

Technologies for Trustworthy Machine Learning

Toreini, Ehsan; Aitken, Mhairi; Coopamootoo, Kovila P. L.; Elliott, Karen; Zelaya, Vladimiro Gonzalez; Missier, Paolo; Ng, Magdalene; Moorsel, Aad van

DOI:

[10.48550/arXiv.2007.08911](https://doi.org/10.48550/arXiv.2007.08911)

Citation for published version (Harvard):

Toreini, E, Aitken, M, Coopamootoo, KPL, Elliott, K, Zelaya, VG, Missier, P, Ng, M & Moorsel, AV 2020
'Technologies for Trustworthy Machine Learning: A Survey in a Socio-Technical Context' arXiv.
<https://doi.org/10.48550/arXiv.2007.08911>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Technologies for Trustworthy Machine Learning: A Survey in a Socio-Technical Context

EHSAN TOREINI, Department of Computer Science, Durham University

MHAIRI AITKEN, Alan Turing Institute

KOVILA P. L. COOPAMOOTOO, School of Computing, Newcastle University

KAREN ELLIOTT, Newcastle University Business School

VLADIMIRO GONZALEZ ZELAYA, School of Computing, Newcastle University

PAOLO MISSIER, School of Computing, Newcastle University

MAGDALENE NG, School of Computing, Newcastle University

AAD VAN MOORSEL, School of Computing, Newcastle University

Concerns about the societal impact of AI-based services and systems has encouraged governments and other organisations around the world to propose AI policy frameworks to address fairness, accountability, transparency and related topics. To achieve the objectives of these frameworks, the data and software engineers who build machine-learning systems require knowledge about a variety of relevant supporting tools and techniques. In this paper we provide an overview of technologies that support building trustworthy machine learning systems, i.e., systems whose properties justify that people place trust in them. We argue that four categories of system properties are instrumental in achieving the policy objectives, namely fairness, explainability, auditability and safety & security (FEAS). We discuss how these properties need to be considered across all stages of the machine learning life cycle, from data collection through run-time model inference. As a consequence, we survey in this paper the main technologies with respect to all four of the FEAS properties, for data-centric as well as model-centric stages of the machine learning system life cycle. We conclude with an identification of open research problems, with a particular focus on the connection between trustworthy machine learning technologies and their implications for individuals and society.

Additional Key Words and Phrases: trust, trustworthiness, machine learning, artificial intelligence

1 INTRODUCTION

Recent years have seen a proliferation of research in technologies for trustworthy machine learning as well as in policy frameworks that introduce requirements around fairness, human rights, privacy and other properties of AI-based services [2]. When building AI-based services, systems and products, the challenge data, software and system engineers face is to design systems that deliver the societal quality objectives articulated in the policy frameworks. In this paper, we consider this challenge in building AI-based systems specifically with respect to machine learning technologies. We do this through a survey that considers technologies for trustworthy machine learning in relation to emerging policy frameworks.

To set the stage, we review and discuss a set of prominent policy frameworks for AI. Such policy frameworks are emerging in many nations, jurisdictions and organisations, to respond to societal concerns about the impact of AI on the quality of life of individuals and on our society, for instance in terms of fairness and justice. There is a noticeable and understandable gap between the subject

Authors' addresses: Ehsan Toreini, ehsan.toreini@durham.ac.uk, Department of Computer Science, Durham University, Durham, United Kingdom; Mhairi Aitken, maitken@turing.ac.uk, Alan Turing Institute, London, United Kingdom; Kovila P. L. Coopamootoo, kovila.coopamootoo@newcastle.ac.uk, School of Computing, Newcastle University, Newcastle upon Tyne, United Kingdom; Karen Elliott, karen.elliott@newcastle.ac.uk, Newcastle University Business School, Newcastle upon Tyne, United Kingdom; Vladimiro Gonzalez Zelaya, c.v.gonzalez-zelaya2@newcastle.ac.uk, School of Computing, Newcastle University, Newcastle upon Tyne, United Kingdom; Paolo Missier, paolo.missier@newcastle.ac.uk, School of Computing, Newcastle University, Newcastle upon Tyne, United Kingdom; Magdalene Ng, magdalene.ng@newcastle.ac.uk, School of Computing, Newcastle University, Newcastle upon Tyne, United Kingdom; Aad van Moorsel, aad.vanmoorsel@newcastle.ac.uk, School of Computing, Newcastle University, Newcastle upon Tyne, United Kingdom.

of the discourse in AI policy frameworks and that of technical approaches in the machine learning literature. Typical concerns in the policy frameworks are expressed in terms such as ethics, justice, fairness, accountability and transparency. These do not map one-to-one to properties of trustworthy systems. (Here, the term trustworthy indicates the existence of system properties that justify parties to place trust in the system [8]).

We argue that the policy framework objectives can be implemented through four key properties of trustworthy machine learning systems: Fairness, Explainability, Auditability and Safety & Security (FEAS). The survey provides an encompassing view of machine learning technologies available with respect to each of the FEAS properties. Of course, each FEAS property deserves its own literature review, the objective in this paper is to highlight the main characteristics and types of approaches for each property to provide a sound foundation for machine learning systems engineering research and development. To further maintain a reasonable scope, within the safety & security category the emphasis is on security and privacy, more so than on safety, reliability and similar concerns.

To build trustworthy machine learning systems, one needs to incorporate concerns about trustworthiness in all stages of the design, from data collection strategies to algorithm selection and visualisation of results. Similarly, these should be considered in all life cycle stages, including design, development and deployment. We use the chain of trust, introduced in [214], to illustrate and discuss the dependencies between stages of the system design and life cycle. In particular, we organise our survey in two classes, Data-Centric Trustworthiness (data collection, pre-processing and feature engineering) and Model-Centric Trustworthiness (model training, testing and inference).

We conclude the paper with a discussion of open problems in building trustworthy machine learning systems. The open problems concentrate on the main research required to establish an implementation pathway for machine-learning based systems to fulfil qualities emerging from the AI policy framework. The future research considers both technology and socio-technical research questions.

The paper is organised as follows. In Section 2, we discuss the relationship between emerging policy framework for AI and the four technology FEAS properties. Section 3 introduces the various stages in the system life cycle at which trustworthiness is to be considered, based on the chain of trust. Section 4 sets the stage for the discussion of FEAS technologies by identifying challenges in trustworthy machine learning and introducing definitions and ways to assess the properties. The core of the survey is in the next two sections, first surveying data-centric technologies in Section 5 and then model-centric technologies in Section 6. Section 7 discusses open research problems to bridge the gap between technologies and the desired societal quality objectives, and Section 8 concludes the paper.

2 RELATING AI POLICY FRAMEWORKS TO TECHNOLOGY PROPERTIES

The widespread integration of AI in services, products and systems has raised various concerns in society about issues such as ethics, employment, legal implications and human rights. To be able to address these concerns a variety of AI policy frameworks has been proposed by governments and organisations world-wide articulated which qualities AI-based systems should exhibit with respect to the identified societal concerns. In this section we discuss the requirements these frameworks pose on AI-based services (Section 2.1) and argue that four technology properties (Fairness, Explainability, Auditability, Safety/Security) should be central to the design of machine-learning systems to fulfil these requirements (Section 2.2).

Table 1. Trustworthy technology classes related to FAT* frameworks. *X*= no mention, *✓*= mentioned, *✓✓*= emphasised

Framework	Year	Document Owner	Entities	Country	Fairness	Explainability	Safety	Auditability
Top 10 principles of ethical AI	2017	UNI Global Union	Ind	Switzerland	✓	✓	✓	✓
Toronto Declaration	2018	Amnesty International	Gov, Ind	Canada	✓✓	✓	X	✓
Future of work and Education for the Digital Age	2018	T20: Think 20	Gov	Argentina	✓✓	✓	✓	✓
Universal Guidelines for AI	2018	The public voice coalition	Ind	Belgium	✓✓	✓	✓	✓
Human Rights in the Age of AI	2018	Access Now	Gov, Ind	United States	✓✓	✓	✓✓	✓
Preparing for the Future of AI	2016	US national Science, and Technology Council	Gov, Ind, Acad	United States	✓	✓	✓✓	✓
Draft AI R&D Guidelines	2017	Japan Government	Gov	Japan	X	✓	✓✓	✓
White Paper on AI Standardization	2018	Standards Administration of China	Gov, Ind	China	✓	X	✓✓	✓✓
Statements on AI, Robotics and 'Autonomous' Systems	2018	European Group on Ethics in Science and New Technologies	Gov, Ind, Acad	Belgium	✓	✓	✓	✓✓
For a Meaningful Artificial Intelligence	2018	Mission assigned by the French Prime Minister	Gov, Ind	France	✓	✓	X	✓✓
AI at the Service of Citizens	2018	Agency for Digital Italy	Gov, Ind	Italy	✓	✓	✓	✓
AI for Europe	2018	European Commission	Gov, Ind	Belgium	✓	✓	✓	✓
AI in the UK	2018	UK House of Lords	Gov, Ind	United Kingdom	✓✓	✓	✓	✓
AI in Mexico	2018	British Embassy in Mexico City	Gov	Mexico	✓	X	✓	✓
Artificial Intelligence Strategy	2018	German Federal Ministries of Education, Economic Affairs, and Labour and Social Affairs	Gov, Ind	Germany	✓	✓	✓	✓✓
Draft Ethics Guidelines for Trustworthy AI	2018	European High Level Expert Group on AI	Gov, Ind, Civ		✓✓	✓	✓✓	✓
AI Principles and Ethics	2019	Smart Dubai	Ind	UAE	✓✓	✓	✓✓	✓
Principles to Promote FEAT AI in the Financial Sector	2019	Monetary Authority of Singapore	Gov, Ind	Singapore	✓	✓	X	✓
Tenets	2016	Partnership on AI	Gov, Ind, Acad	United States	✓	✓	✓	✓
Asilomar AI Principles	2017	Future of Life Institute	Ind	United States	✓	✓	✓	✓
The GNI Principles	2017	Global Network Initiative	Gov, Ind	United States	X	✓	✓	✓
Montreal Declaration	2018	University of Montreal	Gov, Ind, Civ	Canada	✓✓	✓✓	✓✓	✓
Ethically Aligned Design	2019	IEEE	Ind	United States	✓	✓	✓	✓✓
Seeking Ground Rules for AI	2019	New York Times	Ind, GeP	United States	✓	✓	✓	✓
European Ethical Charter on the Use of AI in Judicial Systems	2018	Council of Europe: CEPEJ	Gov	France	✓	✓	✓	✓
AI Policy Principles	2017	ITI	Gov, Ind	United States	✓	✓	✓✓	✓
The Ethics of Code	2017	Sage	Ind	United States	✓	X	X	✓
Microsoft AI Principles	2018	Microsoft	Ind	United States	✓	✓	✓✓	✓
AI at Google: Our Principles	2018	Google	Ind	United States	✓	✓	✓✓	✓
AI Principles of Telefónica	2018	Telefónica	Ind	Spain	✓	✓	✓	X
Guiding Principles on Trusted AI Ethics	2019	Telia Company	Ind	Sweden	✓	✓	✓✓	✓
Declaration of the Ethical Principles for AI	2019	IA Latam	Ind	Chile	✓	✓	✓✓	✓

2.1 AI Policy Frameworks

Table 1 provides a set of policy frameworks that have been proposed up to and including 2019, either exclusively considering AI or for technologies in general, including AI. These frameworks have been introduced by a variety of stakeholders, including technology companies, government, professional and standardisation bodies and academic researchers. The AI elements of these policy frameworks present varying sets of qualities that AI-based systems should exhibit to handle concerns about discrimination, accountability, human rights, etc. By their nature, these documents discuss high-level objectives for the concerned systems as well as the involved science and technology, but do not go into specific technical implementation.

Recent literature provides very interesting reviews of AI policy frameworks. Whittlestone et al. [222] provide a critical analysis of frameworks for ethical AI. They point out the inconsistency across frameworks, for instance, frameworks may confuse the interpretation of qualities, may present conflicting definitions and qualities may be different across different documents. Particularly relevant, also for the underlying technologies, is that the frameworks often fail to realise dependencies between policy objectives (e.g. addressing discrimination may lead to unexpected privacy leakages).

Fjeld et al. [66] also analyse currently available AI policy frameworks, from industry, governments, general public, civil societies and academic bodies. Table 1 includes many of the AI policy frameworks considered by [66]. They identify 47 different objectives, which they categorise into eight qualities. These eight qualities are: (1) privacy, (2) accountability, (3) safety & security, (4) transparency & explainability, (5) fairness & non-discrimination, (6) human control of technology, (7) professional responsibility and (8) promotion of human values. We will map these eight qualities on technology properties in the next section.

2.2 FEAS: Fairness, Explainability, Auditability and Safety

To build AI-based systems that fulfil the above-mentioned qualities distilled from policy framework, we argue that we should consider technologies for four properties: Fairness, Explainability, Auditability and Safety (FEAS).

- **Fairness Technologies:** technologies focused on detection, prevention and mitigation of discrimination and bias
- **Explainability Technologies:** technologies focused on explaining and interpreting the outcomes and actions of AI-based systems to stakeholders (including end-users)
- **Auditability Technologies:** technologies focused on enabling third parties and regulators to supervise, challenge or monitor the operation of the model(s)
- **Safety Technologies:** technologies focused on ensuring the operation of the AI-based system as intended in presence of accidental failures and malicious attackers

Table 2 establishes how the eight qualities derived in [66] for policy frameworks relate to the FEAS properties. In the table, ✓✓ indicates the technology property is necessary in achieving the quality, and, vice versa, the quality sets requirements for the technologies. ✓ indicates the technology property impacts or is impacted by the quality, but not as strong or as comprehensive. If there is no mark, the technology property is not typically necessary and less directly dependent on the requirements of the specific quality. We discuss each of the qualities (as defined in [66]) and explain the relation with the technology properties depicted in Table 2.

Privacy in [66] relates to control over the use of data, the right to restrict processing, etc. Within computing, this usually fits under security, an element of safety. Auditability usually comes with privacy implications, usually negatively impacting on privacy. *Accountability* concerns elements such as assessing impact, legal responsibility, the ability to appeal, creation of a monitoring body, etc.

Table 2. Technology properties to establish the policy framework qualities [136]. ✓ = technology property impacts the quality, ✓✓ = technology property is pivotal to the quality

Policy Framework Quality	Technology Property			
	Fairness	Explainability	Auditability	Safety & Security
Privacy			✓	✓✓
Accountability		✓	✓✓	
Safety & Security				✓✓
Transparency & Explainability		✓✓	✓	
Fairness & Non-Discrimination	✓✓			
Human Control of Technology		✓		
Professional Responsibility				
Promotion of Human Values	✓✓			

If a third-party is involved, this directly relates to auditability (✓✓), and some require the support from technologies for explainability (✓). *Safety & Security* is self-explanatory, as are *Transparency & Explainability* and *Fairness & non-Discrimination*, where we note that *Transparency & Explainability* may require and impact technologies developed for auditability. *Human Control of Technology* includes the ability to opt out and the human review of automated decisions, which primarily relates to technologies for explainability. *Professional Responsibility* refers to ethical design, scientific integrity, which, although relevant for each, does not strongly relate to any of the technology properties specifically. Finally, *Promotion of Human Values* has many dimensions, but many of these are not directly impacted by technology properties. As marked in the table, the exception is the machine learning approaches for fairness that have emerged in recent years.

The mapping in Table 2 between qualities in policy frameworks and concrete technology classes always allows room for debate. To test the mapping is reasonable, we went through each of the policy framework of Table 1 to identify the technology properties that relate strongly with the respective policy frameworks. Table 1 marks frameworks that refer to fairness, explainability, safety and auditability technology properties, using the symbols ✓ and ✓✓ as explained in the table caption. As one can see from Table 1, all four of the FEAS technology properties tend to be mentioned in each of the policy frameworks, with an emphasis on most frequently on Safety & Security and Fairness.

3 TRUSTWORTHY MACHINE LEARNING LIFE CYCLE

The literature presents life cycles for AI-based systems at vastly varying levels of abstractions. In this work, we use the chain of trust for machine learning (Figure 1), which we introduced in [214]. The chain of trust is inspired on the machine learning software pipelines proposed for real-life big data systems [183]. The pipelines combines in a single workflow software packages and tools that perform tasks such as data collection, streaming, analysis, visualization. Examples application for such pipelines include medical imaging [77], IoT [184] or other production systems [183].

The stages identified in pipelines typically encompass the stages identified in Figure 1, worked out in detail for a particular domain. The precise stages that one naturally distinguishes within a life cycle depends on the type of machine learning that is applied. AI applications in robotics are considerably different from the big data oriented systems, systems based on reinforcement learning




Fig. 1. Various stages in the machine learning lifecycle, motivating the notion of a chain of trust.

differ from these based on (un-)supervised learning, and a system that relies on online learning is inherently different from those using a batched learning approach (see, e.g., [219] for a discussion). We believe that the stages in the chain of trust are generic enough to usefully frame the discussion of trustworthy machine learning across the variety of AI-based systems that exists.

The chain of trust in Figure 1 contains two groups of stages, the *data-centric* stages and the *model-centric* stages. In Section 5 and 6 we discuss trustworthiness technologies for these two groups of stages, respectively.

The data-centric stages consist of *Data Collection*, *Data Preparation* and *Feature Engineering*. To establish services that fulfil the qualities proposed in the policy frameworks, it is critical to consider elements such as bias and privacy starting from the initial Data Collection. The data will go through other pre-processing to facilitate the application of machine learning techniques. This is represented by the stages Data Preparation, to ‘clean up’ the data sets to be fit for the application of machine learning approaches, and Feature Extraction, in which we include feature engineering by experts, dimensionality reduction and feature selection. We would not include in this data-centric stage feature learning approaches such as principal component analysis or autoencoders, since these are machine learning techniques in their own right.

The model-centric stages we distinguish are Training, Testing and Inference. *Training* refers to the building of the model based on a subset of the data, the training data. The *Testing* stage refers to evaluating the model arrived at in training, based on some metric. For example, in classification, a typical accuracy metric is the proportion of the test data that has been mapped correctly to their original label. The *Inference* stage refers to the actual use of the model to make decisions or predictions. Delineating between these three stages is natural in batched learning, in which quality data is available to train and test offline. In online or incremental learning approaches, e.g., used if the training data is too large to allow offline training or if the characteristics of the data change relatively rapidly over time, the stages cannot be delineated so clearly since training, testing and inference is intertwined [158].

The different stages in the chain of trust are not independent. For instance, bias in the results of models are impacted by how the data was pre-processed. In Figure 1 we represent this dependency between the stages through the chain. Moreover, systems will often iterate through (some of) the

stages during their lifecycle, particularly in online learning settings where a model is continuously updated. The chain of trust is therefore represented a cycle in Figure 1. We refer to [214] for a detailed discussion of various additional characteristics of the chain of trust.

4 FEAS: DEFINITIONS, CHALLENGES AND ASSESSMENT

In this section we introduce the FEAS properties, providing definitions and discuss how to assess each property.

4.1 Fairness

A decision rule is said to be *unfair* if it treats individuals with similar decision-related features in a different way due to personal characteristics that should not affect the final decision. Recent years have seen several examples of fairness issues in machine learning models [151], for instance Amazon Inc. adopted a ML algorithm for its recruiting system that proved to be biased against women [47]. As a consequence, fairness has gained dramatically increased attention in both academia and industry, aiming to establish variants of fairness-aware machine learning, which Friedler et al. [72] defines as “algorithms to provide methods under which the predicted outcome of a classifier operating on data about people is fair or non-discriminatory”. The literature is extensive, crosses multiple disciplines, and arguable, as a new and fast moving field, still dispersed [192].

Fairness is particularly important if it concerns discrimination in relation to data features about *Protected Attributes* (PAs). These usually include (but are not limited to) ethnicity, gender, age, nationality, religion and socio-economic group. PAs are also referred to as *sensitive* attributes in the literature. California Law [9] distinguishes disparate treatment and disparate impact. Disparate treatment occurs if the PAs affect the process output, while disparate impact happens when the output of the process disproportionately impacts depending on PAs.

Machine learning algorithms are typically not inherently fair or unfair; however, the outcome of a machine learning model *can* be unfair due to several reasons [151]. Collected data is, by its nature, a reflection of the real world. Thus, unfairness or discrimination in the real world would be reflected in related datasets, which would feed into the algorithms and therefore produce unfair outcomes. A dataset doesn’t become fair by simply ignoring the protected attributes [143], since correlations between features could still lead to discrimination. For example, an individual’s post code, degree, hobbies, etc., are likely to be correlated to their ethnicity. In addition, other phenomena may contribute to bias in the data, for instance the lower frequency in data of underrepresented groups (known as *size disparity*), which may cause an inferior predictive performance for these minority groups.

Designing a fair machine learning system requires being able to measure and assess fairness. Researchers have worked on formalising fairness for a long time. Narayanan [166] lists at least 21 different fairness definitions in the literature and this number is growing. A full list of fairness definition can be found in [218]. For the following definitions, let the set A denote the protected attributes of an individual U , and X the rest of attributes that can be observed by the algorithm. Let Y be the actual label of an individual and \hat{Y} the predicted label stemming from the machine learning algorithm. The different fairness definitions may be categorised as suggested by Kusner et al. [122], as follows.

- **Fairness Through Unawareness (FTU).** An algorithm satisfies FTU when its predictions are independent of the PAs [122]:

$$\hat{Y} : X \rightarrow Y.$$

The above definition simply discards the protected attributes (A) in the decision making process. FTU is also referred to as **anti-classification** [40] in the literature.

- **Individual Fairness (IF).** An algorithm satisfies IF when similar samples yield similar outcomes. Generally, given a metric d and some value δ , if $d(U_i, U_j) < \delta$, the predictions for U_i and U_j should be roughly similar:

$$\hat{Y}(X_i, A_i) \approx \hat{Y}(X_j, A_j).$$

The problem with IF is that metric d needs to be adequately defined (or learnt, see [122]). Thus, the implementation of such fairness is prone to increasing unfairness when not chosen correctly.

Joseph et al. [108] propose a metric known as *Rawlsian Fairness* based on Rawl's notion of fair equality for opportunity. This notion emphasises that the individuals with same talents and motivations should have the same prospects of success.

- **Group Fairness.** This refers to a family of definitions, all of which consider the performance of a model on the population groups level. Most of the fairness definitions in this group are focused on keeping positive ratios proportional across groups. The definitions in this category are relevant to the disparate treatment and impact notions. The most popular group fairness definitions—as presented in [64, 95]—follow, where $A = 0$ and $A = 1$ determine membership of a protected group or not.
 - **Demographic Parity (DP).** A classifier satisfies DP when outcomes are equal across groups:

$$P(\hat{Y} = 1 \mid A = 0) = P(\hat{Y} = 1 \mid A = 1).$$

- **Equalised Odds (EO).** A classifier satisfies EO if equality of outcomes happens across both groups and true labels:

$$P(\hat{Y} = 1 \mid A = 0, Y = \gamma) = P(\hat{Y} = 1 \mid A = 1, Y = \gamma),$$

where $\gamma \in \{0, 1\}$.

- **Equality of Opportunity.** is similar to EO, but only requires equal outcomes across subgroups for true positives:

$$P(\hat{Y} = 1 \mid A = 0, Y = 1) = P(\hat{Y} = 1 \mid A = 1, Y = 1).$$

We note that [95] relates these metrics to disparate treatment and impact discussed above. Other metrics of fairness related to disparate treatment may be found in [39, 41].

- **Counterfactual Fairness (CF).** The use of causal models can be a very powerful fairness analysis tool [122, 123]. Given a causal model, CF compares the probabilities of the possible outcomes for different *interventions*—substituting the structural equations of a PA with a fixed value. Formally, a predictor \hat{Y} satisfies CF if under any context $X = x$ and PA value $A = a$,

$$P(\hat{Y}_{A \leftarrow a} = y \mid X = x, A = a) = P(\hat{Y}_{A \leftarrow a'} = y \mid X = x, A = a)$$

for all y and for any value a' attainable by A .

The above definitions of fairness are statistical definitions that can be directly applied to machine learning. Grgic-Hlaca et al. [86] argue that is not always possible or desirable. They propose a procedure to ensure a fair use of attributes through users' feedback and call this **Process Fairness (PF)**. An algorithm satisfies PF if the use of attributes is considered fair in the model assessed through some other means.

It is important to note that establishing fairness with respect to one definition does not necessarily imply into fair output in another notion. In fact, it is theoretically not possible to fulfil all of the definitions at the same time for a model, as shown by Friedler et al. [70]. Similarly, fairness definition

are not necessarily easy to apply correctly, as demonstrated by Kusner et al. [122] who show that some specific definitions of fairness can increase discrimination if applied incorrectly, while [70] show that slight changes in the training set has noticeable impact on the discrimination performance in the output, thus leading to result instability. Clearly, designing discrimination-free models is not straightforward, something we will elaborate on when discussing open problems.

4.2 Explainability

In the machine learning literature, the terms *explainability*, *transparency*, *intelligibility*, *comprehensibility* and *interpretability* are often used interchangeably [90]. Lipton [137] provides a detailed discussion on terminology. The challenge in explainability is to bridge the gap between the operations of the algorithms and human understanding of how the outcomes has been obtained. According to Guidotti et al. [90], explanation is an interface between the decision maker and humans. An explanation should strive to be *accurate*, as it portrays the inner functionality of the machine learning model, as well as *comprehensible*, as it intends to convey the decision process to humans. Herman [98] sees this as a trade-off, between explaining a model's decision faithfully and presenting it a human understandable way. Guidotti et al. [90] suggest an explainable model must have at least two qualities in addition to accuracy, namely interpretability, or the extent to which a consumer of the explanation finds it understandable, and fidelity. The latter refers to the extent that an explainable model operates similar to a black-box model. Ideally, the explainable model must give the same results as the black-box model.

Doshi-Velez et al. [53] introduce two types of explainability, namely explainability with reference to a specific application, for instance a physician being able to understand the reasoning behind a classifier's medical diagnostic, and domain-independent explainability, where it is possible to understand how the classifier has produced one particular prediction. They describe three dimensions for such system: *level* of the explanation, i.e., global vs local *time* limitations, and *user expertise*. Time limitations refer to the amount of time that the user needs to understand the explanation. In emergencies, such as surgical procedures and disasters, the users require a simple, straightforward explanation so they would spend minimum amount of time to understand the decision process. However, in less time-constrained situations, the users might prefer detailed explanations over simplistic ones. Also, the method of explanations is dependent on the user's level of expertise. For instance, a data scientist is more interested to know more about the inner patterns and parameters of the model and the data whereas an executive manager requires a bird-eye explanation of how a specific decision is determined [38].

Technical approaches to providing explanations from predictive machine learning models can be *model-specific* or *model-agnostic*. Examples of model-specific explanations are discussed e.g. in [69]. These include not only linear regression and logistic models, which are arguably the simplest "native explainable" models, but also decision trees, classifiers such as SVM (Support Vector Machines), and Bayesian probabilistic classifiers. Model-specific explanations do not capture the complexity of highly non-linear models, including Neural Networks and Deep Networks in particular.

Model-agnostic explanation approaches address this category of models in a *post hoc* fashion, namely by accepting that the model is a *black box*, and attempting to approximate its behaviour by repeatedly requesting predictions in specific areas of the feature space [220]. A notable example of this approach is Lime [187], which seeks to provide local explanations in the form of linear approximations of the model that are accurate in small regions of the space. This is useful, for example, to explain why a certain individual has been denied a mortgage application. Other post-hoc model-agnostic approaches provide explanations in the form of a ranking of features, even when the underlying model is not-linear. These include, amongst others, InterpretML [129, 173] and Shapley Values [144]. The main characteristic of both these approaches is to provide not only a ranking of

features in order of importance, but at a more granular level, an indication of the importance of the feature as a *function of its value*. For instance, an explanation may indicate that age becomes a more a dominant feature to predict a negative outcome for some disease as it increases, but it may be irrelevant for, say, young patients. Finally, Van der Schaar et al. [3] propose a breakthrough approach to model-agnostic explainability based on the notion that arbitrarily non-linear models can be approximated using a set of equations, from which explanations can be more easily derived.

There is consensus over the lack of evaluation framework to *assess* explainability of a model [53, 137, 155, 200]. Kononenko et al. [116] suggest a qualitative approach to measure explainability. Letham et al. [133] propose an evaluation framework that were limited in terms of scalability or needed prior expertise before using them. Other proposals [144, 197] propose frameworks that involve an explainability dimension along with other parameters. They mainly focus on the effects explainability has on computational efficiency on the model. Finally, Schmidt et al. [200] propose an inclusive quantitative framework to measure explainability in various machine learning algorithms. Their proposed framework measures the effectiveness of the explanation for algorithmic decisions by focusing on the information transfer rate at which humans make the same decisions. They also defined a trust metric which determines how humans are willing to accept the algorithmic explanations.

4.3 Auditability

Auditability refers to the methods that enable *third parties* to verify the operations and outcome of a model. The motivation for auditing very often stems from the other three properties, fairness, explainability and safety & security. For instance, since model fairness is affected by bias in the training set as discussed in Sec. 6.1, one may want to establish an auditable trail to decide whether the bias was native to the raw data, and whether it was affected by any of the operations performed during data pre-processing.

Auditability of the model generation process and the model operation is important for the implementation of regulations such as the European GDPR legislation. At its core, GDPR is about data protection, posing requirements on industry, government and other data owners about safeguarding the date, transparency of the data collection process and explicit consent from an individual for the use of its data. This implies that under GDPR people have the right to challenge the data collection, have control over their collected data and maintains the “right to be forgotten” from the data source. [182] notes a trade-off between keeping data private and the ability of machine-learning systems to perform certain tasks. This trade-off is a key concern behind policy shaping for AI-based systems.

Particularly interesting is that GDPR also requires “suitable explanations” to be provided for any decision made by an algorithm [85]. Providing mechanisms to audit a model, as well as the process used to obtain a model, ensure that the model itself is accountable for the decisions it produces. The case for accountability in algorithmic decision making has also emerged outside GDPR, most notably in [50] and in other countries such as India [28] and China [195] legislation is under review. In turn the need for accountability entails validating the fairness of a model.

4.4 Safety & Security

The category of safety and security properties refers to the robust operation of systems in real-life operation. This covers a wide variety of concerns, from safe vehicles to protection of data against privacy attacks. In this paper we emphasise the review of literature with respect to security and privacy. For a discussion of terminology and a broad perspective on the design of dependable systems, see [8].

Security of AI-based systems and services refers to the protection of these systems against malicious attacks. Traditionally, the objective in securing a system is expressed through the CIA

triad: Confidentiality, Integrity and Availability [208]. *Confidentiality* attacks are designed to expose the structure, configuration of a model or privacy leakage to access the data. Privacy equates to maintaining confidentiality. Attacks to the *integrity* attempt to alter the output of model to the adversary's choice or change the behaviour of the model to the attacker's intention. Finally, *availability* attacks aim to prevent the model's usual service to the client or getting meaningful results.

A perfectly secure system does not exist [5, p. 21] and this holds also for AI-based systems [10]. The security of a machine learning system (see Section 6.4) are impacted by the chosen cryptographic primitives, hardware vulnerabilities, protocol inefficiencies in design and implementation as well as the behaviours of humans implementing, deploying and using it. Clearly, the current large-scale adoption of machine learning motivates attackers to actively engage into detecting design weaknesses or implementation bugs and leveraging them to collect information [209], modify data [14, 20, 22, 238], contaminate the operation of the algorithm [10, 115, 203, 209], exploit leakage of sensitive information [55, 61, 102, 119, 142] and much more [32, 120, 139, 161, 168, 172, 177, 203].

Understanding the attacker is the building block of any defensive mechanisms [5, p. 20]. Attacks on machine learning systems exhibit intention and capability [169] or, similarly, goal, knowledge, capability and attacking strategy, see Baggio et al. [19]. In [19], the *goal* of the attack describes the desired impact of the attack in accordance with the CIA model. The goal's impact can be of two types: targeted attacks, which aim to damage a specific victim, or indiscriminate attacks, which aim for extensive damage with no exception. *Knowledge* determines the degree of intelligence that the attacker has at the time of the execution of the attack. A black-box adversary designs the attack without any knowledge of the model's internal configurations, whereas a white-box adversary knows the settings in which the model is trained and operates. In the real-world, the black-box attacks is reportedly more common [178]. *Capability* determines the degree to which the attacker is able to successfully implement an attack scenario and get access to the data or the model. Capability can be a qualitative feature with regards to three aspects [138]: whether the security threat is causative (i.e. direct manipulation of the model or data) or exploratory (i.e. injection of manipulated data to mislead the model), the percentage of the training or test data that is controlled by the attacker and to which extent the attacker knows the model configurations and its features. Capabilities increase significantly if one considers insiders, that is, individuals that know about one or more of the pipeline stages. To avoid weakening the overall system [178] argues that there should be different levels of access to different actors involved in the model pipeline, i.e. *data-owners*, *system providers*, *consumers and outsiders*. The *attacking strategy* deals with the analysis of the behaviour of the attacker at the time of the exploit. There are usually different ways to implement an attack with a certain goal. Strategy refers to selection of the optimal function that characterises the goal.

Security assessment of a system is usually based on a "what-if" analysis [138]. The appropriate defensive mechanism is implemented based on the conditions of the data, model and the environment that the system is going to be deployed. The designers should provide an assessment of the vulnerabilities of the model, too [17].

When designing a machine learning system for security, one needs to combine a reactive design approach with a proactive approach [138]. In reactive design, the designer implements defensive mechanisms after a vulnerability is detected. However, the current pace of progress in machine learning and applications justifies one to consider security and privacy from the initial stages of designing a model until the model operates in the wild. Thus, the trend is to take a proactive attitude and anticipate common threats through a set of *what-if* queries [191]. For instance, in a proactive design approach, the model designer may invite a red team to attack the system before

deployment. In either approach one needs to realistic analysis of the potential attackers for the goal, knowledge, capability and strategy dimensions of attackers identified by [19].

Both approaches take an empirical way to tackle security. Thus, they rely on a predefined set of *known* attack scenarios. As a result, they lack comprehensive vision for a safe system because they do not take into account creative and novel attacks in their threat model. The more vital applications of machine learning require stricter approach in the design by modelling the behaviour of the system mathematically [74, 124, 125, 135, 202, 225].

5 DATA-CENTRIC TRUSTWORTHINESS TECHNOLOGIES

In this section, we review trustworthiness solutions in the data-centric stages in machine learning pipeline (see Figure 1), that is, data collection, pre-processing and feature selection, respectively. The subsequent sections review each of FEAS properties.

5.1 Fairness in the Data-Centric Stages

Data inevitably reflects the nature of the environment that it has been obtained from as well as the decisions made in that environment. Pedreshi et al. [180] aptly state that “learning from historical data may lean to discover the traditional prejudices that are endemic in reality”.

An important class of bias in ML is *sample bias* [45]. This refers to sampling in such a way that not all PA groups are adequately represented, i.e. to the dataset being *imbalanced* [132]. This kind of bias is not to be confused with *statistical bias*, which refers to a systematic error in the expected result of a process, differing significantly from the true value of the quantity being estimated [51]. The notion of sample bias is in direct relation with group fairness definitions as explained in section 4.1 since they both concern with maintaining a *balanced* ratio of data vs. output in the machine learning system. There are two possible sampling bias directions for PA groups, namely to be *over* or *under* represented [45]. Naturally, one PA group being overrepresented means that some of the remaining groups will be underrepresented. This poses the risk of a learning algorithm being giving up on accuracy for the minority groups for improved overall accuracy [9, 111, 132]. An extreme case of underrepresentation would be the existence of a group being unknown to the algorithm.

A first step to reduce bias is to ensure the data has been collected in an adequate manner [4]. This implies that the collection process should yield representative data from all PA groups. This is not always easy to achieve, since full understanding of the environment and full control over the data-collection processes are needed. In order to prevent sampling bias, a possible approach is *stratified sampling* methods [171].

A potential strategy to avoid bias is to learn about the unintended discrimination that may occur in the specific problem one is resolving. Such *awareness* is vital but not sufficient to prevent discrimination [45]. The active learning approach involves *expert* human supervision in the collection to improve the process [201]. Alessandro et al. [45] suggest that data scientists should also test and document all the features that have high correlation with the minority groups. Also, data scientists should perform a statistical demographic analysis to gain knowledge about groups that are prone to discrimination.

A different approach consists of pre-processing the data, with the aim of feeding an unbiased training set into the model [60, 127]. These pre-processing methods can be based on either class labels or on attributes. In the class labels approach, data is relabelled to reduce disparity; this is sometimes referred to as *massaging* the dataset. These correction methods may be generic [109], or adjusted with respect to a particular combination of *world-views* [71] and fairness definitions [233]. Alternatively, Luong et al. [145] assign a decision value to each data sample and then adjusts this value to eliminate discrimination.

In the *attributes approach* to pre-processing, one either maps the attributes into a transformed space [31] or combines the values of the protected attributes with the non-protected [64]. Feldman et al. [64] propose to modify each attribute so that marginal distribution of the PA and non-PA values remain equal all over the dataset.

Pre-processing techniques have no awareness about the internal operation of a machine learning model and as a result they are not tuned for the optimal functioning of the model, resulting in a trade-off between fairness and accuracy [241]. This, however, makes most of them model-agnostic and highly interpretable. Relabelling of the training data, though, potentially affects the explainability of the model [231] and has a steeper fairness-accuracy trade-off [233].

5.2 Explainability in the Data-Centric Stages

Data has different degrees of explainability to humans [90]. The vast majority of machine learning algorithms use data representation in vector and metrics forms [90]. Johansson et al. [103] suggest that *tables* are the most understandable data format for humans and algorithms. Tables are a familiar format of data representation for humans, while the process of converting them to vectors and metrics is straightforward. *Images* and *texts* are also understandable to humans but conversion to matrices and vectors and back is not straightforward, thus making it more challenging to explain the steps of the algorithm [90]. Explainable algorithms that can process images and texts would require a transformation processes such as equivalences, approximations or heuristics for explaining the decision making process.

Explainability in other types of data such as sequence data, spatio-temporal data and complex network data, has not yet been researched. A potentially effective explainability tool is to visualise the data (and the steps of the algorithms). Visualization could assist in identifying correlations and connections between seemingly unrelated variables. In practice, however, the high dimension of the data does not allow comprehensive visualisation. One solution is transforming high-dimensional data into a two or three-dimensional space, e.g., [146].

Zelaya [82] highlights the impact of the appropriate data-processing methodology before training of the model in the overall explainability of the machine learning system. This includes the data-centric stages of data collection, pre-processing and feature selection. Zelaya [82] defines a metric which compares the outputs of classifiers trained over two training sets, an unprocessed and a pre-processed one. This comparison allows one to analyse the impact of pre-processing on the explainability of the model. Blum et al. [23] present a survey on data selection methods for large datasets containing redundant and irrelevant data samples. They focus on two key issues in dataset, first, selection of features and second, selection of most useful samples from the dataset. Crone et al. [44] study the effects of pre-processing methods such as scaling, sampling and encoding variables over prediction performances. Similar analyses have been performed for text data such as documents [81], text classification [216] and sentiment analysis [56]. Other directions such as [29], followed the same approach.

In our opinion, explainability in the data collection, pre-processing and feature engineering stages of the machine learning life cycle has been underexposed compared to work on algorithm. Arguable, even for explainable algorithms there is a clear need to explain the sources of data and how it was processed. Labelling standards in supervised machine learning are only a small step and the area deserves much more attention.

5.3 Auditability in the Data-Centric Stages

Auditing of the operations of a real-life machine learning system inherently requires one to monitor across all the stages of the chain of trust, from data collection all the way to inference.

technology solutions therefore also tend to span all stages. This is particularly true for provenance solutions that target machine-learning systems [91].

We therefore refer to Section 6.3 for an extensive discussion of provenance technique for machine learning systems.

5.4 Safety & Security in the Data-Centric Stages

There exist a plethora of techniques for safe and secure information systems. In this section, we first discuss how these general concerns apply to systems and services that are based on machine-learning, with a particular emphasis on big data. Then we discuss interesting recent developments in privacy through differential privacy and data-related attacks on machine learning associated with data poisoning.

One of the main differentiators of many modern-day machine-learning systems is the large volume of data that requires to be stored or streamed. The sensitivity of the analysis implies that there is an emphasis to store the data securely to preserve privacy and security against various types of attacks is imminent. Perdosci et al. [181] managed to implement a real-world attack in Polygraph [170] (a worm signature generation tool) by systematically injecting noise into the flow detection dataset and consequently reduce the accuracy of model intentionally.

Lafuente [126] suggests several security strategies for storing big data, both structured and unstructured. First, the data should be *anonymised*, meaning personal identifiable information is removed from the dataset. Note that this alone may not be sufficient for privacy protection since the anonymised data can be cross-referenced with other sources of the data to leak information, a process is known as de-anonymisation or linking attacks in the literature [167]. Second, the data must be *encrypted*. This is particularly the case when it comes to cloud-based machine learning services (i.e. Machine Learning as a Service [99, 186]), both for storage and data transmission. Since the data requires processing, and thus decryption, encryption needs to be complemented by strong security guarantees of the (virtualized) servers. To make sure the data is only decrypted by the party that owns the data, one would utilize reliable key distribution or concepts such as attribute-based encryption [12] which enables a fine-grained access control on the encrypted data. Third, access control and monitoring is essential to protect the data. Usually, a combination of cryptography and access control forms an appropriate approach [126]. However, even popular systems such as Hadoop lack a proper authentication solution [105]. Real-time monitoring is typically accompanied by threat intelligence systems and up-to-date vulnerability detection. The fourth strategy around governance, including implementing regulatory frameworks such as GDPR (see also auditability).

An important recent concept with respect to privacy is to store data using cryptographic primitives that provide *Differential Privacy* [57]. In differential privacy, the calculation of the result by a machine learning algorithm is not sensitive to change in an individual data item. In other words, the output of a machine learning system does not reveal information about an individual. Differential privacy models have two main benefits in comparison with the conventional privacy preserving mechanisms. First, it protects against attackers with full knowledge of the data records except for the target record, i.e., the strongest possible attacker in practice. Second, differential privacy properties of functions are mathematically proven. Differential privacy is considered one of the main trends in the privacy-preserving machine learning literature [1, 62, 121, 193, 221].

To assure integrity of the data, pre-processing techniques benefit from tampering-prevention measurements [178]. The integrity attacks in this category work based on polluting or poisoning the datasets. Poisoning attacks on the training data aim to alter the model that is derived during the training stage. Poisoning attacks on the testing set attempt to manipulate the test data distribution in order to separate it from the training set. This would eventually increase the likelihood of false

recognition in presence of polluted data when the machine learning model is operating in the inference stage [63, 153, 227].

Security architects of a machine learning systems evaluate the degree of malicious modifications of data with probability β in their security assumptions. Kearns et al. [113] determine an upper bound for β . In their analysis, a machine learning algorithm with an error rate of ϵ , requires attackers to modify a fraction less than $\frac{\epsilon}{1+\epsilon}$ of the training set. The attacker with access merely to labels of the training data is constraint to a limited attack surface, requiring one to determine the most useful labels to manipulate [178]. In this case, the attacker can operate by perturbing the labels randomly or heuristically. Baggio et al. [20] demonstrate that to impact the efficiency of SVM classifier, random manipulation of 40 % of labels in a training set is required. In contrast, it requires 10 % of the labels to be manipulated if it is done purposely.

A number of further advances have been made with respect to data poisoning attacks. If attackers gain access to all the features in addition to the labels in training set, then they would be capable of implementing more serious attack scenarios [152, 226]. They would benefit them the most if training stage is performed online [237] (i.e. the training dataset is updated regularly as the data is being collected). In online setting, Kloft et al. [115] poison the training set of a clustering system for anomaly detection and the same approach has been applied to malware clustering techniques [22]. In offline learning, Biggio et al. [21] propose an attack that detect the target sample in the training set and degrades classification accuracy accordingly, assuming the classifier is a SVM. Xiao et al. [226] propose the same strategy to degrade the functionality of feature selection algorithms like LASSO.

The countermeasures to malicious poisoning attack rely on monitoring the distribution characteristics of training and testing datasets. The existence of adversarial samples in a data set deforms the distribution; thus, some researchers propose anomaly detection techniques based on data distributions in adversarial environments [17, 19, 68, 150]. Grosse et al. [87] introduce two statistical metrics, Maximum Mean Discrepancy and Energy Distance, to measure the statistical properties of sample distributions especially on test samples. Their proposal successfully detects malicious injected samples in several data sets.

Mitigation techniques for data integrity focus on ensuring the purity of training and testing data sets [138]. Data sanitization is a dominant technique in this area, e.g., [169] propose Reject on Negative Impact to protect spam filters against malicious poisonings. The basic idea of this method is to eliminate samples with large negative impact in the training set. They also propose an error estimation method in online training mode to compare the error rates of the original classifier and the one after the update. Laishram et al. [128] propose a solution protect SVM classifier from poisoning attacks. Finally, some proposals focus on feature reduction algorithms to prevent poisoning attacks. Rubinstein et al. [194] proposed an PCA-based poisoning detector based on techniques from robust statistics. Their method substantially reduces poisoning in various intrusion detection scenarios.

6 MODEL-CENTRIC TRUSTWORTHINESS TECHNOLOGIES

In this section we discuss technologies for the model-centric stages in the machine learning chain of trust, that is Training, Test and Inference stages. The following sections survey of technologies with respect to each if the four FEAS properties.

6.1 Fairness in the Model-Centric Stages

The research in fairness for machine learning models is developing very fast. In this section we aim to present a high-level perspective on types of approaches presented in the literature, with examples. To that end, we follow [92] and distinguish two types of model-centric solutions for

fairness: *in-processing* solutions aim to design an algorithm that is inherently fair [72, 147], while *post-processing* solutions adjust the outcomes so the outcomes are fair [93]. Within in-processing approaches, one can distinguish between those modifying the model to remove unfairness and those in which an independent agent carries out the detection (or prevention) of unfairness.

Using the in-processing approach various solutions introduce a fairness-regulating term in the algorithm. For instance, Kamishima et al. [111] propose a regularization term for a logistic regression classifier: this term penalizes existing correlations between the protected attributes (PAs) and the predicted labels, leading to fairer classification. Kamiran et al. [110] propose a fair decision-tree algorithm, which defines a criterion to either split or label a leaf node in order to remove discrimination. Other approaches add a set of constraints to the model and then optimise them simultaneously while the model is training. For instance, Goh et al. [80] use a non-convex constraint to optimise a classifier under a fairness constraint based on disparate impact. Quadrianto et al. [185] propose a fairness-definition agnostic unified framework. They tune two ML algorithms (privileged learning and distribution matching) to enforce different notions of fairness. Woodworth et al. [224] present a statistical method used in the learning process and propose a non-discrimination definition that increases the fairness of the learning algorithm. Zafar et al. [232] introduce flexible fair classification and suggest a novel fair boundary decision threshold. They apply their approach to logistic regression and support vector machines. And, as a last example of introducing fairness regulation in the algorithms, Calders et al. [30] train Naive Bayes models for every possible PA value and then adjust the model to balance the outcomes.

A second important in-processing approach is the use of *adversarial learning* approaches. In adversarial learning, two models are simultaneously trained: the first one modelling the data distribution and the second one estimating the probability of a sample coming from the training data rather than from the model, the objective being to maximise the probability of the first model causing the second one to make a mistake [83]. Examples of adversarial learning for fairness are Beutel et al. [13]—who aim at enforcing equality of opportunity—and Zhang et al. [235]; this second method is applicable to multiple fairness definitions and gradient-based learning models.

Post-processing approaches analyze the outcomes of a model and then adjust them to remove unfairness. An important limitation of these solutions is that they require access to PAs at the moment of decision making. Post-processing techniques can therefore not be used when there is no access to the PAs. There are various examples of post-processing approaches. For instance, Hardt et al. [95] design a post-processing approach which adjusts a learned predictor in order to remove discrimination with respect to the equalised odds rate in the results. Corbett-Davis et al. [41] work with three notions of fairness: demographic parity, conditional demographic parity—in which proportions across PA groups should remain constant given a limited set of legitimate risk factors, i. e. attributes naturally related to the classification task of interest—and predictive equality, which is satisfied when there are equal false positive rates across PA groups. They then introduce an algorithm to assign a discrimination score to a classifier, which in turn allows for detecting unfairness.

In addition to the above specifically model-centric approaches, one can combine data-centric and model-centric solutions. Zemel et al. [234] propose an algorithm optimising individual and group fairness at the same time through an adequate representation of the data, obfuscating membership to PA groups. They argue that there is a close connection between individual fairness and the differential privacy technique of hiding attributes of data through the addition of noise without it losing its utility [59]. Finally, Dwork et al. [58] propose a *combined* in-processing and post-processing method based on training different classifiers for different PA groups.

6.2 Explainability in the Model-Centric Stages

Explainability of machine learning models has also experienced a surge in research, and as for fairness, our aim in this paper is to provide a high-level overview of the main types of explainability techniques, including examples. Therefore, we adopt the classification of Guidotti et al. [90], who distinguishes the following types model explainability approaches: explaining the outcome of the model, known as *outcome explanation* and explaining the inner operation of the model, known as *model explanation*. In addition, within model explanation, we distinguish explanation through *transparent models* in which the model is designed for explainability, effectively interpreting its own inner complexities, as well as techniques for *representation of the model* based on documentation, visualization or provenance information.

Outcome explanation does not consider the inner functionality of models and is therefore also referred to as black box interpretability models [200]. These solutions provide an interpretable model that explains prediction of the black box in human-understandable terms [90]. Outcome explainers are further classified into *local* and *global* explainers [189], where local explainers are designed to explain the outcome of a specific model and global explainers (also called agnostic explainer [188]) aim to be independent of the machine learning model type or the data [90].

A number of approaches to global explainers have been proposed in the literature. For instance, Local Interpretable Model-agnostic Explanations (LIME) [189] matches the outcome of any classifier using the local approximation of the model's outcome with an explanation model. In their proposal, the authors define a desired explanation model and provide the outcome explanation in their desired criteria. *Black Box Explanations through Transparent Approximations (BETA)* [129] proposes to explain the outcome of any classifier using a framework to quantify the unambiguity, fidelity and interpretability. It generates explanations in the form of decision sets (i.e. nested if-else statements which divide data into exclusive regions) and then optimizes with respect to a combination of the mentioned quantities. In this manner, BETA ensures the generated explanation is both understandable and comprehensive.

Local explainers are designed for a specific machine learning models or class of machine learning models. Given the popularity of neural networks, it is not surprising that many local explainable solutions focus on neural networks. For instance, *Deep Learning Important Features (DeepLIFT)* [205] explains neural networks by assigning importance to every feature through back propagation. Building on developments such as [205], *SHapley Additive Explanations (SHAP)* [144]) proposes a unified framework for local explainers which generalises six existing methods, including DeepLIFT. We point to Montavon et al. [159] for a more detailed tutorial on explainable neural network.

Local explainers have also been developed for other machine learning models that neural networks. An approach for clustering algorithms is to represent each cluster with a prototypical example [114], which provides insights into the process that formed each cluster. Other local explainers have been developed for ensemble decision trees [37, 49, 52, 76, 94, 199, 210, 213, 239] and SVN models [73, 149, 174]. *Model explanation* aims to explain the internal operation of the model instead of outcomes. From a question about why the model yields specific outcomes, the question becomes: *why does the model behaves the way it does?*. Obviously, the answer to this question must be expressed in a human-understandable manner. The model explanation approaches can be divided into those that use an explainer that mimics the behaviour of the model and the ones that aim to develop a model that can explain itself (transparent models). In model explanation usually the model behaviour is generated either as a *decision tree* or a series of *rule-based* conditions tailored for a specific type of machine learning model [90]. In addition, there exist solutions that follow an agnostic approach through which any type of model gets to be explained [35, 97, 118, 140, 141].

In the decision tree methods, a decision tree defines the conditions according to which the model decides. Such an approach allows one to represent an approximation of the black-box behaviour in a form which is understandable to humans. Craven et al. [42] propose a comprehensive explanation of the neural network by constructing a decision tree using a special set of queries that they called as Trepan, and a similar approach has been applied in [211]. For neural networks, Krishnan et al. [117] and Johansson et al. [107] use genetic programming to determine the decision tree and also for explanation of tree ensembles (such as random forests) methods have been proposed to explain using decision trees [24, 37, 52, 76, 199].

Rule-based approaches aim to explain the model through a *set of rules* derived from the behaviour of the model. This has for instance been done for neural networks in Craven et al. [43]. They propose a solution to map rule extraction (a search problem) into a learning problem [90]. Other solutions in this category are [7, 106, 240], including proposals explaining Support Vector Machines [73, 149] and tree ensembles [49].

The transparent model approach aims to make a specific model inherently explainable. As a general approach, one aims to express the impact of coefficients and features on the explainability of the model [190]. Also, “Generalised additive models with interactions” (*GA2M*), proposed by Lou et al. [141], is one of the more well-known transparent models. Impact can be expressed through different models, including single decision, decision trees and decision lists ([190]), in which the data is split in a threshold-based mechanism. Transparent models have been researched for different machine learning techniques, including linear models [96] and neural networks [159, 196, 197]. In general, transparent models tend to be less accurate than their non-transparent counterparts, so considerable research is still needed in establishing methods that retain the accuracy levels while achieving explainability [90].

Well thought-through representation of the data, the model, the outcomes and the accuracy is another way to improve explainability. Yang et al. [229] introduce “nutritional labels”, which are a collection of metrics on fairness, stability and explainability of individual ranking based systems such as credit scoring. They implement a web application that demonstrates these metrics as visual widgets. Mitchel et al. [157] suggest another documentation solution named “model cards”. These cards are short documents accompanying trained machine learning models that include benchmarks across different populations, together with the kind of model, its target objective and intended use, as well as the applications that the model should not be used for. It also includes a list of the features that affect the model’s outcomes, as well as a list of ethical considerations and usage recommendations. They propose model cards for visual and natural language processing. A related approach is to visualise the decision making process using *provenance graphs*, a causality graph that is a directed cyclic graph, enriched with annotations capturing information pertaining to pipeline execution. In [163], the provenance graph of a model includes an annotated workflow of the methods of data collection, pre-processing techniques, the training process of the model and finally the parameters used for the model. This results in a complete picture of the machine learning pipeline.

6.3 Auditability in the Model-Centric Stages

We introduced model auditability earlier, in Sec. 4.3, where we briefly introduced techniques to achieve accountability in models. Here we briefly review three kinds of approaches to ensure that the *process* used to create a model is auditable.

Firstly, one may try to formally prove that a prediction is correct. This is the approach taken for instance by SafetyNets, a framework that enables an untrusted server (the cloud) to provide a client with a short mathematical proof of the correctness of inference tasks performed by a deep neural network [75], with negligible impact on the accuracy and efficiency of the model.

A second and emerging strand of research is centred on *data provenance*. This is “information about entities, activities, and people involved in producing a piece of data or thing, which can be used to form assessments about its quality, reliability or trustworthiness.”[164]. Here the idea is that provenance may be able to provide a useful account of all data transformations, e.g. across a complex pipeline and further into the learning algorithm. Singh et al. [206] suggest a legal perspective on *decision provenance*, which entails “using provenance methods to provide information exposing decision pipelines”.

On more technical grounds, contributions concerning provenance focus on multiple aspects. Firstly, one key concern is the reproducibility of the model-generation pipeline. For this, systems have been proposed that attempt to automatically capture some of the metadata pertaining to a model, such as the settings of the deep neural network [198]. Pipelines are typically implemented using bespoke scripts. The VAMSA tool aims to “reverse engineer” the script and identify key data processing operations using a knowledge base of python ML APIs [165].

As a new trend is emerging of automating the pipelines and thus script generation, eg the “Automated Machine Learning” or AutoML movement, extracting provenance may become simpler and more precise. Systems in use by commercial providers, such as Databricks’ FlowML, are able to capture elements of model generation, to the extent necessary to make them reproducible.

In the context of Big Data processing, a number of proposals exist to capture fine-grained provenance during data transformation. These include Titian [104], which however is primarily designed to enable debugging of Big Data processing pipelines, with focus on the Spark distributed data processing engine. A recent “roadmap” for the role of provenance in Spark-based scientific workflows is available from Guedes et al [89]. Still in the context of map-reduce based systems, [91] presents state of the art and research directions for provenance in the context of “Hadoop as a Service”.

Finally, it must be noted that the detailed provenance of a process may expose sensitive elements of the data and its transformations. Privacy preservation for provenance is a well-studied topic but few proposals exist that are sufficiently recent or that have come to fruition. Cheney et al. [36] provide a quick survey on provenance sanitisation. Amongst the prototype tools to create views over provenance data (which is represented as a graph), ProvAbs [156] stands out as it is based on the W3 standard PROV data model for describing provenance.

6.4 Safety & Security in the Model-Centric Stages

Whereas for the data-centric stages we discussed issues in data confidentiality and integrity that are relatively well-understood, in the model-centric stages we come across issues that are traditionally less deeply explored in computing. We will look at approaches that are directly related to the models and algorithms of a system. Also, adversarial machine learning attacks and defences are discussed in data-centric stages (see Section 5.4; therefore, we will not discuss them in this section). The technologies mostly relate to the vulnerability as well as the protection in real-life operation, that is, in the Inference stage of the chain of trust in Figure 1. For example, an attacker may target an intrusion detection system that is actively examining a network to detect malicious activities, this means the system is at the inference stage. Such intrusion detection system relies on a set of rules that are fixed and tested based previously in training and testing stages. The attacker is motivated to explore the possibility of finding a network pattern that would bypass the intrusion detection system’s set of rules (the attack that is known as *evasion attacks* [138] in the literature). If the attacker has full knowledge of the intrusion detection system, then the attempts to discover the intrusive network pattern is done under the *white-box* approach. In contrast, the attacker with no knowledge of the model is determined to rely on *black-box* approach. Considering ML model h

trained with parameters θ, h_θ), white-box attacker would know part or all of either h and θ and black-box attacker does not have any prior knowledge about the model or its parameters.

The most common threat in the model-centric stages are spoofing attacks and inversion attacks. *Spoofing attacks* refer to inserting made-up data to change the outcomes of the model in a manner advantageous to the attacker. Compared to the data poisoning attack, the attack we consider under spoofing change the data in a semantically meaningful way, namely to purposely change the outcome. Spoofing attacks come in two incarnations, evasion and impersonation. In case of *evasion*, the attacker generates an adversarial set of samples to evade detection by the model. These attacks mainly compromise information-security focused applications such as malware [207] and spam [223] detection systems. The real-world implementation of these attacks has received considerable attention in the literature [16, 134, 203, 236]. *Impersonation attacks* spoof the data with the aim to pretend that the data input is from a genuine source. In a typical impersonation scenario, attackers covertly pretend to be the victim. Then, they inject malicious samples to the model as if it is generated by the victim; resulting into different classification outcome than the original data [34, 131, 176]. This way, the attackers can gain the same access in the system as the victim's [138]. These attacks has been implemented in image recognition [154], malware detection [148, 228] and intrusion detection system [112]. Furthermore, they have been proved to be effective in deep neural networks (DNN) since they tend to compact the features and thus, the attacker can potentially more easily modify key features to obtain their purpose [14, 160, 162, 172, 203].

In inversion attacks the attacker first gather information about the model operation, e.g., through its application program interface, and then use the insights gained to attack the model [204, 225]. Inversion attacks can be implemented in white-box or black-box settings, depending on the level of information that the attacker is able collect [67]. Such attacks become even more powerful if they utilize data leakage vulnerabilities [67, 68, 215] or misclassification errors [68]. In an interesting demonstration of this attack, Tramèr et al. [215] specifically targeted cloud based machine learning services to design attacks using the confidence values output by the cloud-based models in Amazon and Google.

Countermeasures against such threats are not straightforward. One of the challenges in model-centric security solutions is the diversity of models in machine learning and proposed solutions are usually only applicable to a single model class. However, also the mere application of known techniques such as cryptography to model operations instead of data is not trivial by any stretch. In model-centric solutions calculating over encrypted values relies on *homomorphic encryption* [101]. In homomorphic encryption, one can perform mathematical operations on encrypted data without the need to decrypt them. The model operates based on a joint effort by multiple parties which have partial knowledge about the system. This way, the whole system operates correctly without leaking sensitive information about its internal functionality and the data to any parties, including the malicious ones. Shokri et al. [175] propose a deep-learning privacy persevering based on homomorphic encryption. They propose a system that works on a server that is honest-but-curious. A variety of homomorphic encryption applications, such as multi-party computation [46, 221], classification over homomorphic encryption data [6], homomorphic encryption based neural networks [78] and a homomorphic encryption version of k-means clustering [230]. Homomorphic encryption has large potential in cloud-based machine-learning-as-a-service operations, but efficiency remains a bottleneck. In this context, Ohrimenko et al. [175] is of interest. It proposes an alternative to homomorphic encryption using a secure processor architecture based on SGX-processors in order to perform various types of machine learning models including neural networks, SVM, matrix factorization, decision-trees and k-means clustering. Their evaluations show that their solution bears relatively low overhead compared to solutions based on homomorphic encryption.

In addition to above confidentiality issues, the literature also proposes a number of solutions to maintain the integrity of the model outcomes. Demontis et al. [48] present a secure SVM training procedure they call sec-SVM, which learns more evenly-distributed features in linear classifiers. Biggio et al. [15, 18] have a series of solutions of protective algorithms to safeguard the training stage against data poisoning attacks. They use techniques such as random subspace method, bootstrap aggregating and bagging ensembles to search for outliers during model training. In the testing and inference stages most of the solutions are based on game theory, to model the potential behaviours of attacker and defender [138]. Globerson et al. [79] use a min-max approach to analyse the impact of all possible feature manipulations in the testing stage for SVM algorithm, further improved by Teo et al. [212] to consider additional attack scenarios. Brückner et al. [26] suggest a novel adversarial prediction system called Stackelberg Games and provide an algorithm based on Nash equilibrium [25] called NashSVM. Bulò et al. [27] improve [26] with a randomized strategy selection in their prediction system, thus improving the trade-off options between attack detection and false alarm.

A quite different approach to preserve the integrity of the model operation is to purposely account for *adversarial* settings in the training, testing and inferring stages. In these approaches one retrains the model with adversarial samples [138, 209] so that the retrained model detects anomalies [161, 238]. Goodfellow et al. [84] showed that this can be done reasonably efficiently, reduce the false recognition rate. Of course, the inherent limitation of such approach relies on good adversarial samples, in all permutations, and the methods is therefore not suited for unknown or zero-day attack scenarios.

A final opportunity to protect against manipulated results is to correct the output of the model, or at least warn for suspicious outputs. Gu et al. [88] construct a deep contractive network that adds a penalty score in the model output, and approach mostly aimed at deep learning models. Papernot et al. [179] propose a variant called defensive distillation to defend DNN models, but some concerns about the effectiveness of this method remain [33] arguing about the security of [179].

7 OPEN PROBLEMS

There is a substantial amount of research being invested in each of the four FEAS aspects of building trustworthy machine learning systems: fairness, ethics, auditability and safety/security/privacy. For each of the FEAS properties research challenges are being pursued by researchers world-wide, and in this paper we want to highlight the open problems that exist in combining the resulting research contributions to assist data and software engineers in building trustworthy AI-based systems that fulfil the objectives of the policy frameworks discussed in Section 2.1.

Definitions and Implementation of Fairness. It is probably only natural that, in a fast emerging research area such as trustworthy machine, learning terminology as well as key approaches have not yet converged. This is most apparent with respect to the fairness definitions we surveyed in Section 4.1, of which there are tens available [218]. It is theoretically impossible to satisfy all fairness definitions at the same time [70]. This ambiguity in definitions can have serious implications. A prime example is the well-known COMPAS case, which concerns an approach to prisoner recidivism prediction used in several US states. The reported unfairness across ethnic groups was rebutted by the software company based on its own, different, fairness definition [54, 65, 130]. In addition, reliable practical usage of fairness definitions is not without challenges. Incorrect application of fairness metrics may result in increased discrimination [122] and fairness metrics may be prone to result instability, that is, a small change in training set may have noticeable impact on the discrimination performance [70]. It is clear developers of machine learning systems can easily be left confused about what definitions and associated algorithms and software to apply. There is

a clear open problem in establishing an evidence-based implementation framework for fairness in machine learning, covering definitions of metrics, their semantic subtleties with respect to the discrimination they address, and the associated real-life implementation pitfalls for builders of machine learning systems.

Integration of Trustworthiness and Accuracy. Machine learning algorithms are typically evaluated with respect to accuracy and performance first, which is understandable if models are developed isolated from the context of a real-life system. When using machine learning systems in real-life, trustworthiness often becomes as or more important than accuracy and performance. If the FEAS properties are considered too late in the design of algorithm and system, it may not be possible to adjust the machine learning system satisfactory. There is no general mechanism to balance accuracy and fairness for a set of ML models, so far the research on this topic is limited to specific models and conditions. This implies a clear research challenge in trustworthy machine learning, namely the integrated consideration of accuracy, performance and FEAS properties. This requires development of trade-off approaches with respect to metrics representing the different objectives, and the development of combined metrics. An associated open problem is the application of known trustworthiness techniques to AI-based systems in general and machine-learning pipelines in particular. In particular, systematic treatments such as advocated by [8] for designing dependable systems needs to be further expanded for AI-based systems.

Data Engineering Pathways to Implement Policy Frameworks. The chain of trust introduced in Section 3 asks a data engineer and system builder to consider a sequence of data processing steps, which requires to transform the datasets in ones stage into a form that can be used as input to a later stage. There currently are no well-established data and software engineering approaches for such data transformation in relation to achieving the desired combined accuracy, performance and FEAS properties. Not only are there gaps in knowledge of how individual data or processing choices can affect the outcome, but the formal analysis of how to combine data and processing has not been much explored. There is a need for concerted efforts in creating tools that help service providers and machine learning system builders communicate and to help the latter engineer such requirements into model design. This requires research in formalisation and tooling, which would beneficially impact all stakeholders, from data scientists to service providers and end users.

Integration of Social Science Theories and Perspectives. In this paper, we discussed, on the one hand, the various policy frameworks that are emerging and, on the other hand, the relevant FEAS technologies that are available. To understand best how to bridge societal concerns and technological possibilities research should be informed by multi-disciplinary perspectives. As an example, [214] conducts initial research in this space in relation to trust in AI-based services. In that work, social science theories for trust development are linked up with technology quality requirements. Evidence-based research in social-science, psychology and human-computer interaction has the potential to steer technology development and demonstrate its abilities or limitations in achieving societal goals. Through quantitative as well as qualitative research the impact of technologies with respect to their societal impact can be measured and better understood.

Education. Veale et al. [217] interview 27 public sector machine learning practitioners in five countries about their understanding of public values such as fairness and accountability. Their results confirm the disconnect between individual perspectives and technological properties and point out the challenges that come with this in developing ethical machine learning systems. Similarly, Holstein et al. [100] interview 25 machine learning practitioners and survey 267 more to systematically investigate the implementation of fair machine learning by commercial product teams. They also highlight the difference between real-world challenges and the provided solutions.

Improving this situation could start with education of all stakeholders in machine learning systems, including software and data engineers, for instance through curriculum development such as demonstrated by [11].

8 CONCLUSION

This paper reviews the landscape and the most pertinent technologies required for building machine learning systems that aim to fulfil the objectives of the many AI-related policy frameworks that are emerging in relation to fairness and ethics. The paper identifies four main system properties to implement such systems, namely fairness, explainability, auditability and safety/security (FEAS), and provides a review of the main trends and technology categories in each of these. It surveys the FEAS technologies with respect to all stages of the machine learning chain of trust, subdivided in data-centric and model-centric technologies, respectively. The paper argues that in future work, the emphasis should not only be on individual FEAS technologies but particularly on how these technologies integrate with accuracy and efficiency consideration, can be translated in data and software engineering tools that integrate concerns across the various stages of the machine learning life cycle and in cross-disciplinary research to more profoundly connect FEAS technologies with the AI policy framework objectives they ultimately set out to achieve.

ACKNOWLEDGMENTS

Research supported by UK EPSRC, under grant EP/R033595, "Trust Engineering for the Financial Industry", <https://gow.epsrc.ukri.org/NGBOViewGrant.aspx?GrantRef=EP/R033595/1>.

REFERENCES

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 308–318.
- [2] Mhairi Aitken, Ehsan Toreini, Peter Carmichael, Kovila Coopamootoo, Karen Elliott, and Aad van Moorsel. 2020. Establishing a social licence for Financial Technology: Reflections on the role of the private sector in pursuing ethical data practices. *Big Data & Society* 7, 1 (2020), 2053951720908892.
- [3] Ahmed M Alaa and Mihaela van der Schaar. 2019. Demystifying Black-box Models with Symbolic Metamodels. In *Advances in Neural Information Processing Systems 32*, H Wallach, H Larochelle, A Beygelzimer, F d\textquotesingle Alché-Buc, E Fox, and R Garnett (Eds.). Curran Associates, Inc., 11304–11314. <http://papers.nips.cc/paper/9308-demystifying-black-box-models-with-symbolic-metamodels.pdf>
- [4] Jeanne Altmann. 1974. Observational study of behavior: sampling methods. *Behaviour* 49, 3-4 (1974), 227–266.
- [5] Ross Anderson. 2008. *Security engineering*. John Wiley & Sons.
- [6] Louis J M Aslett, Pedro M Esperança, and Chris C Holmes. 2015. Encrypted statistical machine learning: new privacy preserving methods. *arXiv preprint arXiv:1508.06845* (2015).
- [7] M Gethsiyal Augasta and Thangairulappan Kathirvalavakumar. 2012. Reverse engineering the neural networks for rule extraction in classification problems. *Neural processing letters* 35, 2 (2012), 131–150.
- [8] A Avizienis, J . Laprie, B Randell, and C Landwehr. 2004. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing* 1, 1 (1 2004), 11–33. <https://doi.org/10.1109/TDSC.2004.2>
- [9] Solon Barocas and Andrew D Selbst. 2016. Big data’s disparate impact. *Calif. L. Rev.* 104 (2016), 671.
- [10] Marco Barreno, Blaine Nelson, Russell Sears, Anthony D Joseph, and J Doug Tygar. 2006. Can machine learning be secure?. In *Proceedings of the 2006 ACM Symposium on Information, computer and communications security*. 16–25.
- [11] Jo Bates, David Cameron, Alessandro Checco, Paul Clough, Frank Hopfgartner, Suvodeep Mazumdar, Laura Scaffi, Peter Stordy, and Antonio de la Vega de León. 2020. Integrating FATE/critical data studies into data science curricula: where are we going and how do we get there?. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 425–435.
- [12] John Bethencourt, Amit Sahai, and Brent Waters. 2007. Ciphertext-policy attribute-based encryption. In *2007 IEEE symposium on security and privacy (SP’07)*. IEEE, 321–334.
- [13] Alex Beutel, Jilin Chen, Zhe Zhao, and Ed H Chi. 2017. Data decisions and theoretical implications when adversarially learning fair representations. *arXiv preprint arXiv:1707.00075* (2017).

- [14] Battista Biggio, Samuel Rota Bulò, Ignazio Pillai, Michele Mura, Eyasu Zemene Mequanint, Marcello Pelillo, and Fabio Roli. 2014. Poisoning complete-linkage hierarchical clustering. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*. Springer, 42–52.
- [15] Battista Biggio, Iginio Corona, Giorgio Fumera, Giorgio Giacinto, and Fabio Roli. 2011. Bagging classifiers for fighting poisoning attacks in adversarial classification tasks. In *International workshop on multiple classifier systems*. Springer, 350–359.
- [16] Battista Biggio, Iginio Corona, Davide Maiorca, Blaine Nelson, Nedim Šrndić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. 2013. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases*. Springer, 387–402.
- [17] Battista Biggio, Iginio Corona, Blaine Nelson, Benjamin I P Rubinstein, Davide Maiorca, Giorgio Fumera, Giorgio Giacinto, and Fabio Roli. 2014. Security evaluation of support vector machines in adversarial environments. In *Support Vector Machines Applications*. Springer, 105–153.
- [18] Battista Biggio, Giorgio Fumera, and Fabio Roli. 2010. Multiple classifier systems for robust classifier design in adversarial environments. *International Journal of Machine Learning and Cybernetics* 1, 1-4 (2010), 27–41.
- [19] Battista Biggio, Giorgio Fumera, and Fabio Roli. 2014. Security evaluation of pattern classifiers under attack. *IEEE transactions on knowledge and data engineering* 26, 4 (2014), 984–996.
- [20] Battista Biggio, Blaine Nelson, and Pavel Laskov. 2011. Support vector machines under adversarial label noise. In *Asian Conference on Machine Learning*. 97–112.
- [21] Battista Biggio, Blaine Nelson, and Pavel Laskov. 2012. Poisoning attacks against support vector machines. *arXiv preprint arXiv:1206.6389* (2012).
- [22] Battista Biggio, Konrad Rieck, Davide Ariu, Christian Wressnegger, Iginio Corona, Giorgio Giacinto, and Fabio Roli. 2014. Poisoning Behavioral Malware Clustering. In *Proceedings of the 2014 Workshop on Artificial Intelligent and Security Workshop (AISec '14)*. ACM, New York, NY, USA, 27–36. <https://doi.org/10.1145/2666652.2666666>
- [23] Avrim L Blum and Pat Langley. 1997. Selection of relevant features and examples in machine learning. *Artificial Intelligence* 97, 1-2 (1997), 245–271.
- [24] Leo Breiman. 1984. *Classification and regression trees*. Routledge.
- [25] Michael Brückner, Christian Kanzow, and Tobias Scheffer. 2012. Static prediction games for adversarial learning problems. *Journal of Machine Learning Research* 13, Sep (2012), 2617–2654.
- [26] Michael Brückner and Tobias Scheffer. 2011. Stackelberg games for adversarial prediction problems. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 547–555.
- [27] Samuel Rota Bulò, Battista Biggio, Ignazio Pillai, Marcello Pelillo, and Fabio Roli. 2016. Randomized prediction games for adversarial machine learning. *IEEE transactions on neural networks and learning systems* 28, 11 (2016), 2466–2478.
- [28] ANIRUDH BURMAN. 2019. Will a GDPR-Style Data Protection Law Work for India? (2019).
- [29] Jie Cai, Jiawei Luo, Shulin Wang, and Sheng Yang. 2018. Feature selection in machine learning: A new perspective. *Neurocomputing* 300 (2018), 70–79. <https://doi.org/10.1016/j.neucom.2017.11.077>
- [30] Toon Calders and Sico Verwer. 2010. Three naive Bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery* 21, 2 (2010), 277–292.
- [31] Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R Varshney. 2017. Optimized pre-processing for discrimination prevention. In *Advances in Neural Information Processing Systems*. 3992–4001.
- [32] Nicholas Carlini, Pratyush Mishra, Tavish Vaidya, Yuankai Zhang, Micah Sherr, Clay Shields, David Wagner, and Wencho Zhou. 2016. Hidden voice commands. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*. 513–530.
- [33] Nicholas Carlini and David Wagner. 2016. Defensive distillation is not robust to adversarial examples. *arXiv preprint arXiv:1607.04311* (2016).
- [34] Nicholas Carlini and David Wagner. 2017. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 39–57.
- [35] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. 2015. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1721–1730.
- [36] James Cheney and Roly Perera. 2015. An Analytical Survey of Provenance Sanitization. In *Provenance and Annotation of Data and Processes*, Bertram Ludäscher and Beth Plale (Eds.). Springer International Publishing, Cham, 113–126.
- [37] H A Chipman, E I George, and R E McCulloh. 1998. Making sense of a forest of trees. *Computing Science and Statistics* (1998), 84–92.
- [38] Jaegul Choo and Shixia Liu. 2018. Visual analytics for explainable deep learning. *IEEE computer graphics and applications* 38, 4 (2018), 84–92.

- [39] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data* 5, 2 (2017), 153–163.
- [40] Sam Corbett-Davies and Sharad Goel. 2018. The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning. *arXiv preprint arXiv:1808.00023* (2018).
- [41] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. 2017. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 797–806.
- [42] Mark Craven and Jude W Shavlik. 1996. Extracting tree-structured representations of trained networks. In *Advances in neural information processing systems*. 24–30.
- [43] Mark W Craven and Jude W Shavlik. 1994. Using sampling and queries to extract rules from trained neural networks. In *Machine learning proceedings 1994*. Elsevier, 37–45.
- [44] Sven F Crone, Stefan Lessmann, and Robert Stahlbock. 2006. The impact of preprocessing on data mining: An evaluation of classifier sensitivity in direct marketing. *European Journal of Operational Research* 173, 3 (2006), 781–800. <https://doi.org/10.1016/j.ejor.2005.07.023>
- [45] Brian d'Alessandro, Cathy O'Neil, and Tom LaGatta. 2017. Conscientious classification: A data scientist's guide to discrimination-aware classification. *Big data* 5, 2 (2017), 120–134.
- [46] Ivan Damgård, Valerio Pastro, Nigel Smart, and Sarah Zakarias. 2012. Multiparty computation from somewhat homomorphic encryption. In *Annual Cryptology Conference*. Springer, 643–662.
- [47] Jeffrey Dastin. 2018. Amazon scraps secret AI recruiting tool that showed bias against women. *San Fransico, CA: Reuters*. Retrieved on October 9 (2018), 2018.
- [48] Ambra Demontis, Marco Melis, Battista Biggio, Davide Maiorca, Daniel Arp, Konrad Rieck, Igino Corona, Giorgio Giacinto, and Fabio Roli. 2017. Yes, machine learning can be more secure! a case study on android malware detection. *IEEE Transactions on Dependable and Secure Computing* (2017).
- [49] Houtao Deng. 2019. Interpreting tree ensembles with intrees. *International Journal of Data Science and Analytics* 7, 4 (2019), 277–287.
- [50] Nicholas Diakopoulos. 2016. Accountability in Algorithmic Decision Making. *Commun. ACM* 59, 2 (Jan. 2016), 56–62. <https://doi.org/10.1145/2844110>
- [51] Thomas G Dietterich and Eun Bae Kong. 1995. *Machine learning bias, statistical bias, and statistical variance of decision tree algorithms*. Technical Report. Technical report, Department of Computer Science, Oregon State University.
- [52] Pedro Domingos. 1998. Knowledge discovery via multiple models. *Intelligent Data Analysis* 2, 1-4 (1998), 187–202.
- [53] Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608* (2017).
- [54] Julia Dressel and Hany Farid. 2018. The accuracy, fairness, and limits of predicting recidivism. *Science advances* 4, 1 (2018), eaao5580.
- [55] Harris Drucker, Donghui Wu, and Vladimir N Vapnik. 1999. Support vector machines for spam categorization. *IEEE Transactions on Neural networks* 10, 5 (1999), 1048–1054.
- [56] Rehab Duwairi and Mahmoud El-Orfali. 2014. A study of the effects of preprocessing strategies on sentiment analysis for Arabic text. *Journal of Information Science* 40, 4 (2014), 501–513.
- [57] Cynthia Dwork. 2011. Differential privacy. *Encyclopedia of Cryptography and Security* (2011), 338–340.
- [58] Cynthia Dwork, Nicole Immorlica, Adam Tauman Kalai, and Max Leiserson. 2018. Decoupled classifiers for group-fair and efficient machine learning. In *Conference on Fairness, Accountability and Transparency*. 119–133.
- [59] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.
- [60] Bradley Efron. 1994. Missing data, imputation, and the bootstrap. *J. Amer. Statist. Assoc.* 89, 426 (1994), 463–475.
- [61] Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, Pascal Vincent, and Samy Bengio. 2010. Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research* 11, Feb (2010), 625–660.
- [62] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*. ACM, 1054–1067.
- [63] Reuben Feinman, Ryan R Curtin, Saurabh Shintre, and Andrew B Gardner. 2017. Detecting adversarial samples from artifacts. *arXiv preprint arXiv:1703.00410* (2017).
- [64] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 259–268.
- [65] Avi Feller, Emma Pierson, Sam Corbett-Davies, and Sharad Goel. 2016. A computer program used for bail and sentencing decisions was labeled biased against blacks. It's actually not that clear. *The Washington Post* (2016).

- [66] Jessica Fjeld, Hannah Hilligoss, Nele Achten, Maia Levy Daniel, Sally Kagay, and Joshua Feldman. 2019. Principled Artificial Intelligence: Mapping Consensus and Divergence in Ethical and Rights-Based Approaches. (2019). <https://ai-hr.cyber.harvard.edu/>
- [67] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1322–1333.
- [68] Matthew Fredrikson, Eric Lantz, Somesh Jha, Simon Lin, David Page, and Thomas Ristenpart. 2014. Privacy in pharmacogenetics: An end-to-end case study of personalized warfarin dosing. In *23rd USENIX Security Symposium (USENIX Security 14)*. 17–32.
- [69] Alex A Freitas. 2014. Comprehensive classification models: a position paper. *ACM SIGKDD explorations newsletter* 15, 1 (2014), 1–10.
- [70] Sorelle A Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. 2016. On the (im) possibility of fairness. *arXiv preprint arXiv:1609.07236* (2016).
- [71] Sorelle A Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. 2016. On the (im) possibility of fairness. *arXiv preprint arXiv:1609.07236* (2016).
- [72] Sorelle A Friedler, Carlos Scheidegger, Suresh Venkatasubramanian, Sonam Choudhary, Evan P Hamilton, and Derek Roth. 2019. A comparative study of fairness-enhancing interventions in machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, 329–338.
- [73] Glenn Fung, Sathyakama Sandilya, and R Bharat Rao. 2005. Rule extraction from linear support vector machines. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*. ACM, 32–40.
- [74] Pratik Gajane and Mykola Pechenizkiy. 2017. On formalizing fairness in prediction with machine learning. *arXiv preprint arXiv:1710.03184* (2017).
- [75] Zahra Ghodsi, Tianyu Gu, and Siddharth Garg. 2017. Safetynets: Verifiable execution of deep neural networks on an untrusted cloud. In *Advances in Neural Information Processing Systems*. 4672–4681.
- [76] Robert D Gibbons, Giles Hooker, Matthew D Finkelman, David J Weiss, Paul A Pilkonis, Ellen Frank, Tara Moore, and David J Kupfer. 2013. The CAD-MDD: A computerized adaptive diagnostic screening tool for depression. *The Journal of clinical psychiatry* 74, 7 (2013), 669.
- [77] Eli Gibson, Wenqi Li, Carole Sudre, Lucas Fidon, Dzshoskun I Shakir, Guotai Wang, Zach Eaton-Rosen, Robert Gray, Tom Doel, Yipeng Hu, et al. 2018. NiftyNet: a deep-learning platform for medical imaging. *Computer methods and programs in biomedicine* 158 (2018), 113–122.
- [78] Ran Gilad-Bachrach, Nathan Dowlin, Kim Laine, Kristin Lauter, Michael Naehrig, and John Wernsing. 2016. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *International Conference on Machine Learning*. 201–210.
- [79] Amir Globerson and Sam Roweis. 2006. Nightmare at test time: robust learning by feature deletion. In *Proceedings of the 23rd international conference on Machine learning*. ACM, 353–360.
- [80] Gabriel Goh, Andrew Cotter, Maya Gupta, and Michael P Friedlander. 2016. Satisfying real-world goals with dataset constraints. In *Advances in Neural Information Processing Systems*. 2415–2423.
- [81] Carlos Adriano Gonçalves, Celia Talma Gonçalves, Rui Camacho, and Eugenio C Oliveira. 2010. The impact of Pre-Processing on the Classification of MEDLINE Documents. *Pattern Recognition in Information Systems, Proceedings of the 10th International Workshop on Pattern Recognition in Information Systems, PRIS 2010, In conjunction with ICEIS 2010* (2010), 10.
- [82] Carlos Vladimiro González Zelaya. 2019. Towards Explaining the Effects of Data Preprocessing on Machine Learning. *2019 IEEE 35th International Conference on Data Engineering (ICDE)* (2019).
- [83] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- [84] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
- [85] Bryce Goodman and Seth Flaxman. 2016. European Union regulations on algorithmic decision-making and a "right to explanation". *Proceedings of the 2016 ICML Workshop on Human Interpretability in Machine Learning (WHI 2016)* (jun 2016). arXiv:1606.08813 <http://arxiv.org/abs/1606.08813>
- [86] Nina Grgic-Hlaca, Muhammad Bilal Zafar, Krishna P Gummadi, Adrian Weller, Nina Grgic-Hlaca, Muhammad Bilal Zafar, Krishna P Gummadi, and Adrian Weller. 2016. The case for process fairness in learning: Feature selection for fair decision making. In *NIPS Symposium on Machine Learning and the Law*, Vol. 1. 2. <https://doi.org/10.2174/1381612821666150416100516>
- [87] Kathrin Grosse, Praveen Manoharan, Nicolas Papernot, Michael Backes, and Patrick McDaniel. 2017. On the (statistical) detection of adversarial examples. *arXiv preprint arXiv:1702.06280* (2017).

- [88] Shixiang Gu and Luca Rigazio. 2014. Towards deep neural network architectures robust to adversarial examples. *arXiv preprint arXiv:1412.5068* (2014).
- [89] T. Guedes, V. Silva, M. Mattoso, M. V. N. Bedo, and D. de Oliveira. 2018. A Practical Roadmap for Provenance Capture and Data Analysis in Spark-Based Scientific Workflows. In *2018 IEEE/ACM Workflows in Support of Large-Scale Science (WORKS)*. 31–41. <https://doi.org/10.1109/WORKS.2018.00009>
- [90] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. 2018. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)* 51, 5 (2018), 93.
- [91] Himanshu Gupta, Sameep Mehta, Sandeep Hans, Bapi Chatterjee, Pranay Lohia, and C Rajmohan. 2017. Provenance in context of Hadoop as a Service (HaaS)-State of the Art and Research Directions. In *2017 IEEE 10th International Conference on Cloud Computing (CLOUD)*. IEEE, 680–683.
- [92] Sara Hajian, Francesco Bonchi, and Carlos Castillo. 2016. Algorithmic bias: From discrimination discovery to fairness-aware data mining. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 2125–2126.
- [93] Sara Hajian, Josep Domingo-Ferrer, and Oriol Farràs. 2014. Generalization-based privacy preservation and discrimination prevention in data publishing and mining. *Data Mining and Knowledge Discovery* 28, 5-6 (2014), 1158–1188.
- [94] Satoshi Hara and Kohei Hayashi. 2016. Making tree ensembles interpretable. *arXiv preprint arXiv:1606.05390* (2016).
- [95] Moritz Hardt, Eric Price, Nati Srebro, and others. 2016. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*. 3315–3323.
- [96] Stefan Haufe, Frank Meinecke, Kai Görgen, Sven Dähne, John-Dylan Haynes, Benjamin Blankertz, and Felix Bießmann. 2014. On the interpretation of weight vectors of linear models in multivariate neuroimaging. *Neuroimage* 87 (2014), 96–110.
- [97] Andreas Henelius, Kai Puolamäki, Henrik Boström, Lars Asker, and Panagiotis Papapetrou. 2014. A peek into the black box: exploring classifiers by randomization. *Data mining and knowledge discovery* 28, 5-6 (2014), 1503–1529.
- [98] Bernease Herman. 2017. The promise and peril of human evaluation for model interpretability. *arXiv preprint arXiv:1711.07414* (2017).
- [99] Ehsan Hesamifard, Hassan Takabi, Mehdi Ghasemi, and Rebecca N Wright. 2018. Privacy-preserving machine learning as a service. *Proceedings on Privacy Enhancing Technologies* 2018, 3 (2018), 123–142.
- [100] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé III, Miro Dudik, and Hanna Wallach. 2019. Improving fairness in machine learning systems: What do industry practitioners need?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 600.
- [101] Weiwei Hu and Ying Tan. 2017. Generating adversarial malware examples for black-box attacks based on GAN. *arXiv preprint arXiv:1702.05983* (2017).
- [102] Ling Huang, Anthony D Joseph, Blaine Nelson, Benjamin I P Rubinstein, and J D Tygar. 2011. Adversarial machine learning. In *Proceedings of the 4th ACM workshop on Security and artificial intelligence*. ACM, 43–58.
- [103] Johan Huysmans, Karel Dejaeger, Christophe Mues, Jan Vanthienen, and Bart Baesens. 2011. An empirical evaluation of the comprehensibility of decision table, tree and rule based predictive models. *Decision Support Systems* 51, 1 (2011), 141–154.
- [104] Matteo Interlandi, Kshitij Shah, Sai Deep Tetali, Muhammad Ali Gulzar, Seunghyun Yoo, Miryung Kim, Todd Millstein, and Tyson Condie. 2015. Titian: Data Provenance Support in Spark. *Proc. VLDB Endow.* 9, 3 (nov 2015), 216–227. <https://doi.org/10.14778/2850583.2850595>
- [105] Masoumeh Rezaei Jam, Leili Mohammad Khanli, Morteza Sargolzaei Javan, and Mohammad Kazem Akbari. 2014. A survey on security of Hadoop. In *2014 4th International Conference on Computer and Knowledge Engineering (ICCKE)*. IEEE, 716–721.
- [106] U Johansson, R König, and L Niklasson. 2003. Rule extraction from trained neural networks using genetic programming. In *13th International Conference on Artificial Neural Networks*. 13–16.
- [107] Ulf Johansson and Lars Niklasson. 2009. Evolving decision trees using oracle guides. In *2009 IEEE Symposium on Computational Intelligence and Data Mining*. IEEE, 238–244.
- [108] Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. 2016. Rawlsian Fairness for Machine Learning. *FATML* (2016), 1–26. <http://arxiv.org/abs/1610.09559>
- [109] Faisal Kamiran and Toon Calders. 2012. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems* (2012). <https://doi.org/10.1007/s10115-011-0463-8>
- [110] Faisal Kamiran, Toon Calders, and Mykola Pechenizkiy. 2010. Discrimination aware decision tree learning. In *2010 IEEE International Conference on Data Mining*. IEEE, 869–874.
- [111] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2012. Fairness-aware classifier with prejudice remover regularizer. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 35–50.

- [112] H Gunes Kayacik, A Nur Zincir-Heywood, and Malcolm I Heywood. 2007. Automatically evading IDS using GP authored attacks. *IEEE*, 153–160.
- [113] Michael Kearns and Ming Li. 1993. Learning in the presence of malicious errors. *SIAM J. Comput.* 22, 4 (1993), 807–837.
- [114] Been Kim, Cynthia Rudin, and Julie A Shah. 2014. The bayesian case model: A generative approach for case-based reasoning and prototype classification. In *Advances in Neural Information Processing Systems*. 1952–1960.
- [115] Marius Kloft and Pavel Laskov. 2010. Online anomaly detection under adversarial impact. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. 405–412.
- [116] Igor Kononenko and others. 2010. An efficient explanation of individual classifications using game theory. *Journal of Machine Learning Research* 11, Jan (2010), 1–18.
- [117] R Krishnan, G Sivakumar, and P Bhattacharya. 1999. Extracting decision trees from trained neural networks. *Pattern recognition* 32, 12 (1999).
- [118] Sanjay Krishnan and Eugene Wu. 2017. Palm: Machine learning explanations for iterative debugging. In *Proceedings of the 2nd Workshop on Human-In-the-Loop Data Analytics*. 1–6.
- [119] Alex Krizhevsky, Geoffrey Hinton, and others. 2009. *Learning multiple layers of features from tiny images*. Technical Report. Citeseer.
- [120] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2016. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533* (2016).
- [121] Matt Kusner, Jacob Gardner, Roman Garnett, and Kilian Weinberger. 2015. Differentially private Bayesian optimization. In *International Conference on Machine Learning*. 918–927.
- [122] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual fairness. In *Advances in Neural Information Processing Systems*. 4066–4076.
- [123] Matt J Kusner and Joshua R Loftus. 2020. The long road to fairer algorithms.
- [124] Sergei O Kuznetsov. 2001. Machine learning on the basis of formal concept analysis. *Automation and Remote Control* 62, 10 (2001), 1543–1564.
- [125] Sergei O Kuznetsov. 2004. Machine learning and formal concept analysis. In *International Conference on Formal Concept Analysis*. Springer, 287–312.
- [126] Guillermo Lafuente. 2015. The big data security challenge. *Network security* 2015, 1 (2015), 12–14.
- [127] Soumendra Nath Lahiri. 2013. *Resampling methods for dependent data*. Springer Science & Business Media.
- [128] Ricky Laishram and Vir Virander Phoha. 2016. Curie: A method for protecting SVM Classifier from Poisoning Attack. *arXiv preprint arXiv:1606.01584* (2016).
- [129] Himabindu Lakkaraju, Ece Kamar, Rich Caruana, and Jure Leskovec. 2017. Interpretable & Explorable Approximations of Black Box Models. (jul 2017). *arXiv:1707.01154* <http://arxiv.org/abs/1707.01154>
- [130] Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. 2016. How we analyzed the COMPAS recidivism algorithm. *ProPublica* (5 2016) 9 (2016).
- [131] Pavel Laskov and others. 2014. Practical evasion of a learning-based classifier: A case study. In *2014 IEEE symposium on security and privacy*. IEEE, 197–211.
- [132] Guillaume Lemaître, Fernando Nogueira, and Christos K Aridas. 2017. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *The Journal of Machine Learning Research* 18, 1 (2017), 559–563.
- [133] Benjamin Letham, Cynthia Rudin, Tyler H McCormick, David Madigan, and others. 2015. Interpretable classifiers using rules and bayesian analysis: Building a better stroke prediction model. *The Annals of Applied Statistics* 9, 3 (2015), 1350–1371.
- [134] Bo Li and Yevgeniy Vorobeychik. 2014. Feature cross-substitution in adversarial classification. In *Advances in neural information processing systems*. 2087–2095.
- [135] Meng Li, Jiajun Wei, Xi Zheng, and Matthew L Bolton. 2017. A formal machine-learning approach to generating human-machine interfaces from task models. *IEEE Transactions on Human-Machine Systems* 47, 6 (2017), 822–833.
- [136] Xin Li, Traci J Hess, and Joseph S Valacich. 2008. Why do we trust new technology? A study of initial trust formation with organizational information systems. *The Journal of Strategic Information Systems* 17, 1 (2008), 39–71.
- [137] Zachary C Lipton. 2016. The mythos of model interpretability. *arXiv preprint arXiv:1606.03490* (2016).
- [138] Qiang Liu, Pan Li, Wentao Zhao, Wei Cai, Shui Yu, and Victor C M Leung. 2018. A survey on security threats and defensive techniques of machine learning: A data driven view. *IEEE access* 6 (2018), 12103–12117.
- [139] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. 2016. Delving into transferable adversarial examples and black-box attacks. *arXiv preprint arXiv:1611.02770* (2016).
- [140] Yin Lou, Rich Caruana, and Johannes Gehrke. 2012. Intelligible models for classification and regression. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. 150–158.
- [141] Yin Lou, Rich Caruana, Johannes Gehrke, and Giles Hooker. 2013. Accurate intelligible models with pairwise interactions. *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining* -

- KDD '13* (2013), 623. <https://doi.org/10.1145/2487575.2487579>
- [142] Daniel Lowd. 2005. Good word attacks on statistical spam filters. In *In Proceedings of the Second Conference on Email and Anti-Spam (CEAS)*. Citeseer.
 - [143] Kristian Lum and James Johndrow. 2016. A statistical framework for fair predictive algorithms. *arXiv preprint arXiv:1610.08077* (2016).
 - [144] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*. 4765–4774.
 - [145] Binh Thanh Luong, Salvatore Ruggieri, and Franco Turini. 2011. k-NN as an implementation of situation testing for discrimination discovery and prevention. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 502–510.
 - [146] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, Nov (2008), 2579–2605.
 - [147] David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. 2019. Fairness through causal awareness: Learning causal latent-variable models for biased data. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 349–358.
 - [148] Davide Maiorca, Igino Corona, and Giorgio Giacinto. 2013. Looking at the bag is not enough to find the bomb: an evasion of structural methods for malicious pdf files detection. In *Proceedings of the 8th ACM SIGSAC symposium on Information, computer and communications security*. ACM, 119–130.
 - [149] David Martens, Bart Baesens, Tony Van Gestel, and Jan Vanthienen. 2007. Comprehensible credit scoring models using rule extraction from support vector machines. *European journal of operational research* 183, 3 (2007), 1466–1476.
 - [150] Frank McSherry. 2017. Statistical inference considered harmful. <https://github.com/frankmcsherry/blog/blob/master/posts/2016-06-14.md>
 - [151] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2019. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635* (2019).
 - [152] Shike Mei and Xiaojin Zhu. 2015. The security of latent dirichlet allocation. In *Artificial Intelligence and Statistics*. 681–689.
 - [153] Jan Hendrik Metzen, Tim Genewein, Volker Fischer, and Bastian Bischoff. 2017. On detecting adversarial perturbations. *arXiv preprint arXiv:1702.04267* (2017).
 - [154] Andrew Milich and Michael Karr. [n. d.]. Eluding Mass Surveillance: Adversarial Attacks on Facial Recognition Models. ([n. d.]).
 - [155] Tim Miller. 2018. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* (2018).
 - [156] P. Missier, J. Bryans, C. Gamble, V. Curcin, and R. Danger. 2015. *Provabs: Model, policy, and tooling for abstracting PROV graphs*. Vol. 8628. 3–15 pages. https://doi.org/10.1007/978-3-319-16462-5_1
 - [157] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, 220–229.
 - [158] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. 2018. *Foundations of machine learning*. MIT press.
 - [159] Grégoire Montavon, Wojciech Samek, and Klaus-Robert Müller. 2018. Methods for interpreting and understanding deep neural networks. *Digital Signal Processing* 73 (2018), 1–15.
 - [160] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. 2017. Universal adversarial perturbations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1765–1773.
 - [161] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. 2016. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2574–2582.
 - [162] Konda Reddy Mopuri, Utsav Garg, and R Venkatesh Babu. 2017. Fast Feature Fool: {A} data independent approach to universal adversarial perturbations. *CoRR* abs/1707.0 (2017). <http://arxiv.org/abs/1707.05572>
 - [163] Luc Moreau, Ben Clifford, Juliana Freire, Joe Futrelle, Yolanda Gil, Paul Groth, Natalia Kwasnikowska, Simon Miles, Paolo Missier, Jim Myers, Beth Plale, Yogesh Simmhan, Eric Stephan, and Jan Van Den Bussche. 2011. The Open Provenance Model core specification (v1.1). *Future Generation Computer Systems* 27, 6 (2011), 743–756. <https://doi.org/10.1016/j.future.2010.07.005>
 - [164] Luc Moreau, Paolo Missier, Khalid Belhajjame, Reza B'Far, James Cheney, Sam Coppens, Stephen Cresswell, Yolanda Gil, Paul Groth, Graham Klyne, Timothy Lebo, Jim McCusker, Simon Miles, James Myers, Satya Sahoo, and Curt Tilmel. 2012. *PROV-DM: The PROV Data Model*. Technical Report. World Wide Web Consortium. <http://www.w3.org/TR/prov-dm/>
 - [165] Mohammad Hossein Namaki, Avriilia Floratou, Fotis Psallidas, Subru Krishnan, Ashvin Agrawal, and Yinghui Wu. 2020. Vamsa: Tracking Provenance in Data Science Scripts. *arXiv:cs.LG/2001.01861*

- [166] Arvind Narayanan. 2018. Translation tutorial: 21 fairness definitions and their politics. In *Proc. Conf. Fairness Accountability Transp., New York, USA*.
- [167] Arvind Narayanan and Vitaly Shmatikov. 2008. Robust de-anonymization of large datasets (how to break anonymity of the Netflix prize dataset). (2008).
- [168] Nina Narodytska and Shiva Kasiviswanathan. 2017. Simple black-box adversarial attacks on deep neural networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE, 1310–1318.
- [169] Blaine Nelson, Marco Barreno, Fuching Jack Chi, Anthony D Joseph, Benjamin I P Rubinstein, Udam Saini, Charles Sutton, J D Tygar, and Kai Xia. 2009. Misleading learners: Co-opting your spam filter. In *Machine learning in cyber trust*. Springer, 17–51.
- [170] James Newsome, Brad Karp, and Dawn Song. 2005. Polygraph: Automatically generating signatures for polymorphic worms. In *2005 IEEE Symposium on Security and Privacy (S&P'05)*. IEEE, 226–241.
- [171] Jerzy Neyman. 1992. On the two different aspects of the representative method: the method of stratified sampling and the method of purposive selection. In *Breakthroughs in Statistics*. Springer, 123–150.
- [172] Anh Nguyen, Jason Yosinski, and Jeff Clune. 2015. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 427–436.
- [173] Harsha Nori, Samuel Jenkins, Paul Koch, and Rich Caruana. 2019. InterpretML: A Unified Framework for Machine Learning Interpretability. (sep 2019). arXiv:1909.09223 <http://arxiv.org/abs/1909.09223>
- [174] Haydymar Núñez, Cecilio Angulo, and Andreu Català. 2002. Rule extraction from support vector machines.. In *Esann*. 107–112.
- [175] Olga Ohrimenko, Felix Schuster, Cédric Fournet, Aastha Mehta, Sebastian Nowozin, Kapil Vaswani, and Manuel Costa. 2016. Oblivious multi-party machine learning on trusted processors. In *25th USENIX Security Symposium (USENIX Security 16)*. 619–636.
- [176] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. 2017. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*. ACM, 506–519.
- [177] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z Berkay Celik, and Ananthram Swami. 2016. The limitations of deep learning in adversarial settings. In *2016 IEEE European symposium on security and privacy (EuroS&P)*. IEEE, 372–387.
- [178] Nicolas Papernot, Patrick McDaniel, Arunesh Sinha, and Michael P Wellman. 2018. SoK: Security and privacy in machine learning. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 399–414.
- [179] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. 2016. Distillation as a defense to adversarial perturbations against deep neural networks. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 582–597.
- [180] Dino Pedreshi, Salvatore Ruggieri, and Franco Turini. 2008. Discrimination-aware data mining. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 560–568.
- [181] Roberto Perdisci, David Dagon, Wenke Lee, Prahlad Fogla, and Monirul Sharif. 2006. Misleading worm signature generators using deliberate noise injection. In *2006 IEEE Symposium on Security and Privacy (S&P'06)*. IEEE, 15–pp.
- [182] Lucille Perreault. 2015. Big Data and Privacy: Control and Awareness Aspects.. In *CONF-IRM*. 15.
- [183] Neoklis Polyzotis, Sudip Roy, Steven Euijong Whang, and Martin Zinkevich. 2018. Data lifecycle challenges in production machine learning: a survey. *ACM SIGMOD Record* 47, 2 (2018), 17–28.
- [184] Bin Qian, Jie Su, Zhenyu Wen, Devki Nandan Jha, Yinhao Li, Yu Guan, Deepak Puthal, Philip James, Renyu Yang, Albert Y Zomaya, et al. 2019. Orchestrating Development Lifecycle of Machine Learning Based IoT Applications: A Survey. *arXiv preprint arXiv:1910.05433* (2019).
- [185] Novi Quadrianto and Viktoriia Sharmanska. 2017. Recycling privileged learning and distribution matching for fairness. In *Advances in Neural Information Processing Systems*. 677–688.
- [186] Mauro Ribeiro, Katarina Grolinger, and Miriam AM Capretz. 2015. Mlaas: Machine learning as a service. In *2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA)*. IEEE, 896–902.
- [187] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. Model-Agnostic Interpretability of Machine Learning. In *2016 ICML Workshop on Human Interpretability in Machine network Learning (WHI 2016)*. <https://doi.org/10.1145/2858036.2858529> arXiv:1606.05386
- [188] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. Model-agnostic interpretability of machine learning. *arXiv preprint arXiv:1606.05386* (2016).
- [189] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. ACM, 1135–1144.
- [190] Ronald L Rivest. 1987. Learning decision lists. *Machine learning* 2, 3 (1987), 229–246.
- [191] S Rizzi. 2009. What-if analysis. *Encyclopedia of database systems*.

- [192] Andrea Romei and Salvatore Ruggieri. 2014. A multidisciplinary survey on discrimination analysis. *The Knowledge Engineering Review* 29, 5 (2014), 582–638.
- [193] Benjamin I P Rubinstein, Peter L Bartlett, Ling Huang, and Nina Taft. 2009. Learning in a large function space: Privacy-preserving mechanisms for SVM learning. *arXiv preprint arXiv:0911.5708* (2009).
- [194] Benjamin I P Rubinstein, Blaine Nelson, Ling Huang, Anthony D Joseph, Shing-hon Lau, Satish Rao, Nina Taft, and J Doug Tygar. 2009. Antidote: understanding and defending against poisoning of anomaly detectors. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement*. ACM, 1–14.
- [195] Samm Sacks. 2018. New China Data Privacy Standard Looks More Far-Reaching than GDPR. *Center for Strategic and International Studies* 29 (2018).
- [196] Wojciech Samek. 2019. *Explainable AI: interpreting, explaining and visualizing deep learning*. Vol. 11700. Springer Nature.
- [197] Wojciech Samek, Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, and Klaus-Robert Müller. 2016. Evaluating the visualization of what a deep neural network has learned. *IEEE transactions on neural networks and learning systems* 28, 11 (2016), 2660–2673.
- [198] Sebastian Schelter, Joos-Hendrik Boese, Johannes Kirschnick, Thoralf Klein, and Stephan Seufert. 2017. Automatically tracking metadata and provenance of machine learning experiments. In *Machine Learning Systems Workshop at NIPS*. 27–29.
- [199] Vitaly Schetinin, Jonathan E Fieldsend, Derek Partridge, Timothy J Coats, Wojtek J Krzanowski, Richard M Everson, Trevor C Bailey, and Adolfo Hernandez. 2007. Confident interpretation of Bayesian decision tree ensembles for clinical applications. *IEEE Transactions on Information Technology in Biomedicine* 11, 3 (2007), 312–319.
- [200] Philipp Schmidt and Felix Biessmann. 2019. Quantifying Interpretability and Trust in Machine Learning Systems. *arXiv preprint arXiv:1901.08558* (2019).
- [201] Burr Settles. 2009. *Active learning literature survey*. Technical Report. University of Wisconsin-Madison Department of Computer Sciences.
- [202] Valerie L Shalin, Edward J Wisniewski, Keith R Levi, and Paul D Scott. 1988. A formal analysis of machine learning systems for knowledge acquisition. *International Journal of Man-Machine Studies* 29, 4 (1988), 429–446.
- [203] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K Reiter. 2016. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1528–1540.
- [204] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 3–18.
- [205] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. 2017. Learning important features through propagating activation differences. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 3145–3153.
- [206] Jatinder Singh, Jennifer Cobbe, and Chris Norval. 2018. Decision Provenance: Harnessing data flow for accountable systems. *IEEE Access* 7 (2018), 6562–6574.
- [207] Nedim Šrndić and Pavel Laskov. 2013. Detection of malicious pdf files based on hierarchical document structure. In *Proceedings of the 20th Annual Network & Distributed System Security Symposium*. Citeseer, 1–16.
- [208] Mark Stamp. 2011. *Information security: principles and practice*. Wiley Online Library.
- [209] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013).
- [210] Hui Fen Tan, Giles Hooker, and Martin T Wells. 2016. Tree space prototypes: Another look at making tree ensembles interpretable. *arXiv preprint arXiv:1611.07115* (2016).
- [211] Pang-Ning Tan. 2018. *Introduction to data mining*. Pearson Education India.
- [212] Choon H Teo, Amir Globerson, Sam T Roweis, and Alex J Smola. 2008. Convex learning with invariances. In *Advances in neural information processing systems*. 1489–1496.
- [213] Gabriele Tolomei, Fabrizio Silvestri, Andrew Haines, and Mounia Lalmas. 2017. Interpretable predictions of tree-based ensembles via actionable feature tweaking. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*. ACM, 465–474.
- [214] Ehsan Toreini, Mhairi Aitken, Kovila Coopamootoo, Karen Elliott, Carlos Gonzalez Zelaya, and Aad van Moorsel. 2020. The relationship between trust in AI and trustworthy machine learning technologies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 272–283.
- [215] Florian Tramèr, Fan Zhang, Ari Juels, Michael K Reiter, and Thomas Ristenpart. 2016. Stealing machine learning models via prediction apis. In *25th USENIX Security Symposium (USENIX Security 16)*. 601–618.
- [216] Alper Kursat Uysal and Serkan Gunal. 2014. The impact of preprocessing on text classification. *Information Processing and Management* 50, 1 (2014), 104–112. <https://doi.org/10.1016/j.ipm.2013.08.006>

- [217] Michael Veale, Max Van Kleek, and Reuben Binns. 2018. Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. In *Proceedings of the 2018 chi conference on human factors in computing systems*. ACM, 440.
- [218] Sahil Verma and Julia Rubin. 2018. Fairness definitions explained. In *2018 IEEE/ACM International Workshop on Software Fairness (FairWare)*. IEEE, 1–7.
- [219] Walther Wachenfeld and Hermann Winner. 2016. Do autonomous vehicles learn? In *Autonomous Driving*. Springer, 451–471.
- [220] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2017. Counterfactual Explanations without Opening the Black Box: Automated Decisions and the GDPR. (nov 2017). arXiv:1711.00399 <http://arxiv.org/abs/1711.00399>
- [221] Qia Wang, Wenjun Zeng, and Jun Tian. 2014. Compressive sensing based secure multiparty privacy preserving framework for collaborative data-mining and signal processing. In *2014 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 1–6.
- [222] Jess Whittlestone, Rune Nyrup, Anna Alexandrova, and Stephen Cave. 2019. The Role and Limits of Principles in AI Ethics: Towards a Focus on Tensions. In *Proceedings of the AAAI/ACM Conference on AI Ethics and Society, Honolulu, HI, USA*. 27–28.
- [223] Gregory L Wittel and Shyhtsun Felix Wu. 2004. On Attacking Statistical Spam Filters.. In *CEAS*.
- [224] Blake Woodworth, Suriya Gunasekar, Mesrob I Ohannessian, and Nathan Srebro. 2017. Learning non-discriminatory predictors. *arXiv preprint arXiv:1702.06081* (2017).
- [225] Xi Wu, Matthew Fredrikson, Somesh Jha, and Jeffrey F Naughton. 2016. A methodology for formalizing model-inversion attacks. In *2016 IEEE 29th Computer Security Foundations Symposium (CSF)*. IEEE, 355–370.
- [226] Huang Xiao, Battista Biggio, Gavin Brown, Giorgio Fumera, Claudia Eckert, and Fabio Roli. 2015. Is feature selection secure against training data poisoning?. In *International Conference on Machine Learning*. 1689–1698.
- [227] Weilin Xu, David Evans, and Yanjun Qi. 2017. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155* (2017).
- [228] W Xu, Y Qi, and D Evans. 2016. Automatically Evading Classifiers: A Case Study on PDF Malware Classifiers. NDSS.
- [229] Ke Yang, Julia Stoyanovich, Abolfazl Asudeh, Bill Howe, H V Jagadish, and Gerome Miklau. 2018. A nutritional label for rankings. In *Proceedings of the 2018 International Conference on Management of Data*. ACM, 1773–1776.
- [230] Yao Yucheng, Song Ling, and E Chi. 2017. Investigation on distributed K-means clustering algorithm of homomorphic encryption. *Computer Technology and Development* 27, 2 (2017), 81–85.
- [231] Muhammad Bilal Zafar. 2019. Discrimination in Algorithmic Decision Making: From Principles to Measures and Mechanisms. (2019).
- [232] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. 2015. Fairness constraints: Mechanisms for fair classification. *arXiv preprint arXiv:1507.05259* (2015).
- [233] Vladimiro G Zelaya, Paolo Missier, and Dennis Prangle. 2019. Parametrised Data Sampling for Fairness Optimisation. (2019).
- [234] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. 2013. Learning fair representations. In *International Conference on Machine Learning*. 325–333.
- [235] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. 2018. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, 335–340.
- [236] Fei Zhang, Patrick P K Chan, Battista Biggio, Daniel S Yeung, and Fabio Roli. 2015. Adversarial feature selection against evasion attacks. *IEEE transactions on cybernetics* 46, 3 (2015), 766–777.
- [237] Mengchen Zhao, Bo An, Wei Gao, and Teng Zhang. 2017. Efficient Label Contamination Attacks Against Black-Box Learning Models.. In *IJCAI*. 3945–3951.
- [238] Wentao Zhao, Jun Long, Jianping Yin, Zhiping Cai, and Geming Xia. 2012. Sampling attack against active learning in adversarial environment. In *International Conference on Modeling Decisions for Artificial Intelligence*. Springer, 222–233.
- [239] Yichen Zhou and Giles Hooker. 2016. Interpreting models via single tree approximation. *arXiv preprint arXiv:1610.09036* (2016).
- [240] Zhi-Hua Zhou, Yuan Jiang, and Shi-Fu Chen. 2003. Extracting symbolic rules from trained neural network ensembles. *Ai Communications* 16, 1 (2003), 3–15.
- [241] Indrè Žliobaitė. 2015. On the relation between accuracy and fairness in binary classification. *arXiv preprint arXiv:1505.05723* (2015).