instance, treatment of missing data and normalization. Preferably through a flow chart
<b>Modeling</b>	5a. Detailed descriptions of all models trained
	5b. Justification for the choice of model types implemented
	5c. Method for evaluating the model(s) reported in the paper, including details of train-test splits or cross-validation folds
	5d. Method for selecting the model(s) reported in the paper
	5e. For the model(s) reported in the paper, specify details about the hyperparameter tuning
	5f. Justification that model comparisons are against appropriate baselines
<b>Data leakage</b>	6a. Justification that pre-processing (Section 4) and modeling (Section 5) steps only use information from the training dataset (and not the test dataset)
	6b. Methods to address dependencies or duplicates between the training and test datasets (e.g. different samples from the same patients are kept in the same dataset partition)
	6c. Justification that each feature or input used in the model is legitimate for the task at hand and does not lead to leakage
<b>Metrics and uncertainty</b>	7a. All metrics used to assess and compare model performance (e.g., accuracy, AUROC etc.). Justify that the metric used to select the final model is suitable for the task
	7b. Uncertainty estimates (e.g., confidence intervals, standard deviations), details of how these are calculated
	7c. Justification for the choice of statistical tests (if used) and a check for the assumptions of the statistical test
<b>Generalizability and limitations</b>	8a. Evidence of external validity
	8b. Contexts in which the authors do not expect the study's findings to hold

**Table 2.** The REFORMS checklist for ML-based science. Some items are shortened for brevity; see Appendix A for a complete template. Alongside each item, authors should report the section or page number where the item is reported. Some items in the REFORMS checklist could be hard to report for specific studies. Instead of requiring strict adherence for each item, authors and referees should decide which items are relevant for a study and how details can be reported better. To that end, we hope that the checklist can offer a useful starting point for authors and referees working on ML-based science.

past literature on best practices and common errors in ML-based science. We also occasionally draw on literature about science with traditional statistical methods, as best practices

and shortcomings are shared in many aspects of ML-based science and other quantitative science. In Appendix A, we provide a template for the REFORMS checklist. In Appendix

## Box 2: Goals of the REFORMS checklist

**Goal 1: Establish the scientific claim and its relation to the ML task.** A key feature of our checklist, distinguishing it from those used in ML methods research, is its focus on using ML to support scientific claims. In such research, it is necessary to establish the intended scientific claim as well as how the performance of the ML model supports that claim. For example, it is necessary to state the scientific claim's population of interest and then justify why the dataset used in the ML task represents this population. This should be compared to ML methods research, where the performance of a model is often itself the claim.

**Goal 2: Ensure that the ML task is executed correctly and that the performance is reported in sufficient detail.** To establish that the performance of the specified ML model supports the intended scientific claim, it is necessary to ensure that the performance of the model is calculated correctly. There are many ways in which the performance of a model can be misleading. For example, a common error in ML research is evaluating a model on data it was trained on, resulting in overly optimistic results. Additionally, carefully reporting uncertainty is necessary to interpret model performance correctly.

**Goal 3: Enable an independent scientist to verify results.** Finally, our checklist is designed to help ensure that all resources and descriptions needed for verifying a study are provided alongside the paper. Thus, our checklist helps ensure that independent researchers can understand and evaluate a given study.

These goals are not intended to be disjoint and often support one another. They can help orient the reader when navigating our checklist, and they reveal how our checklist is tailored to ML-based science.

B, we distill the guidelines for filling out the checklist as a standalone document. In Appendix C, we present a table with additional details on the references used in our review.

## Module 1: Study goals

This section focuses on stating a study's goals. This is motivated by recent research which shows that reporting study goals in adequate depth and clarity is not trivial or common [56]. Studies that appear to ask the same research question may actually have subtle differences in their questions which lead to substantially different findings [57].

**1a) Population or distribution about which the scientific claim is made.** The "population of interest" is the group to which the researchers intend the findings of the study to generalize. This is typically broader than the sampling frame and sample, which are discussed in Module 3. Defining the population of interest is important because it shapes the article's conclusions, places boundaries on those conclusions, and provides the basis for metrics of significance and uncertainty that derive from the concept of sampling from a population [56, 58, 59]. Unfortunately, research articles do not always clearly state their population of interest [56, 60, 61].

**1b) Motivation for choosing this population or distribution.** The choice of a particular population of interest may be motivated by pure scientific interest or by a need for applied knowledge. Explaining the motivation for studying the population of interest helps the reader understand the importance of the study and contextualize its results.

We acknowledge that the motivation for choosing a particular population of interest may arise from what data is available. The development of the motivation depends on whether researchers followed a deductive, inductive, or iterative approach [62]. In a deductive approach, researchers begin with a theory and a population of interest, then select or create a dataset based on the data's ability to test that theory for that population. In an inductive approach, researchers begin with a dataset, then determine what research questions and populations of interest that dataset can address. In an iterative approach, researchers iterate between data collection, data analysis, and theory until they develop a hypothesis, and then collect additional data to test that hypothesis [62]. The deductive approach is the most widely accepted approach in the scientific community, but Grimmer et al. argue that there is also great value in inductive and iterative approaches, and that it is important to communicate which approach was used [62].

**1c) Motivation for the use of ML methods in the study.** The research questions asked in ML-based studies often differ from those asked with traditional statistical methods. Breiman famously argued that there are "two cultures" in statistical modeling: the "data modeling" culture and the "algorithmic modeling" culture. In the "data modeling" culture, which is aligned with traditional statistical methods, researchers' focus is on estimating the parameters of a function that is meant to meaningfully represent the process by which input data produces output data in the world. For example, researchers might ask, "what is the relationship between household income and the likelihood of experiencing clinical depression?," and estimate the magnitude and direction of that relationship using a logistic regression model. In the "algorithmic modeling" culture, which aligns with much of ML-based research, the focus is on building a model that reliably maps input data to output data. In this culture, the

parameters do not necessarily need to provide a faithful and interpretable description of patterns in the world; rather, the goal is to accurately predict output data when given a new sample of input data that is separate from the original data sample in which the model was trained. For example, researchers might ask, "given all the data available about a person in a particular dataset, how accurately can we predict whether that person will experience clinical depression?," and test a variety of models to find the one with the best predictive accuracy. Breiman's proposition was that more scientific research should use "algorithmic modeling" [63]. Several responses to Breiman have argued for the value of merging the "two cultures" or moving iteratively between them [49, 64–68].

Since non-ML ("data modeling") methods are currently the standard approach in many scientific disciplines, explaining the motivation for using ML methods will help readers better understand a study's goals. Considering the differences and similarities between Breiman's "two cultures" may be useful to researchers when motivating the use of ML methods. Additionally, many recent articles provide guidance on the value of ML methods for science [1, 49, 64, 69–74].

## Module 2: Computational reproducibility

Computational reproducibility refers to the ability of an independent researcher to get the same results as reported in a paper or manuscript. It is an essential part of computational research [75].

Computational reproducibility can help independent researchers evaluate the findings in a paper and verify if they hold up under scrutiny. The availability of reproducibility materials has led to several errors being spotted [76–81]. Conversely, if the code and data for reproducing all results in a study are not available, identifying the precise sources of errors in a study becomes hard [82–84].

**Current computational reproducibility standards fall short.** Some journals require authors to make their computational reproducibility materials available post-publication without requiring these materials at the time of publication. However, such measures can miss the mark. Stodden et al. [85] attempted to contact the authors of 204 papers published in the journal *Science* to obtain reproducibility materials. Only 44% of authors responded. Similarly, Gabelica et al. [86] studied papers published in 333 open-access journals indexed on BioMed Central in January 2019. Of the 1,792 papers that claimed they would share data upon request, 1,669 did not share the data. That is, they were unable to get the data for 93% of the papers. This indicates the importance of requiring computational materials at the time of publication rather than at the authors' discretion later.

Vasilevsky et al. [87] studied the data-sharing policies at 318 biomedical journals. They found that almost a third of these journals had no data-sharing policies in place. Even the journals that did have data-sharing policies did not have clear guidelines for authors to comply with their policies.

ML methods research has also struggled to ensure computational reproducibility. Gundersen and Kjensmo [18] systematically analyzed 400 papers that were published at leading conferences. In addition to code and dataset availability, they evaluated the documentation of methods in a paper's text, for instance, whether the experimental setup is described. They found that *none* of the 400 papers satisfied all of their reproducibility criteria, and in general, papers only satisfied 20–30% of the criteria.

Pineau et al. [31] found that only around half of the papers submitted to NeurIPS 2018, a leading ML conference, contained the code and data needed to reproduce results. To improve reproducibility, they introduced a checklist that was used in NeurIPS 2019. In the checklist, including reproducibility materials was optional but recommended. Still, after the checklist was introduced, over 75% of the papers included reproducibility materials along with the submissions. Similar checklists have become the standard at several ML conferences [88–90].

While ML methods research differs from ML-based science in its goals and practices, we can learn from these experiences to emphasize the importance of computational reproducibility in ML-based science.

**Ensuring computational reproducibility in ML-based science is challenging.** ML methods used in scientific research can be complex and often require several packages and dependencies. This makes computational reproducibility challenging [26]. Liu and Salganik describe their experiences ensuring computational reproducibility while editing a special issue in *Socius* on the Fragile Families Challenge [24]. The Fragile Families Challenge was a prediction competition where multiple participants tried to predict children's life outcomes on the same dataset [33]. The special issue published papers based on a few of the resulting models. Liu and Salganik wanted to ensure that the code and data alongside every publication were verified to ensure that they produced the same results as those presented in the paper. Despite spending 13 months working on achieving computational reproducibility and exchanging dozens of emails with the authors, they were unable to verify the computational reproducibility of all papers. Even preliminary steps, like installing the correct versions of each package, were non-trivial when a large number of packages were used in a study. They could eventually verify the computational reproducibility of 7 of the 12 papers. They published the rest with the code and data available at the time of publication.

**Interventions adopted by journals.** Journals have adopted several measures to improve computational reproducibility [53, 54, 91]. The Transparency and Openness Promotion (TOP) standards introduced by the Open Science Foundation [92] have a few sections that focus on computational reproducibility in scientific research. They are divided into three levels: Level 1 requires stating if the computational reproducibility materials are available. If they are, authors should provide details about how to access them. Level 2 requires the materials to be available in a trusted repository at the time of publication. Level 3 requires the materials to be verified by the journal to ensure they generate the results reported in the paper prior to publication.

Several social science journals use the Data and Code Availability Standards for computational reproducibility [93]. Other journals have taken additional measures to verify if the code is correct. For instance, *Nature Methods* conducts code reviews in addition to peer reviews for papers that provide computational artifacts [94].

In our checklist, we include the basic details that can enable independent researchers to verify the computational reproducibility of a result. Our checklist requires information about the dataset, code, computing environment, documentation for how to get the results in a study, and a reproduction script to automatically run the code and generate results. We acknowledge that computational reproducibility is hard and that some items in this module are more challenging compared to others. For instance, providing a reproduction script is not always possible. There may be other challenges. For instance, it is not always possible to release private datasets. We provide some options to address these challenges below, such as providing a synthetic imitation of the data when the real data cannot be publicly released.

**2a) Dataset.** Peng et al. [95] and Nosek et al. [92] highlight the importance of citing datasets with permanent links to clarify which version of a dataset is used in a study. If an original dataset is provided alongside a study, documentation for the dataset is also important. For instance, authors can include data dictionaries [96] or datasheets [97]. Such documentation should report basic details about the properties and format of the data.

Some datasets could contain sensitive information and cannot be publicly released. To address this, authors have previously released synthetic datasets when working with sensitive data. For instance, Obermeyer et al. released a synthetic dataset alongside their study of racial bias in a healthcare algorithm [98]. They used the *synthpop* package in R to generate the synthetic data [99].

**2b) Code.** Similarly, it is important to report the exact version of code used for running the experiments and producing the results in a paper [100]. Authors can accomplish this by providing a DOI, commit tag (for instance, from

code repositories such as GitHub, GitLab, or BitBucket), or other documentation to precisely identify the version of the code used to train and evaluate the model and produce the results reported in the paper. Note that using archiving systems that provide permanent identifiers for the code used, like Dataverse, is likely to aid long-term reproducibility [101, 102].

**2c) Computing environment.** Different computational experiments require different amounts of computing resources. To help readers understand the precise computing requirements for reproducing the study, authors should report details about the hardware (CPU, RAM, disk space), software (operating system, programming language, version number for each package used), and computing resources (time taken to generate the results) used to generate their results. Stodden & Miguez [103] provide best practices to document computing infrastructure.

**2d) Documentation.** Good documentation helps researchers unfamiliar with a project by walking them through the steps of setting up and running the code provided, starting from environment requirements and installation, to examples of usage and expected results [104, 105].

**2e) Reproduction script.** A script to produce all results reported in the paper using the reproducibility materials can significantly reduce the time it takes for an independent researcher to reproduce the results reported in a study. Reproduction scripts can download packages with the version numbers needed to run the code, set the right dependencies, download and store datasets in the correct location, set up the computing environment, and run the code to produce the results reported in the paper.

Authors can implement such scripts in several ways, such as using a bash script [106] or using an online reproducibility platform such as CodeOcean [107]. Note that this is a high bar for computational reproducibility. In some cases, it might not be possible to provide a script that would allow an independent researcher to reproduce all results—for instance, if the analysis is run on an academic high-performance computing cluster, or if the dataset does not allow for programmatic download.

### Module 3: Data quality

This module helps readers and referees understand and evaluate the quality of the data used in the study. Using poor quality data or data that is not suitable for answering a research question can lead to results that are meaningless or misleading.

**3a) Data source(s).** Describing a study’s data source(s) allows readers to evaluate the data’s strengths and weaknesses and to judge whether the data are appropriate for the study’s goals. Most studies provide descriptions of their



data source, but those descriptions sometimes lack important details [108–110]. Additionally, the quality of reporting about ground-truth annotation methods in ML-based science varies widely [111]. To ensure a minimum level of information about data sources is reported, our checklist asks researchers to report when, where, and how data were collected, and how ground-truth annotations were performed on the dataset, if applicable.

**3b) Sampling frame.** A sampling frame is a list of people or units from which a sample is drawn. Due to practical limitations, the sampling frame in many studies does not include all members of the target population. It is important to describe a study’s sampling frame so that readers understand the boundaries of the study’s sample and how that sample relates to the target population (discussed further under checklist item 3g: *Dataset for evaluation is representative*).

Some research papers do not provide a clear description of their sampling frame or eligibility criteria for inclusion in the sample [60, 108, 112]. Furthermore, one review article found that when papers stated information about their sampling frame or eligibility criteria was available in a prior publication, the prior publication was not always accessible and the relevant information was often extremely difficult or impossible to find [112].

**3c) Justification for why the dataset is useful for the modeling task.** Our checklist asks researchers to justify why the dataset is useful for the modeling task because the appropriateness of a data source will depend on the research question. For example, while biased or incomplete data is inappropriate for some research questions, such data can work well for other research questions as long as the researcher understands how these shortcomings impact the analysis and communicates the limitations [62]. A broad claim like “this is the best dataset available on this topic” does not help readers understand the strengths and weaknesses of the data for the study’s research question; researchers should be specific about *why* the dataset is well-suited to the question.

Modern ML-based research often relies on repurposed data sources, which are sometimes termed “big data:” for example, social media data, digital trace data, or digital administrative records [113]. Salganik describes ten common characteristics of “big data” that result in differing strengths and weaknesses compared to traditional data sources [113]. Researchers who are using ML with repurposed data can use these ten characteristics as a guide when justifying why the data source is appropriate for their research goals and identifying shortcomings of the data.

**3d) Outcome variable.** Our checklist asks researchers to report how their outcome variable is defined. The outcome

or target variable is the quantity that the model is used to predict, detect, classify, or estimate.

The outcome variable is typically an empirical proxy for an unobservable theoretical construct [114]. For example, researchers might pose a question about the construct “academic performance,” and use grade point average as measured in school administrative data as the empirical proxy for this construct. The outcome variable is usually not a perfect match for the theoretical construct it represents. Thus, in order to allow readers to evaluate a paper’s claims, authors should describe precisely how their outcome variable is measured, and note any ways in which this outcome variable might not align with the associated construct. This is especially important because mismatches between variables and the constructs they are purported to represent can create fairness issues [115].

Our checklist also asks for descriptive statistics about the outcome variable. Reviews of prior literature have found that descriptive statistics are not always sufficiently reported [116, 117]. Descriptive statistics about the outcome variable help readers to understand the context being studied and to identify concerns related to rare or skewed values.

**3e) Number of samples in the dataset.** Reporting sample size is important because a study must have sufficient sample size to achieve its objectives. Some research objectives can be achieved with small to moderate-sized samples: for example, detecting large differences between groups. Other objectives generally require large samples: for example, studying rare events, studying heterogeneity, or detecting small differences [113].

Note that there can be downsides to large sample sizes. When the sampling frame is unrepresentative, increasing the sample size can shrink confidence intervals without decreasing bias in estimates, thus giving false confidence [118]. Furthermore, if a study exposes participants to any level of risk, larger samples may magnify harms [113].

Scientific literature is generally consistent about reporting total sample size [60, 119]. However, reporting on sample size for subgroups or sample size after attrition in longitudinal research is less consistent [60]. To ensure clear reporting, our checklist asks that in addition to the total sample size, researchers who are conducting a classification task report the number of samples in each class. This follows recent calls for more granular details about data, such as by Gebru et al. [97]. We also ask researchers to distinguish between the number of individuals in the dataset and the number of rows in the dataset, in cases where an individual can appear in more than one row.

**3f) Missingness.** Missing data is highly prevalent in many research domains and can impact the results [120–122]. It is important for researchers to report the prevalence of missing data in their dataset and to specify how they handled

missingness [120]. Missingness is particularly important to address carefully when it is non-random [123]. Extensive literature across multiple fields has established that research articles frequently provide insufficient information about the presence and handling of missing data [60, 108, 121, 124–128].

Checklist item *3f* focuses on reporting the prevalence of missing data. Reporting how missing data is handled is covered in item *4c*: *Data transformations*.

**3g) Dataset for evaluation is representative.** This item asks researchers to justify that their sample is representative of the target population defined in Module 1. Representativeness is important for the study’s ability to generalize from the sample to the target population. Lack of representativeness in the sampling process is sometimes underreported; for example, one review of past literature found underreporting of information about selection bias [60].

Probability sampling is a common approach for achieving representativeness. However, probability sampling is not always necessary. Due to coverage errors and nonresponse in probability sampling, the differences between probability sampling and non-probability sampling are not always as large as they first appear [113]. For studies that use non-probability sampling, researchers may be able to make a reasonable argument for representativeness by comparing sample characteristics with population characteristics [58]. Researchers can also use statistical methods to adjust for non-probability sampling (or to adjust for errors in probability sampling), such as post-stratification, sample matching, propensity score weighting, and calibration [113].

In some studies, the dataset will not be representative of the target population. Reasons a sample might fail to be representative of the target population depend on the type of data collected. For example, Salganik describes three types of representation error in survey data [113], and Grimmer et al. describe four sources of bias in sample selection for text analysis [62]. A non-representative sample is okay for certain research goals. For example, Salganik argues that research that aims to make out-of-sample generalizations generally requires representative data, but research that aims to make within-sample comparisons can be well-suited to non-representative data [113]. Concerns about non-representativeness should be noted under checklist item *8a*: *Evidence of external validity*.

## Module 4: Data preprocessing

Preprocessing is the series of steps taken to convert the dataset from its rawest available form into the final form used in the modeling process. This includes data selection

(i.e., selecting a set of samples from the dataset to be included in the modeling process) as well as other transformations of the data, such as imputing missing data and normalizing feature values.

Our checklist focuses on two broad components of preprocessing: first, the subset of data to consider (i.e., which rows of a dataset are considered), and second, the transformations that are subsequently applied to the data (i.e., how entries of a dataset might be altered). Each of these components have implications on the scope and validity of resulting scientific claims, and are essential for ensuring the reproducibility of the results. As discussed in Module 3, preprocessing methods are often not specified in papers.

**4a) Excluded data and rationale.** Researchers might exclude some samples from the dataset—for instance, to remove outliers or to only focus on certain subsets. Thus, the resulting scientific claims should be made in relation to the particular subset of a dataset that is ultimately used. This type of preprocessing is closely related to our discussion on data quality (Module 3) and generalizability (Module 8). This item underscores the importance of reporting the specific subset of the dataset used, in addition to details about the overall data.

Hofman et al. [129] note how the choice of the subset can significantly affect the performance of the resulting model. A specific example they use is the prediction of the reach (“cascade size”) of a social media post as a function of the poster’s past success. They show that the threshold of popularity used to determine the subset of posts considered plays a large role in influencing the predictability of cascade size. The scientific claim—regarding the predictability of the success of a post—depends on what data is included and excluded. We ask authors to justify why the particular subset of data used was chosen.

**4b) How impossible or corrupt samples are dealt with.** A dataset may contain erroneous or undesirable data points. Data may be impossible (e.g., a person whose height is recorded as 10 feet) or corrupted (e.g., a survey response filled out by a bot). Different techniques have been developed to detect such cases [130, 131]. Attempting to filter such data can be important to ensure the dataset represents the intended population of the study. Data detected as impossible or corrupt may be removed or transformed—our checklist asks authors to report such steps in items *4a* and *4c* respectively.

**4c) Data transformations.** Once the set of data points to be used is decided upon, researchers could transform the data in various ways: for example, by normalizing or augmenting data, or imputing missing data. Both imputing missing data and under/over-sampling from a subpopulation—if done improperly—can harm the validity of model performance in relation to the scientific claim

made. For example, mean imputation and oversampling must be done separately on each fold of a dataset. Failing to do so can result in overoptimistic results [132].

Specifying preprocessing steps is also important for ensuring the reproducibility of the results. Choices in preprocessing technique can significantly affect the properties of the resulting model, including accuracy and interpretability [133, 134].

## Module 5: Modeling

Researchers make several choices in creating an ML model. The exact specification of the model is important to consider with respect to the particular scientific claim being made. In addition, due to the large number of choices involved in creating an ML model, it is important to report exact details of how an ML model is created—otherwise, reproducibility by independent researchers could be hindered. Raff [135] attempted to reproduce ML results from 255 papers using only the paper’s text (i.e., without using the code accompanying a paper), and found that 93 could not be reproduced.

Modeling choices are also closely related to choices involving evaluation. Indeed, model selection—the process of choosing the model(s) whose results are reported in a study—often depends on the evaluation setup. An improper approach to evaluation or model selection can result in exaggerated performance estimates.

Major machine learning conferences, such as NeurIPS and ICML, incorporate checklists that ask authors to verify that they have included training details within their paper, similar to the specific items we propose below.

A note on terminology: there is no agreement on the use of different terms related to the modeling process. For the purpose of the discussion in this module, we make one conceptual distinction. Items *5a* and *5b* focuses on aspects of the model that are specified prior to training. This includes, for example, the type of model and the loss function. Meanwhile, items *5c*, *5d*, *5e*, and *5f* focus on fitted (trained) models.

**5a) Model description.** Specific details about all models trained are essential for ensuring the reproducibility of a paper. This includes specifying the input and output of a model, the type of model (e.g., Random Forest, Neural Network), and the loss function and algorithm used to train the model.

**5b) Justification for the choice of model types implemented.** With many different possible types of machine learning models to choose from, the choice of types to consider can be dependent on the intended scientific claim of a paper. For example, if a scientific claim aims to establish how predictive a set of features can be (e.g., [33, 136]), then

it may be appropriate to consider a wide range of types—the most important criteria is the resulting accuracy of the learned models. If the scientific claim aims to establish the potential usability of a model in practical settings, there may be additional desiderata. For example, a model that is used in high-stakes settings like clinical decision-making may need to be interpretable or explainable [137], so it may be more appropriate to choose a model type that is likelier to have these properties.

It is not always clear what the best type of model to use is for a given scientific question. For example, there is an active debate on how or if interpretability or explainability should be implemented in the context of different applications (see e.g., [138]). Here, we ask authors to provide reasons for why the set of model types they implemented is appropriate.

The focus of this item is on the types of models considered, but some claims may depend on the choice of a specific learned model (and its performance). The following items address the process of evaluating, selecting, and comparing specific learned models.

**5c) Model evaluation method.** We ask authors to report details about the model evaluation procedure. We include evaluation in this ‘modeling’ module since evaluation is often used as part of the model selection process (see item *5d*). That is, it is common to consider many models and select the best-performing one based on the evaluation setup.

Evaluation of ML models must be done on test data separate from training data. To ensure reproducibility and verify validity, it is necessary to report how data was split and used. We ask authors to report whether models are evaluated using cross-validation, a holdout test set, or an external validation test set. We also ask for the sample size within each split of the data, including the number of samples of each class for classification tasks.

**5d) Model selection method.** There are many possible learned models that can arise depending on specific choices in the modeling process. Even holding the type of model fixed, differences can arise due to the specific choice of hyperparameters, for example. Researchers should report how the final model(s) reported in a paper are selected.

A common goal of model selection is to select the model with the best performance on a hold-out set, but improper model selection can result in misleading performance [139, 140].

Testing multiple models on the holdout set and choosing the one with the best performance on the holdout set can result in an over-optimistic estimate of performance. Neunhoffer and Sternberg [79] consider this type of error in the context of political science, showing how a study on civil war prediction failed to reproduce because of improper model selection.

Model selection is not limited to choosing the model with the highest accuracy. Multiple models can achieve the same accuracy as one another, yet make different predictions for the same individuals and possess different characteristics regarding issues like fairness and interpretability. This phenomenon is known as "model multiplicity." Black et al. provide guidance on opportunities and concerns that arise from model multiplicity [141].

**5e) Hyperparameter selection.** Training procedures are often dependent on the choice of model hyperparameters, such as regularization weight, the number of training epochs, or the learning rate for a model. The choice of hyperparameters is part of the model selection process.

In the context of natural language processing research, Dodge et al. [142] show that details about hyperparameter search impact the resulting accuracy. More resources devoted to searching over hyperparameters can improve performance significantly. Islam et al. [143] show that due to variance in performance depending on hyperparameters, specific choices of hyperparameters can be misleading when interpreting results, especially when comparing different models. One method may perform better than another with one set of hyperparameters while performing worse given a different set of hyperparameters.

Finally, note that the degree to which hyperparameter optimization affects results can vary depending on the model type [144]. This may be a reason to prefer one model type over another.

**5f) Appropriate baselines.** It can be important to compare the performance of a model to baselines, especially when the scientific claim argues that a particular ML method can outperform existing approaches on a task. To clearly establish the performance of an ML model, it is necessary to detail how baseline models were trained, and how the baseline methods are optimized. For example, if the baselines were not chosen using the same model selection methods, this could result in the baselines being "weak." As a result, the apparent benefits of the new method may be misleading. Sculley et al. [23] and Lin [145] detail examples in which machine learning models were compared to weak baselines.

Lones [146] discusses how to compare models more broadly. See also item 7c: *Appropriate statistical tests* for a discussion on the use of statistical testing to compare models.

## Module 6: Data leakage

Leakage is a spurious relationship between the features and the target variable that arises as an artifact of the data collection, sampling, pre-processing, or modeling steps. For example, normalizing features in the training and test data together leads to leakage since information about the test data features is included in the training data [147].

Data leakage is a common error in the use of ML methods. Epic, a U.S. healthcare technology company, released a sepsis prediction model in hospitals nationwide. However, one of the features used in the model was whether a patient had been prescribed antibiotics. This is an error because antibiotics would typically be prescribed after the diagnosis of sepsis, so they act as a proxy for the outcome variable. Consequently, the model's performance was inflated due to having access to information it would not have in a real-world scenario [148].

Leakage has caused widespread reproducibility errors in ML-based science. In a survey of leakage across ML-based science, Kapoor and Narayanan found that leakage affects hundreds of papers across 17 fields [38].

We ask authors to justify that their study does not suffer from major sources of leakage, which we elaborate on in this section. Kapoor and Narayanan [38] offer model info sheets to help detect and prevent leakage before publication. Our checklist focuses on the three main types of leakage in ML-based science found in their survey.

**6a) Train-test separation is maintained.** When information from the test set is used during the training process, it leads to over-optimistic performance results due to data leakage.

Not using a held-out test set is a textbook error in ML [149]. Still, it is widespread. For example, Poldrack et al. [150] find that of the 100 neuropsychiatry studies that claimed to predict patient outcomes, 45 only reported in-sample statistical fit as evidence for predictive accuracy.

There can be other, more subtle variations of this error. For example, if the train-test split occurs after any of the other pre-processing or modeling steps (Modules 4 and 5), this also results in leakage. Vandewiele et al. [81] found overoptimistic results in 21 papers that claimed to predict the risk of pre-term births. These papers suffered from the same error: oversampling data before partitioning it into training and test sets. This resulted in the test set becoming artificially similar to the training set, and led to exaggerated performance claims across the literature on pre-term risk prediction.

**6b) Dependencies or duplicates between datasets.** In some cases, samples in the dataset might have dependencies. For example, a clinical dataset might have many samples from the same patient. Oner et al. [151] find that when image data from the same patient is present in training and test sets, it leads to overoptimistic results. Similarly, for time-series forecasting models, randomly splitting a time-series dataset into training and test sets is likely to lead to overoptimism, because the training data has information "from the future" [152].

In such cases, the train-test split or cross-validation (CV) split should take these dependencies into account—for instance, by including all samples from each patient in the same train-test split or CV fold. There are several ways to avoid such dependencies. Bergmeir & Benítez [153] find that blocked CV for time series evaluation deals with temporal autocorrelation. Hammerla & Plotz [154] demonstrate how “neighborhood bias” can affect data recordings close in time. They introduce “meta-segmented CV” to deal with such dependencies. Roberts et al. [155] describe block CV strategies for a number of structures with dependencies, including temporal, spatial, and hierarchical dependencies.

Duplicates in the datasets can also spread across training and test sets if the dataset is split randomly. This should be avoided, as it leaks information across the train-test split. Roberts et al. [15] outline this error with Frankenstein datasets: datasets that combine multiple other sources of data can end up using the same data twice—for instance, if two datasets rely on the same underlying data source are combined into a larger dataset.

**6c) Feature legitimacy.** If any of the features used in a model is a proxy for the outcome, this can result in leakage. Filho et al. [156] found that a prominent paper on hypertension prediction [157] suffered from data leakage due to illegitimate features. The model included the use of anti-hypertensive drugs as a feature to predict hypertension. But this feature is not available when predicting if a new patient suffers from hypertension in a clinical setting, so it artificially inflates the performance of the ML model. Similarly, Epic’s sepsis prediction tool used “antibiotics” as a feature to predict if someone would get sepsis [148].

This type of leakage is more likely when there are a large number of features, due to the increased likelihood of including one or more illegitimate features. The sheer volume of features can make it challenging to scrutinize each one for potential leakage.

## Module 7: Metrics and uncertainty quantification

The performance of ML models is key to the scientific claims of interest. Since there are many possible choices authors can make when choosing performance metrics, it is important to reason why the metrics used are appropriate for the task [8]. Additionally, communicating and reasoning about uncertainty is important but under-reported. For example, Simmonds et al. [37] find that studies often do not report the various kinds of uncertainty in the modeling process.

We ask authors to report their performance metrics and uncertainty estimates for those metrics in enough detail to enable a judgment about whether they made valid choices for evaluating the model’s performance.

The checklist requires authors to detail how they evaluate the performance of their model on its own and in relation to baselines. In addition, in this section, we ask authors to reason about their measurement of model performance and uncertainty in relation to their scientific claim.

**7a) Performance metrics used.** There are many possible ways to measure the performance of an ML model (see Table 4 in [8] for an overview). Certain metrics can be misleading or inappropriate. For example, accuracy might not be suitable to measure the performance of an ML model in the presence of heavy class imbalance: when most data points have positive labels, it is easy to obtain high accuracy simply by predicting positive for all cases [158]. Other common metrics, like AUC, also have limitations [159].

The proper choice of metric is often influenced by the particular application being studied. In some domains, certain errors may be more costly than others. For example, false positives are more costly than false negatives in email spam detection [160]. We ask authors to justify the use of a particular performance metric in relation to the scientific claim.

**7b) Uncertainty estimates.** Uncertainty in the performance of ML models can arise in many ways. For example, it could arise from randomness in the training data, evaluation data, or in the training process itself. Uncertainty is important to capture when evaluating the strength of a scientific claim. Because a dataset represents only a finite sample from a population, there is uncertainty in this sampling. Since it is generally infeasible to obtain several independent samples, methods such as bootstrapping (evaluating model performance by re-sampling from the same dataset) are often used to generate confidence intervals [139].

In a systematic review of uncertainty quantification across seven scientific fields, Simmonds et al. [161] found that scientific fields differ in what kinds of uncertainty they report. They propose best practices and a checklist to help researchers account for uncertainty.

**7c) Appropriate statistical tests.** Comparing the performance of different models can be key to some scientific claims. For example, a scientific claim may argue that one ML method outperforms others. Statistical testing is one tool to evaluate such differences in performance. Raschka [139] gives an overview of statistical tests for ML models. Still, in some or even most cases, an appropriate and valid statistical test may not be known; this can be an area for further development. The reliance on statistical significance testing has also led to misinterpretations and false conclusions. As a result, reporting uncertainty is better than performing statistical tests alone [162]. Still, if a statistical test must be performed, it should be appropriate for the comparison.

## Module 8: Generalizability and limitations

External validity (or "generalizability") refers to the extent to which the findings from a study's sample apply to the target population, as well as the extent to which the findings apply to other populations, outcomes, and contexts [163, 164].

ML-based science faces a number of threats to external validity [13, 165]. Since studies that use ML methods are often unaccompanied by external (i.e., out-of-distribution) validation [166], it is important to reason about these threats. Additionally, authors are best positioned to identify the boundaries of applicability of their claims in order to prevent misunderstandings about the claims made in their study.

**External validity of three types of claims in ML-based science.** We distinguish among three types of external validity, corresponding to three types of claims made in ML-based science.

*External validity of claims about observed patterns.* In some studies, researchers use ML methods to make a claim about the presence of a pattern in the world. For example, Mathur et al. use structural topic modeling to study manipulative tactics in a corpus of U.S. political campaign emails from the 2020 election cycle. One of their findings is that the median active sender of campaign emails uses "sensationalist clickbait" in 37% of those emails [167]. A question about the external validity of this claim might be: is the frequency of sensationalist clickbait in campaign emails similar in other U.S. election cycles?

*External validity of claims about learned models.* A "model" is a function that an algorithm learns when the algorithm is applied to data [168]. In some studies, researchers train an ML model and make a claim about the model's performance when deployed in real-world settings. For example, Dugas et al. build a model to forecast influenza outbreaks at the city level using Google Flu Trends and other readily accessible data, with the goal that the model can be deployed by medical centers to provide warning of upcoming outbreaks. They find that "the model, on the average, predicts weekly influenza cases during 7 out-of-sample outbreaks within 7 cases for 83% of estimates" (p. 1) [169]. A question about the external validity of their claim, which they note in their discussion, is: does the model achieve similar performance in other years? External validity of claims about learned models is more widely known as "domain generalization" or "robustness to distribution shift" and is a well-studied phenomenon in ML [170, 171].

*External validity of claims about learning algorithms.* An "algorithm" is a procedure for learning from data [168]. In some studies, researchers make a claim about the usefulness of ML algorithms in a particular context. For example, Bansak et al. develop an ML algorithm to assign refugees

to resettlement sites by "leverag[ing] synergies between refugee characteristics and resettlement sites." They test the algorithm in retrospective data from the U.S. and Switzerland, and find that their "approach led to gains of roughly 40 to 70%, on average, in refugees' employment outcomes relative to current assignment practices" (p. 1). They claim that governments can use this algorithmic assignment approach to improve resettlement outcomes for refugees. However, unlike the influenza forecasting example above, they do not claim that the specific model they trained can be deployed directly into use [172]. A question about external validity of their claim might be: does algorithmic assignment of refugee resettlement locations work similarly well for other time periods?

**Reporting on external validity falls short in past literature.** Reviews of past literature have found that scientific papers sometimes lack discussion about external validity [60] or lack information about sample demographics that would help readers draw their own conclusions about external validity [173]. Furthermore, many ML-based studies make claims about the ability of a model to deploy in real-world settings, but do not report whether their evaluation sample matches the population in the real-world setting [173], and do not conduct external validation [173, 174].

**8a) Evidence of external validity.** In this checklist item, we ask researchers to discuss the ability to generalize their claims from the sample to the target population and to other populations, outcomes, or contexts. Researchers can use a mix of quantitative and theoretical approaches to make arguments regarding their findings' external validity. They can report quantitative evidence by testing their claims in out-of-distribution data. They can make theoretical arguments about their expectations of external validity by referring to prior literature and reasoning about the level of similarity between contexts [61].

Several threats to the external validity of ML models have been documented in past literature. Finlayson et al. outline external validity failures due to dataset shifts in clinical settings [175]. Hullman et al. discuss the threats to external validity that arise in different phases of an ML research project [13]. Liao et al. outline a taxonomy of evaluation failures in ML, including failures in external validity [166]. Geirhos et al. discuss the phenomenon of shortcut learning in ML models, a phenomenon where models rely on shortcuts (such as the background color in an image) instead of detecting patterns that actually relate to the phenomena of interest [176].

Researchers who make claims about learned models should be aware that even if their model currently generalizes to their target population, that performance may degrade over time due to temporal distribution shift (or "temporal drift"). Causes of temporal drift include changes in technology, changes in population and setting, and changes in behavior

[175]. Concerns about the risk of drift should be communicated in the paper, when applicable.

**8b) Contexts in which the authors do not expect the study's findings to hold.** Clarifying the circumstances under which scientific conclusions or model performance are not expected to hold helps to set clear expectations and avoid unjustified hype. Raji et al. find that flaws in ML models deployed in real-world settings stem in part from a lack of focus on identifying when models are not expected to work [165]. Simons et al. argue that making an explicit "constraints on generality" statement that identifies boundaries of the circumstances where findings are expected to hold has several benefits, including helping to ensure that a study's conclusions accurately reflect its evidence, increasing the likelihood of successful replication, and inspiring follow-up studies that test the findings in new populations [61].

## Conclusion

ML methods present an exciting advance for scientific research. Done right, they can allow researchers to analyze complex data and work with new modalities such as images and video. Yet, recent failures of ML-based science reveal the urgent necessity of improving standards of transparency across fields that use ML methods. Our paper provides a cross-disciplinary bar for reporting standards in ML-based science.

**Supplementary materials.** All appendices are available at <https://reforms.cs.princeton.edu/>.

## References

- [1] Susan Athey and Guido W. Imbens. Machine learning methods that economists should know about. *Annual Review of Economics*, 11(1):685–725, 2019. \_eprint: <https://doi.org/10.1146/annurev-economics-080217-053433>.
- [2] Daniel R. Schrider and Andrew D. Kern. Supervised machine learning for population genetics: A new paradigm. *Trends in Genetics*, 34(4):301–312, April 2018.
- [3] John Joseph Valletta, Colin Torney, Michael Kings, Alex Thornton, and Joah Madden. Applications of machine learning in animal behaviour studies. *Animal Behaviour*, 124:203–220, February 2017.
- [4] R. Iniesta, D. Stahl, and P. McGuffin. Machine learning, statistical learning and the future of biological research in psychiatry. *Psychological Medicine*, 46(12):2455–2465, September 2016. Publisher: Cambridge University Press.
- [5] Scott Tonidandel, Eden B. King, and Jose M. Cortina. Big data methods: Leveraging modern data analytic techniques to build organizational science. *Organizational Research Methods*, 21(3):525–547, July 2018. Publisher: SAGE Publications Inc.
- [6] Tal Yarkoni and Jacob Westfall. Choosing prediction Over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, 12(6):1100–1122, November 2017. Publisher: SAGE Publications Inc.
- [7] Justin Grimmer, Margaret E. Roberts, and Brandon M. Stewart. Machine learning for social science: An agnostic approach. *Annual Review of Political Science*, 24(1):395–419, 2021. \_eprint: <https://doi.org/10.1146/annurev-polisci-053119-015921>.
- [8] Anja K. Leist, Matthias Klee, Jung Hyun Kim, David H. Rehkopf, Stéphane P. A. Bordas, Graciela Muniz-Terrera, and Sara Wade. Mapping of machine learning approaches for description, prediction, and causal inference in the social and health sciences. *Science Advances*, 8(42):eabk1942, October 2022.
- [9] Timothy L. Wiemken and Robert R. Kelley. Machine learning in epidemiology and health outcomes research. *Annual Review of Public Health*, 41(1):21–36, April 2020.
- [10] Hal R. Varian. Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2):3–28, May 2014.
- [11] Sendhil Mullainathan and Jann Spiess. Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2):87–106, May 2017.
- [12] Leo Breiman. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical Science*, 16(3):199–231, August 2001. Publisher: Institute of Mathematical Statistics.
- [13] Jessica Hullman, Sayash Kapoor, Priyanka Nanayakkara, Andrew Gelman, and Arvind Narayanan. The worst of both worlds: A comparative analysis of errors in learning from data in psychology and machine learning. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 335–348, Oxford United Kingdom, July 2022. ACM.
- [14] Andrew L. Beam, Arjun K. Manrai, and Marzyeh Ghassemi. Challenges to the reproducibility of machine learning models in health care. *JAMA*, 323(4):305, January 2020.
- [15] Michael Roberts, Derek Driggs, Matthew Thorpe, Julian Gilbey, Michael Yeung, Stephan Ursprung, Angelica I. Aviles-Rivero, Christian Etmann, Cathal McCague, Lucian Beer, Jonathan R. Weir-McCall, Zhongzhao Teng, Effrossyni Gkrania-Klotsas, James H. F. Rudd, Evis Sala, and Carola-Bibiane Schönlieb. Common pitfalls and recommendations for using machine learning to detect and prognosticate for COVID-19 using chest radiographs and CT scans. *Nature Machine Intelligence*, 3(3):199–217, March 2021. Number: 3 Publisher: Nature Publishing Group.
- [16] Xavier Bouthillier, César Laurent, and Pascal Vincent. Unreproducible research is reproducible. In *International Conference on Machine Learning*, pages 725–734. PMLR, May 2019. ISSN: 2640-3498.
- [17] Matthew B. A. McDermott, Shirly Wang, Nikki Marinsek, Rajesh Ranganath, Luca Foschini, and Marzyeh Ghassemi. Reproducibility in machine learning for health research: Still a ways to go. *Science Translational Medicine*, 13(586), March 2021. Publisher: American Association for the Advancement of Science Section: Perspective.
- [18] Odd Erik Gundersen and Sigbjørn Kjensmo. State of the art: Reproducibility in artificial intelligence. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), April 2018.
- [19] Gaël Varoquaux and Veronika Cheplygina. Machine learning for medical imaging: Methodological failures and recommendations for the future. *NPJ digital medicine*, 5(1):48, April 2022.
- [20] Robin M. Schmidt, Frank Schneider, and Philipp Hennig. Descending through a crowded valley - benchmarking deep learning optimizers. In *Proceedings of the 38th International Conference on Machine Learning*, pages 9367–9376. PMLR, July 2021. ISSN: 2640-3498.
- [21] Xavier Bouthillier, Pierre Delaunay, Mirko Bronzi, Assya Trofimov, Brennan Nichyporuk, Justin Szeto, Naz Sepah, Edward Raff, Kanika Madan, Vikram Voleti, Samira Ebrahimi Kahou, Vincent Michalski, Dmitriy Serdyuk, Tal Arbel, Chris Pal, Gaël Varoquaux, and Pascal Vincent. Accounting for variance in machine learning benchmarks. page 23, 2020.
- [22] Orianna DeMasi, Konrad Kording, and Benjamin Recht. Meaningless comparisons lead to false optimism in medical machine learning. *PLOS ONE*, 12(9):e0184604, September 2017.
- [23] D. Sculley, Jasper Snoek, Alex Wiltschko, and Ali Rahimi. Winner's curse? on pace, progress, and empirical rigor. June 2018.

- [24] David M. Liu and Matthew J. Salganik. Successes and struggles with computational reproducibility: Lessons from the fragile families challenge. *Socius: Sociological Research for a Dynamic World*, 5:237802311984980, January 2019.
- [25] Sören Dittmer, Michael Roberts, Julian Gilbey, Ander Biguri, AIX-COVNET Collaboration, Ian Selby, Anna Breger, Matthew Thorpe, Jonathan R. Weir-McCall, Effrossyni Gkrania-Klotsas, Anna Korhonen, Emily Jefferson, Georg Langs, Guang Yang, Helmut Prosch, Jan Stanczuk, Jing Tang, Judith Babar, Lorena Escudero Sánchez, Philip Teare, Mishal Patel, Marcel Wassin, Markus Holzer, Nicholas Walton, Pietro Lió, Tolou Shadbahr, Evis Sala, Jacobus Preller, James H. F. Rudd, John A. D. Aston, and Carola-Bibiane Schönlieb. Navigating the development challenges in creating complex data systems. *Nature Machine Intelligence*, 5(7):681–686, June 2023.
- [26] Odd Erik Gundersen, Saeid Shamsaliei, and Richard Juul Isdahl. Do machine learning platforms provide out-of-the-box reproducibility? *Future Generation Computer Systems*, 126:34–47, January 2022.
- [27] Jake M. Hofman, Amit Sharma, and Duncan J. Watts. Prediction and explanation in social systems. *Science*, 355(6324):486–488, February 2017.
- [28] John Banja. AI hype and radiology: A plea for realism and accuracy. *Radiology: Artificial Intelligence*, 2(4):e190223, July 2020.
- [29] Valen E. Johnson, Richard D. Payne, Tianying Wang, Alex Asher, and Soutrik Mandal. On the reproducibility of psychological science. *Journal of the American Statistical Association*, 112(517):1–10, January 2017.
- [30] Samuel J. Bell and Onno P. Kampman. Perspectives on machine learning from psychology’s reproducibility crisis, April 2021. arXiv:2104.08878 [cs].
- [31] Joelle Pineau, Philippe Vincent-Lamarre, Koustuv Sinha, Vincent Larivière, Alina Beygelzimer, Florence d’Alché Buc, Emily Fox, and Hugo Larochelle. Improving reproducibility in machine learning research (a report from the NeurIPS 2019 reproducibility program). *The Journal of Machine Learning Research*, 22(1):164:7459–164:7478, July 2022.
- [32] Marta Serra-Garcia and Uri Gneezy. Nonreplicable publications are cited more than replicable ones. *Science Advances*, 7(21):eabd1705, May 2021.
- [33] Matthew J Salganik, Ian Lundberg, Alexander T Kindel, et al. Measuring the predictability of life outcomes with a scientific mass collaboration. *Proceedings of the National Academy of Sciences*, 117(15):8398–8403, 2020.
- [34] John Mongan, Linda Moy, and Charles E. Kahn. Checklist for artificial intelligence in medical imaging (CLAIM): A guide for authors and reviewers. *Radiology: Artificial Intelligence*, 2(2):e200029, March 2020. Publisher: Radiological Society of North America.
- [35] Patrick M Bossuyt, Johannes B Reitsma, David E Bruns, Constan-tine A Gatsonis, Paul P Glasziou, Les Irwig, Jeroen G Lijmer, David Moher, Drummond Rennie, Henrica C W de Vet, Herbert Y Kressel, Nader Rifai, Robert M Golub, Douglas G Altman, Lotty Hooff, Daniël A Korevaar, and Jérémie F Cohen. STARD 2015: An updated list of essential items for reporting diagnostic accuracy studies. *BMJ*, page h5527, October 2015.
- [36] Richard G. White, Avi J. Hakim, Matthew J. Salganik, Michael W. Spiller, Lisa G. Johnston, Ligia Kerr, Carl Kendall, Amy Drake, David Wilson, Kate Orroth, Matthias Egger, and Wolfgang Hladik. Strengthening the Reporting of Observational Studies in Epidemiology for respondent-driven sampling studies: “STROBE-RDS” statement. *Journal of Clinical Epidemiology*, 68(12):1463–1471, December 2015.
- [37] Emily G. Simmonds, Kwaku Peprah Adjei, Christoffer Wold Andersen, Janne Cathrin Hetle Aspheim, Claudia Battistin, Nicola Bulso, Hannah Christensen, Benjamin Cretois, Ryan Cubero, Ivan A. Davidovich, Lisa Dickel, Benjamin Dunn, Etienne Dunn-Sigouin, Karin Dyrstad, Sigurd Einum, Donata Giglio, Haakon Gjerlow, Amelie Godefroidt, Ricardo Gonzalez-Gil, Soledad Gonzalo Cagno, Fabian Grosse, Paul Halloran, Mari F. Jensen, John James Kennedy, Peter Egge Langsaether, Jack H. Laverick, Debora Lederberger, Camille Li, Elizabeth Mandeville, Caitlin Mandeville, Espen Moe, Tobias Navarro Schroder, David Nunan, Jorge Sicacha Parada, Melanie Rae Simpson, Emma Sofie Skarstein, Clemens Spensberger, Richard Stevens, Aneesh Subramanian, Lea Svendsen, Ole Magnus Theisen, Connor Watret, and Robert B. OHara. How is model-related uncertainty quantified and reported in different disciplines? 2022.
- [38] Sayash Kapoor and Arvind Narayanan. Leakage and the reproducibility crisis in ML-based science, July 2022. arXiv:2207.07048 [cs, stat].
- [39] Gary S. Collins, Johannes B. Reitsma, Douglas G. Altman, and Karel GM Moons. Transparent Reporting of a Multivariable Prediction Model for Individual Prognosis or Diagnosis (TRIPOD): The TRIPOD statement. *BMC Medicine*, 13(1):1, January 2015.
- [40] Amy C. Plint, David Moher, Andra Morrison, Kenneth Schulz, Douglas G. Altman, Catherine Hill, and Isabelle Gaboury. Does the CONSORT checklist improve the quality of reports of randomised controlled trials? A systematic review. *The Medical Journal of Australia*, 185(5):263–267, September 2006.
- [41] SeungHye Han, Tolani F. Olonisakin, John P. Pribis, Jill Zupetic, Joo Heung Yoon, Kyle M. Holleran, Kwonho Jeong, Nader Shaikh, Doris M. Rubio, and Janet S. Lee. A checklist is associated with increased quality of reporting preclinical biomedical research: A systematic review. *PLoS ONE*, 12(9):e0183591, September 2017.
- [42] Principles and guidelines for reporting preclinical research. <https://www.nih.gov/research-training/rigor-reproducibility/principles-guidelines-reporting-preclinical-research>, August 2015.
- [43] Reporting guidelines. The EQUATOR Network. <https://www.equator-network.org/reporting-guidelines/>.
- [44] Anna Rogers, Timothy Baldwin, and Kobi Leins. ‘just what do you think you’re doing, Dave?’ A checklist for responsible data use in NLP. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4821–4833, Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- [45] Odd Erik Gundersen, Yolanda Gil, and David W. Aha. On reproducible AI: Towards reproducible research, open science, and digital scholarship in AI publications. *AI Magazine*, 39(3):56–68, September 2018.
- [46] David Donoho. 50 years of data science. *Journal of Computational and Graphical Statistics*, 26(4):745–766, October 2017.
- [47] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, et al. On the opportunities and risks of foundation models, 2022.
- [48] Sayash Kapoor and Arvind Narayanan. OpenAI’s policies hinder reproducible research on language models, 2022.
- [49] Jake M. Hofman, Duncan J. Watts, Susan Athey, Filiz Garip, Thomas L. Griffiths, Jon Kleinberg, Helen Margetts, Sendhil Mullainathan, Matthew J. Salganik, Simine Vazire, Alessandro Vespignani, and Tal Yarkoni. Integrating explanation and prediction in computational social science. *Nature*, 595(7866):181–188, July 2021. Bandiera\_abtest: a Cg\_type: Nature Research Journals Number: 7866 Primary\_atype: Reviews Publisher: Nature Publishing Group Subject\_term: Interdisciplinary studies;Scientific community Subject\_term\_id: interdisciplinary-studies;scientific-community.
- [50] Eric Winsberg. *Science in the age of computer simulation*. University of Chicago Press, 2010.
- [51] Jürgen Pfeffer and Momin M. Malik. Simulating the dynamics of socio-economic systems. In Betina Hollstein, Wenzel Matiaske, and Kai-Uwe Schnapp, editors, *Networked governance: New research perspectives*, pages 143–161. Springer International Publishing, 2017.



- [52] Juan Carrasquilla and Roger G. Melko. Machine learning phases of matter. *Nature Physics*, 13(5):431–434, May 2017.
- [53] Nature. Reporting standards and availability of data, materials, code and protocols.
- [54] Science. Science journals: Editorial policies.
- [55] Malcolm Macleod, Andrew M. Collings, Chris Graf, Veronique Kiermer, David Mellor, Sowmya Swaminathan, Deborah Sweet, and Valda Vinson. The MDAR (Materials Design Analysis Reporting) framework for transparent reporting in the life sciences. *Proceedings of the National Academy of Sciences*, 118(17):e2103238118, April 2021.
- [56] Ian Lundberg, Rebecca Johnson, and Brandon M. Stewart. What is your estimand? Defining the target quantity connects statistical evidence to theory. *American Sociological Review*, 86(3):532–565, June 2021.
- [57] Jake M. Hofman, Amit Sharma, and Duncan J. Watts. Prediction and explanation in social systems. *Science*, 355(6324):486–488, February 2017.
- [58] Leland Wilkinson. Statistical methods in psychology journals: Guidelines and explanations. *American Psychologist*, 54:594–604, 1999.
- [59] Alex Casteel and Nancy Bridier. Describing populations and samples in doctoral student research. *International Journal of Doctoral Studies*, 16:339–362, 2021.
- [60] Leigh Tooth, Robert Ware, Chris Bain, David M. Purdie, and Annette Dobson. Quality of reporting of observational longitudinal research. *American Journal of Epidemiology*, 161(3):280–288, February 2005.
- [61] Daniel J. Simons, Yuichi Shoda, and D. Stephen Lindsay. Constraints on generality (COG): A proposed addition to all empirical papers. *Perspectives on Psychological Science*, 12(6):1123–1128, August 2017.
- [62] Justin Grimmer, Margaret E. Roberts, and Brandon Stewart. *Text as data: A new framework for Machine Learning and the Social Sciences*. Princeton University Press, 2022.
- [63] Leo Breiman. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical Science*, 16(3):199 – 231, 2001.
- [64] Ian Lundberg, Jennie E. Brand, and Nanum Jeon. Researcher reasoning meets computational capacity: Machine learning for social science. *Social Science Research*, 108:102807, 2022.
- [65] Anna Neufeld and Daniela Witten. Discussion of breiman's "two cultures": From two cultures to one. *Observational Studies*, 7(1):171–174, 2021.
- [66] Galit Shmueli. Comment on breiman's "two cultures" (2002): From two cultures to multicultural. *Observational Studies*, 7(1):197–201, 2021.
- [67] Michael Baiocchi and Jordan Rodu. Reasoning using data: Two old ways and one new. *Observational Studies*, 7(1):3–12, 2021.
- [68] Elizabeth L. Ogburn and Ilya Shpitser. Causal modelling: The two cultures. *Observational Studies*, 7(1):179–183, 2021.
- [69] Mario Molina and Filiz Garip. Machine learning for sociology. *Annual Review of Sociology*, 45(1):27–45, 2019.
- [70] Justin Grimmer, Margaret Roberts, and Brandon Stewart. Machine learning for social science: An agnostic approach. *Annual Review of Political Science*, 24, 05 2021.
- [71] Hooman H. Rashidi, Nam Tran, Samer Albahra, and Luke T. Dang. Machine learning in health care and laboratory medicine: General overview of supervised learning and auto-ML. *International Journal of Laboratory Hematology*, 43(S1):15–22, July 2021.
- [72] Andrew L. Beam and Isaac S. Kohane. Big data and machine learning in health care. *JAMA*, 319(13):1317, April 2018.
- [73] Liam G. McCoy, Connor T.A. Brenna, Stacy S. Chen, Karina Vold, and Sunit Das. Believing in black boxes: machine learning for health-care does not need explainability to be evidence-based. *Journal of Clinical Epidemiology*, 142:252–257, 2022.
- [74] Adi L Tarca, Vincent J Carey, Xue wen Chen, Roberto Romero, and Sorin Drăghici. Machine learning and its applications to biology. *PLoS Computational Biology*, 3(6):e116, June 2007.
- [75] Victoria Stodden. Reproducing statistical results. *Annual Review of Statistics and Its Application*, 2(1):1–19, April 2015.
- [76] T. Herndon, M. Ash, and R. Pollin. Does high public debt consistently stifle economic growth? A critique of Reinhart and Rogoff. *Cambridge Journal of Economics*, 38(2):257–279, March 2014.
- [77] Rubén Herzog, Pedro A. M. Mediano, Fernando E. Rosas, Robin Carhart-Harris, Yonatan Sanz Perl, Enzo Tagliazucchi, and Rodrigo Cofre. Retraction note: A mechanistic model of the neural entropy increase elicited by psychedelic drugs. *Scientific Reports*, 12(1):15500, September 2022.
- [78] May R. Berenbaum. Retraction for Shu et al., Signing at the beginning makes ethics salient and decreases dishonest self-reports in comparison to signing at the end. *Proceedings of the National Academy of Sciences*, 118(38):e2115397118, September 2021.
- [79] Marcel Neunhoffer and Sebastian Sternberg. How cross-validation can go wrong and what to do about it. *Political Analysis*, 27(1):101–106, January 2019.
- [80] Jake M. Hofman, Daniel G. Goldstein, Siddhartha Sen, Forough Poursabzi-Sangdeh, Jennifer Allen, Ling Liang Dong, Brenda Fried, Harpreet Gaur, Adnan Hoq, Emeka Mbazon, Naomi Moreira, Cindy Muso, Etta Rapp, and Roymil Terrero. Expanding the scope of reproducibility research through data analysis replications. *Organizational Behavior and Human Decision Processes*, 164:192–202, May 2021.
- [81] Gilles Vandewiele, Isabelle Dehaene, György Kovács, Lucas Sterckx, Olivier Janssens, Femke Ongenaes, Femke De Backere, Filip De Turck, Kristien Roelens, Johan Decruyenaere, Sofie Van Hoecke, and Thomas Demeester. Overly optimistic prediction results on imbalanced data: A case study of flaws and benefits when applying over-sampling. *Artificial Intelligence in Medicine*, 111:101987, January 2021.
- [82] John P. A. Ioannidis, David B. Allison, Catherine A. Ball, Issa Coulibaly, Xiangqin Cui, Aedín C. Culhane, Mario Falchi, Cesare Furlanello, Laurence Game, Giuseppe Jurman, Jon Mangion, Tapan Mehta, Michael Nitzberg, Grier P. Page, Enrico Petretto, and Vera van Noort. Repeatability of published microarray gene expression analyses. *Nature Genetics*, 41(2):149–155, February 2009.
- [83] Timothy Verstynen and Konrad Paul Kording. Overfitting to 'predict' suicidal ideation. *Nature Human Behaviour*, pages 1–2, April 2023.
- [84] Benjamin Haibe-Kains, George Alexandru Adam, Ahmed Hosny, Farnoosh Khodakarami, Levi Waldron, Bo Wang, Chris McIntosh, Anna Goldenberg, Anshul Kundaje, Casey S. Greene, Tamara Broderick, Michael M. Hoffman, Jeffrey T. Leek, Keegan Korthauer, Wolfgang Huber, Alvis Brazma, Joelle Pineau, Robert Tibshirani, Trevor Hastie, John P. A. Ioannidis, John Quackenbush, and Hugo J. W. L. Aerts. Transparency and reproducibility in artificial intelligence. *Nature*, 586(7829):E14–E16, October 2020.
- [85] Victoria Stodden, Jennifer Seiler, and Zhaokun Ma. An empirical analysis of journal policy effectiveness for computational reproducibility. *Proceedings of the National Academy of Sciences*, 115(11):2584–2589, March 2018.
- [86] Mirko Gabelica, Ružica Bojčić, and Livia Puljak. Many researchers were not compliant with their published data sharing statement: A mixed-methods study. *Journal of Clinical Epidemiology*, 150:33–41, October 2022.
- [87] Nicole A. Vasilevsky, Jessica Minnier, Melissa A. Haendel, and Robin E. Champieux. Reproducible and reusable research: Are journal data sharing policies meeting the mark? *PeerJ*, 5:e3208, April 2017.

- [88] AAAI reproducibility checklist. <https://aaai.org/conference/aaai/aaai-23/reproducibility-checklist/>.
- [89] NeurIPS 2023 paper guidelines. <https://neurips.cc/public/guides/PaperChecklist>.
- [90] ICML 2023 paper guidelines. <https://icml.cc/Conferences/2023/PaperGuidelines>.
- [91] The Journal of Politics: Guidelines for data replication <https://www.journals.uchicago.edu/journals/jop/data-replication>.
- [92] B. A. Nosek, G. Alter, G. C. Banks, D. Borsboom, S. D. Bowman, S. J. Breckler, S. Buck, C. D. Chambers, G. Chin, G. Christensen, M. Contestabile, A. Dafoe, E. Eich, J. Freese, R. Glennerster, D. Goroff, D. P. Green, B. Hesse, M. Humphreys, J. Ishiyama, D. Karlan, A. Kraut, A. Lupia, P. Mabry, T. Madon, N. Malhotra, E. Mayo-Wilson, M. McNutt, E. Miguel, E. Levy Paluck, U. Simonsohn, C. Soderberg, B. A. Spellman, J. Turitto, G. VandenBos, S. Vazire, E. J. Wagenmakers, R. Wilson, and T. Yarkoni. Promoting an open research culture. *Science*, 348(6242):1422–1425, June 2015.
- [93] Koren, Miklós, Connolly, Marie, Lull, Joan, and Vilhuber, Lars. Data and code availability standard. December 2022.
- [94] Reviewing computational methods. *Nature Methods*, 12(12):1099–1099, December 2015.
- [95] Kenneth Peng, Arunesh Mathur, and Arvind Narayanan. Mitigating datasets harms requires stewardship: Lessons from 1000 papers. *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, 1, December 2021.
- [96] Data dictionaries, U.S. Geological Survey. <https://www.usgs.gov/data-management/data-dictionaries>.
- [97] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, December 2021.
- [98] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, October 2019.
- [99] Beata Nowok, Gillian M. Raab, and Chris Dibben. synthpop : Bespoke creation of synthetic data in R. *Journal of Statistical Software*, 74(11), 2016.
- [100] Geir Kjetil Sandve, Anton Nekrutenko, James Taylor, and Eivind Hovig. Ten simple rules for reproducible computational research. *PLoS Computational Biology*, 9(10):e1003285, October 2013.
- [101] Timothy H. Vines, Arianne Y.K. Albert, Rose L. Andrew, Florence Débarre, Dan G. Bock, Michelle T. Franklin, Kimberly J. Gilbert, Jean-Sébastien Moore, Sébastien Renaut, and Diana J. Rennison. The availability of research data declines rapidly with article age. *Current Biology*, 24(1):94–97, 2014.
- [102] Elizabeth Gibney and Richard Van Noorden. Scientists losing data at a rapid rate. *Nature*, December 2013.
- [103] Victoria Stodden and Sheila Míguez. Best practices for computational science: Software infrastructure and environments for reproducible and extensible research. 2(1):e21, July 2014.
- [104] Vilhuber, Lars, Connolly, Marie, Koren, Miklós, Lull, Joan, and Morrow, Peter. A template README for social science replication packages. December 2020.
- [105] Matias Singers. Awesome README.
- [106] Harbert. Bash scripting, October 2018.
- [107] April Clyburne-Sherin, Xu Fei, and Seth Ariel Green. Computational reproducibility via containers in psychology. *Meta-Psychology*, 3, November 2019.
- [108] Constanza L. Andaur Navarro, Johanna A. A. Damen, Toshihiko Takada, Steven W. J. Nijman, Paula Dhiman, Jie Ma, Gary S. Collins, Ram Bajpai, Richard D. Riley, Karel G. M. Moons, and Lotty Hooft. Completeness of reporting of clinical prediction models developed using supervised machine learning: A systematic review. *BMC Medical Research Methodology*, 22(1), January 2022.
- [109] Mohamed Yusuf, Ignacio Atal, Jacques Li, Philip Smith, Philippe Ravaud, Martin Fergie, Michael Callaghan, and James Selfe. Reporting quality of studies using machine learning models for medical diagnosis: A systematic review. *BMJ Open*, 10(3), 2020.
- [110] Yoosang Kim, Jidong Huang, and Sherry Emery. Garbage in, garbage out: Data collection, quality assessment and reporting standards for social media data use in health research, infodemiology and digital disease detection. *Journal of Medical Internet Research*, 18(2):e41, February 2016.
- [111] R. Stuart Geiger, Kevin Yu, Yanlai Yang, Mindy Dai, Jie Qiu, Rebekah Tang, and Jenny Huang. Garbage in, garbage out? In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. ACM, January 2020.
- [112] Franz Porzolt, Felicitas Wiedemann, Susanne I. Becker, and C. J. Rhoads. Inclusion and exclusion criteria and the problem of describing homogeneity of study populations in clinical trials. *BMJ Evidence - Based Medicine*, 24(3):92, 06 2019.
- [113] Matthew Salganik. *Bit by bit: Social research in the Digital age*. Princeton University Press, 2019.
- [114] Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning: Limitations and Opportunities*. fairmlbook.org, 2019. <http://www.fairmlbook.org>.
- [115] Abigail Z. Jacobs and Hanna Wallach. Measurement and fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 375–385, New York, NY, USA, 2021. Association for Computing Machinery.
- [116] Marcus Créde and P. D. Harms. Three cheers for descriptive statistics—and five more reasons why they matter. *Industrial and Organizational Psychology*, 14(4):486–488, 2021.
- [117] Jenifer Larson-Hall and Luke Plonsky. Reporting and interpreting quantitative research findings: What gets reported and recommendations for the field. *Language Learning*, 65(S1):127–159, May 2015.
- [118] Valerie C. Bradley, Shiro Kuriwaki, Michael Isakov, Dino Sejdicinovic, Xiao-Li Meng, and Seth Flaxman. Unrepresentative big surveys significantly overestimated US vaccine uptake. *Nature*, 600(7890):695–700, December 2021.
- [119] Luke Plonsky. Study quality in SLA: An assessment of designs, analyses, and reporting practices in quantitative L2 research. *Studies in Second Language Acquisition*, 35(4):655–687, 2013.
- [120] Patrick McKnight, Katherine McKnight, Souraya Sidani, and Aurelio José Figueredo. *Missing Data : A Gentle Introduction*. Guilford Publications, 2007.
- [121] James L. Peugh and Craig K. Enders. Missing data in educational research: A review of reporting practices and suggestions for improvement. *Review of Educational Research*, 74(4):525–556, 2004.
- [122] Matthew J. Salganik, Ian Lundberg, Alexander T. Kindel, et al. Supplementary information for: Measuring the predictability of life outcomes with a scientific mass collaboration. *Proceedings of the National Academy of Sciences*, 117(15):8398–8403, March 2020.
- [123] Christina Mack, Zhaohui Su, and Daniel Westreich. Managing missing data in patient registries: Addendum to registries for evaluating patient outcomes: A user’s guide (Third edition). 2018. <https://www.ncbi.nlm.nih.gov/books/NBK493611/>.
- [124] SWJ Nijman, AM Leeuwenberg, I Beekers, I Verkouter, JIL Jacobs, ML Bots, FW Asselbergs, KGM Moons, and TPA Debray. Missing data is poorly handled and reported in prediction model studies using machine learning: A literature review. *Journal of Clinical Epidemiology*, 142:218–229, February 2022.
- [125] Todd D. Little, Terrence D. Jorgensen, Kyle M. Lang, and E. Whitney G. Moore. On the joys of missing data. *Journal of Pediatric Psychology*, 39(2):151–162, July 2013.
- [126] Jody S. Nicholson, Pascal R. Deboeck, and Waylon Howard. Attrition in developmental psychology. *International Journal of Behavioral Development*, 41(1):143–153, July 2016.

- [127] William R. Sterner. What is missing in counseling research? reporting missing data. *Journal of Counseling and Development*, 89(1):56–62, January 2011.
- [128] Jamilla A. Hussain, Martin Bland, Dean Langan, Miriam J. Johnson, David C. Currow, and Ian R. White. Quality of missing data reporting and handling in palliative care trials demonstrates that further development of the CONSORT statement is required: a systematic review. *Journal of Clinical Epidemiology*, 88:81–91, August 2017.
- [129] Jake M Hofman, Amit Sharma, and Duncan J Watts. Prediction and explanation in social systems. *Science*, 355(6324):486–488, 2017.
- [130] Xu Chu, Ihab F Ilyas, Sanjay Krishnan, and Jiannan Wang. Data cleaning: Overview and emerging challenges. In *Proceedings of the 2016 international conference on management of data*, pages 2201–2206, 2016.
- [131] Erin M Buchanan and John E Scofield. Methods to detect low quality data and its implication for psychological research. *Behavior research methods*, 50:2586–2596, 2018.
- [132] Gilles Vandewiele, Isabelle Dehaene, György Kovács, Lucas Sterckx, Olivier Janssens, Femke Ongenaes, Femke De Backere, Filip De Turck, Kristien Roelens, Johan Decruyenaere, et al. Overly optimistic prediction results on imbalanced data: A case study of flaws and benefits when applying over-sampling. *Artificial Intelligence in Medicine*, 111:101987, 2021.
- [133] Tolou Shadbahr, Michael Roberts, Jan Stanczuk, Julian Gilbey, Philip Teare, Sören Dittmer, Matthew Thorpe, Ramon Vinas Torne, Evis Sala, Pietro Lio, et al. Classification of datasets with imputed missing values: Does imputation quality matter? *arXiv preprint arXiv:2206.08478*, 2022.
- [134] Emilia Gryska, Isabella Björkman-Burtscher, Asgeir Store Jakola, Tora Dunås, Justin Schneiderman, and Rolf A Heckemann. Deep learning for automatic brain tumour segmentation on mri: Evaluation of recommended reporting criteria via a reproduction and replication study. *BMJ open*, 12(7):e059000, 2022.
- [135] Edward Raff. A step toward quantifying independently reproducible machine learning research. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [136] Jon Kleinberg, Annie Liang, and Sendhil Mullainathan. The theory is predictive, but is it complete? An application to human perception of randomness. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 125–126, 2017.
- [137] Ribana Roscher, Bastian Bohn, Marco F Duarte, and Jochen Garcke. Explainable machine learning for scientific insights and discoveries. *Ieee Access*, 8:42200–42216, 2020.
- [138] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature machine intelligence*, 1(5):206–215, 2019.
- [139] Sebastian Raschka. Model evaluation, model selection, and algorithm selection in machine learning, November 2020. arXiv:1811.12808 [cs, stat].
- [140] Gavin C. Cawley and Nicola L. C. Talbot. On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11(70):2079–2107, 2010.
- [141] Emily Black, Manish Raghavan, and Solon Barocas. Model multiplicity: Opportunities, concerns, and solutions. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, June 2022.
- [142] Jesse Dodge, Suchin Gururangan, Dallas Card, Roy Schwartz, and Noah A Smith. Show your work: Improved reporting of experimental results. *arXiv preprint arXiv:1909.03004*, 2019.
- [143] Riashat Islam, Peter Henderson, Maziar Gomrokchi, and Doina Precup. Reproducibility of benchmarked deep reinforcement learning tasks for continuous control. *arXiv preprint arXiv:1708.04133*, 2017.
- [144] Philipp Probst, Anne-Laure Boulesteix, and Bernd Bischl. Tunability: Importance of hyperparameters of machine learning algorithms. *The Journal of Machine Learning Research*, 20(1):1934–1965, 2019.
- [145] Jimmy Lin. The neural hype and comparisons against weak baselines. *ACM SIGIR Forum*, 52(2):40–51, January 2019.
- [146] Michael A. Lones. How to avoid machine learning pitfalls: A guide for academic researchers, February 2023. arXiv:2108.02497 [cs].
- [147] Shachar Kaufman, Saharon Rosset, Claudia Perlich, and Ori Stitelman. Leakage in data mining: Formulation, detection, and avoidance. *ACM Transactions on Knowledge Discovery from Data*, 6(4):15:1–15:21, December 2012.
- [148] Casey Ross. Epic’s sepsis algorithm is going off the rails in the real world. The use of these variables may explain why, September 2021.
- [149] Max Kuhn and Kjell Johnson. *Applied Predictive Modeling*. Springer-Verlag, New York, 2013.
- [150] Russell A. Poldrack, Grace Huckins, and Gael Varoquaux. Establishment of best practices for evidence for prediction: A review. *JAMA psychiatry*, 77(5):534–540, May 2020.
- [151] Mustafa Umit Oner, Yi-Chih Cheng, Hwee Kuan Lee, and Wing-Kin Sung. Training machine learning models on patient level data segregation is crucial in practical clinical applications, April 2020.
- [152] Momin M. Malik. A hierarchy of limitations in machine learning, February 2020. arXiv:2002.05193 [cs, econ, math, stat].
- [153] Christoph Bergmeir and José M. Benítez. On the use of cross-validation for time series predictor evaluation. *Information Sciences*, 191:192–213, May 2012.
- [154] Nils Y. Hammerla and Thomas Plötz. Let’s (not) stick together: Pairwise similarity biases cross-validation in activity recognition. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 1041–1051, Osaka Japan, September 2015. ACM.
- [155] David R. Roberts, Volker Bahn, Simone Ciuti, Mark S. Boyce, Jane Elith, Gurutzeta Guillera-Aroita, Severin Hauenstein, José J. Lahoz-Monfort, Boris Schröder, Wilfried Thuiller, David I. Warton, Brendan A. Wintle, Florian Hartig, and Carsten F. Dormann. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, 40(8):913–929, August 2017.
- [156] Alexandre Chiavegatto Filho, André Filipe De Moraes Batista, and Hellen Geremias Dos Santos. Data leakage in health outcomes prediction with machine learning. Comment on “Prediction of incident hypertension within the next year: Prospective study using statewide electronic health records and machine learning”. *Journal of Medical Internet Research*, 23(2):e10969, February 2021.
- [157] Chengyin Ye, Tianyun Fu, Shiyong Hao, Yan Zhang, Oliver Wang, Bo Jin, Minjie Xia, Modi Liu, Xin Zhou, Qian Wu, Yanting Guo, Chunqing Zhu, Yu-Ming Li, Devore S Culver, Shaun T Alfreds, Frank Stearns, Karl G Sylvester, Eric Widen, Doff McElhinney, and Xuefeng Ling. Prediction of incident hypertension within the next year: Prospective study using statewide electronic health records and machine learning. *Journal of Medical Internet Research*, 20(1):e22, January 2018.
- [158] Maria Carolina Monard and GEAPA Batista. Learning with skewed class distributions. *Advances in Logic, Artificial Intelligence and Robotics*, 85:173–180, 2002.
- [159] Jorge M Lobo, Alberto Jiménez-Valverde, and Raimundo Real. Auc: A misleading measure of the performance of predictive distribution models. *Global ecology and Biogeography*, 17(2):145–151, 2008.
- [160] Alexy Bhowmick and Shyamanta M Hazarika. E-mail spam filtering: A review of techniques and trends. *Advances in Electronics, Communication and Computing: ETAEERE-2016*, pages 583–590, 2018.
- [161] Emily G Simmonds, Kwaku Peparah Adjei, Christoffer Wold Andersen, Janne Cathrin Hetle Aspheim, Claudia Battistin, Nicola Bulso, Hannah Christensen, Benjamin Cretois, Ryan Cubero, Ivan A Davidovich, et al. How is model-related uncertainty quantified and reported in different disciplines? *arXiv preprint arXiv:2206.12179*, 2022.
- [162] Valentin Amrhein, Sander Greenland, and Blake McShane. Scientists rise up against statistical significance. *Nature*, 567(7748):305–307,

- 2019.
- [163] William R. Shadish, Thomas D. Cook, and Donald T. Campbell. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Houghton Mifflin, 2 edition, 2001.
  - [164] Naoki Egami and Erin Hartman. Elements of external validity: Framework, design, and analysis. *American Political Science Review*, 117(3):1070–1088, October 2022.
  - [165] Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. The fallacy of AI functionality. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 959–972, Seoul Republic of Korea, June 2022. ACM.
  - [166] Thomas Liao, Rohan Taori, Deborah Raji, and Ludwig Schmidt. Are we learning yet? A meta review of evaluation failures across machine learning. *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, 1, December 2021.
  - [167] Arunesh Mathur, Angelina Wang, Carsten Schwemmer, Maia Hamin, Brandon M Stewart, and Arvind Narayanan. Manipulative tactics are the norm in political emails: Evidence from 300k emails from the 2020 US election cycle. *Big Data & Society*, 10(1):205395172211453, January 2023.
  - [168] Jason Brownlee. Difference between algorithm and model in machine learning, April 2020. <https://machinelearningmastery.com/difference-between-algorithm-and-model-in-machine-learning/>.
  - [169] Andrea Freyer Dugas, Mehdi Jalalpour, Yulia Gel, Scott Levin, Fred Torcaso, Takeru Igusa, and Richard E. Rothman. Influenza forecasting with Google Flu trends. *PLoS ONE*, 8(2):e56176, February 2013.
  - [170] Olivia Wiles, Sven Gowal, Florian Stimberg, Sylvestre Alvisé-Rebuffi, Ira Ktena, Krishnamurthy Dvijotham, and Taylan Cemgil. A fine-grained analysis on distribution shift, 2021. <https://arxiv.org/abs/2110.11328>.
  - [171] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton A. Earnshaw, Imran S. Haque, Sara Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. WILDS: A benchmark of in-the-wild distribution shifts, 2020. <https://arxiv.org/abs/2012.07421>.
  - [172] Kirk Bansak, Jeremy Ferwerda, Jens Hainmueller, Andrea Dillon, Dominik Hangartner, Duncan Lawrence, and Jeremy Weinstein. Improving refugee integration through data-driven algorithmic assignment. *Science*, 359(6373):325–329, January 2018.
  - [173] Selen Bozkurt, Eli M Cahan, Martin G Seneviratne, Ran Sun, Juan A Lossio-Ventura, John P A Ioannidis, and Tina Hernandez-Boussard. Reporting of demographic data and representativeness in machine learning models using electronic health records. *Journal of the American Medical Informatics Association*, 27(12):1878–1884, September 2020.
  - [174] Constanza L. Andaur Navarro, Johanna A.A. Damen, Toshihiko Takada, Steven W.J. Nijman, Paula Dhiman, Jie Ma, Gary S. Collins, Ram Bajpai, Richard D. Riley, Karel G.M. Moons, and Lotty Hoofst. Systematic review finds “spin” practices and poor reporting standards in studies on machine learning-based prediction models. *Journal of Clinical Epidemiology*, 158:99–110, June 2023.
  - [175] Samuel G. Finlayson, Adarsh Subbaswamy, Karandeep Singh, John Bowers, Annabel Kupke, Jonathan Zittrain, Isaac S. Kohane, and Suchi Saria. The clinician and dataset shift in artificial intelligence. *New England Journal of Medicine*, 385(3):283–286, July 2021.
  - [176] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, November 2020.