

Socially Responsible AI Algorithms: Issues, Purposes, and Challenges

Lu Cheng

Computer Science and Engineering, Arizona State University

LCHENG35@ASU.EDU

Kush R. Varshney

IBM Research – Thomas J. Watson Research Center

KRVARSHN@US.IBM.COM

Huan Liu

Computer Science and Engineering, Arizona State University

HUANLIU@ASU.EDU

Abstract

In the current era, people and society have grown increasingly reliant on artificial intelligence (AI) technologies. AI has the potential to drive us towards a future in which all of humanity flourishes. It also comes with substantial risks for oppression and calamity. Discussions about whether we should *(re)trust* AI have repeatedly emerged in recent years and in many quarters, including industry, academia, healthcare, services, and so on. Technologists and AI researchers have a responsibility to develop trustworthy AI systems. They have responded with great effort to design more responsible AI algorithms. However, existing technical solutions are narrow in scope and have been primarily directed towards algorithms for scoring or classification tasks, with an emphasis on fairness and unwanted bias. To build long-lasting trust between AI and human beings, we argue that the key is to think beyond algorithmic fairness and connect major aspects of AI that potentially cause AI's indifferent behavior. In this survey, we provide a systematic framework of *Socially Responsible AI Algorithms* that aims to examine the *subjects* of AI indifference and the need for socially responsible AI algorithms, define the *objectives*, and introduce the *means* by which we may achieve these objectives. We further discuss how to leverage this framework to improve societal well-being through *protection, information, and prevention/mitigation*.

1. Introduction

Artificial intelligence (AI) has had and will continue to have a central role in countless aspects of life, livelihood, and liberty. AI is bringing forth a sea-change that is not only limited to technical domains, but is a truly sociotechnical phenomenon affecting healthcare, education, commerce, finance, and criminal justice, not to mention day-to-day life. AI offers both promise and perils. A report published by Martha Lane Fox's Doteveryone think tank (Miller C, 2019) reveals that 59% of tech workers have worked on products they felt harmful to society, and more than 25% of workers in AI who had such an experience quit their jobs as a result. This was particularly marked in relation to AI products. The rise of activism – which has been regarded as one of the current few mechanisms to keep big tech companies in check (Schwab, 2021) – against negative social impacts of big tech have brought Social Responsibility of AI into the spotlight of the media, the general public, and AI technologists and researchers (Abdalla & Abdalla, 2020). Even researchers in universities and research institutes are trying hard to rectify the mistakes made by algorithms. Stanford's COVID-19 vaccine allocation algorithm, for example, prioritizes older employees over front-line workers

(Asimov, 2020), turning much of our attention again to the transparency and fairness of AI.

Research directed towards developing fair, transparent, accountable, and ethical AI algorithms has burgeoned with a focus on decision-making algorithms such as scoring or classification to mitigate unwanted bias and achieve fairness (Jagadish, 2019). However, this narrow subset of research risks blinding us to the challenges and opportunities that are presented by the full scope of AI. To identify potential higher-order effects on safety, privacy, and society at large, it is critical to think beyond algorithmic bias, to capture all the connections among different aspects related to AI algorithms. Therefore, this survey complements prior work through a holistic understanding of the relations between AI systems and humans. In this work, we begin by introducing an inclusive definition of *Social Responsibility of AI*. Drawing on theories in business research, we then present a pyramid of Social Responsibility of AI that outlines four specific AI responsibilities in a hierarchy. This is adapted from the pyramid proposed for Corporate Social Responsibility (CSR) by Carroll et al. (1991). In the second part of the survey, we review major aspects of AI algorithms and provide a systematic framework – *Socially Responsible AI Algorithms* (SRAs) – that aims to understand the connections among these aspects. In particular, we examine the *subjects* and *causes* of socially indifferent AI algorithms¹, define the *objectives*, and introduce the *means* by which we can achieve SRAs. We further discuss how to leverage SRAs to improve daily life of human beings and address challenging societal issues through *protecting*, *informing*, and *preventing/mitigating*. We illustrate these ideas using recent studies on several emerging societal challenges. The survey concludes with open problems and challenges in SRAs.

Differences from Existing Surveys. Some recent surveys focus on specific topics such as bias and fairness (Mehrabani, Morstatter, Saxena, Lerman, & Galstyan, 2019; Caton & Haas, 2020), interpretability/explainability (Carvalho, Pereira, & Cardoso, 2019; Tjoa & Guan, 2020), and privacy-preservation (Beigi & Liu, 2020; Dwork, 2008). These surveys successfully draw great attention to the social responsibility of AI, leading to further developments in this important line of research. However, as indispensable components of socially responsible AI, these topics have been presented in their own self-contained ways. These works pave the way for looking at socially responsible AI holistically. Therefore, our survey aims to frame socially responsible AI with a more systematic view that goes beyond discussion of each independent line of research. We summarize our *contributions* as follows:

- We formally define social responsibility of AI with three specified dimensions: *principles*, *means*, and *objectives*. We then propose the pyramid of social responsibility of AI, describing its four fundamental responsibilities: *functional*, *legal*, *ethical*, and *philanthropic* responsibilities. The pyramid embraces the entire range of AI responsibilities involving efforts from various disciplines.
- We propose a systematic framework that discusses the essentials of socially responsible AI algorithms (SRAs) – including its *subjects*, *causes*, *means*, and *objectives* – and the roles of SRAs in *protecting*, *informing* users, and *preventing* them from negative impact of AI. This framework subsumes existing topics such as fairness and interpretability.

1. We define “indifferent” as the complement of responsible rather than “irresponsible”.

- We look beyond prior research in socially responsible AI and identify an extensive list of open problems and challenges, ranging from understanding *why* we need AI systems to showing the need to define new AI ethics principles and policies. We hope our discussions can spark future research on SRAs.

Intended Audience and Paper Organization. This survey is intended for AI researchers, AI technologists, researchers, and practitioners from other disciplines who would like to contribute to making AI more socially responsible with their expertise. The rest of the survey is organized as follows: Section 2 introduces the definition and the pyramid of social responsibility of AI, and compares definitions of similar concepts. Section 3 discusses the framework of socially responsible algorithms and its essentials, followed by Section 4 that illustrates the roles of SRAs using several emerging societal issues as examples. Section 5 details the open problems and challenges that socially responsible AI currently confronts. The last section concludes the survey.

2. Social Responsibility of AI

Social Responsibility of AI includes efforts devoted to addressing both technical and societal issues. While similar concepts (e.g., “Ethical AI”) repeatedly appear in the news, magazines, and scientific articles, “Social Responsibility of AI” has yet to be properly defined. In this section, we first attempt to provide an inclusive definition and then propose the *Pyramid* of Social Responsibility of AI to outline the various responsibilities of AI in a hierarchy: functional responsibilities, legal responsibilities, ethical responsibilities, and philanthropic responsibilities. At last, we compare “Socially Responsible AI” with similar concepts.

2.1 What is Social Responsibility of AI?

Definition 1 (Social Responsibility of AI) *Social Responsibility of AI refers to a human value-driven process where values such as Fairness, Transparency, Accountability, Reliability and Safety, Privacy and Security, and Inclusiveness are the principles; designing Socially Responsible AI Algorithms is the means; and addressing the social expectations of generating shared value – enhancing both AI’s ability and benefits to society – is the main objective.*

Here, we define three dimensions of Social Responsibility of AI: the *principles* lay the foundations for ethical AI systems; the *means* to reach the overarching goal of Social Responsibility of AI is to develop Socially Responsible AI Algorithms; and the *objective* of Social Responsibility of AI is to improve both AI’s capability and humanity with the second being the proactive goal.

2.2 The Pyramid of Social Responsibility of AI

Social Responsibility of AI should be framed in such a way that the entire range of AI responsibilities are embraced. Adapting Carroll’s Pyramid of CSR (Carroll et al., 1991) in the AI context, we suggest four kinds of social responsibilities that constitute the Social Responsibility of AI: *functional*, *legal*, *ethical*, and *philanthropic* responsibilities, as shown in Figure 1. By modularizing AI responsibilities, we hope to help AI technologists and




Figure 1: The pyramid of Social Responsibility of AI, adapted from the Pyramid of CSR by Carroll et al. (1991).

researchers to reconcile these obligations and simultaneously fulfill all the components in the pyramid. All of these responsibilities have always existed, but functional responsibilities have been the main consideration until recently. Each type of responsibility requires close consideration.

The pyramid portrays the four components of Social Responsibility of AI, beginning with the basic building block notion that the functional competence of AI undergirds all else. *Functional responsibilities* require AI systems to perform in a manner consistent with profits maximization, operating efficiency, and other key performance indicators. Meanwhile, AI is expected to obey the law, which codifies the acceptable and unacceptable behaviors in our society. That is, *legal responsibilities* require AI systems to perform in a manner consistent with expectations of government and law. All AI systems should at least meet the minimal legal requirements. At its most fundamental level, *ethical responsibilities* are the obligation to do what is right, just, and fair, and to prevent or mitigate negative impact on stakeholders (e.g., users, the environment). To fulfill its ethical responsibilities, AI systems need to perform in a manner consistent with societal expectations and ethical norms, which cannot be compromised in order to achieve AI’s functional responsibilities. Finally, in *philanthropic responsibilities*, AI systems are expected to be *good AI citizens* and to contribute to tackling societal challenges such as cancer and climate change. Particularly, it is important for AI systems to perform in a manner consistent with the philanthropic and charitable expectations of society to enhance people’s quality of life. The distinguishing feature between ethical and philanthropic responsibilities is that the latter are not expected in an ethical sense. For example, while communities desire AI systems to be applied to

Concepts	Definitions
Robust AI	AI systems with the ability “to cope with errors during execution and cope with erroneous input” (Wikipedia, 2021a).
Ethical AI	AI systems that do what is right, fair, and just. Prevent harm.
Trustworthy AI	AI systems that achieve their full potential if <i>trust</i> can be established in the development, deployment, and use (Thiebes, Lins, & Sunyaev, 2020).
Fair AI	AI systems absent from “any prejudice or favoritism toward an individual or a group based on their inherent or acquired characteristics” (Mehrabi et al., 2019).
Safe AI	AI systems deployed in ways that do not harm humanity (Feige, 2019).
Dependable AI	AI systems that focus on reliability, verifiability, explainability, and security (Singh, Vatsa, & Ratha, 2021).
Human-centered AI	AI systems that are “continuously improving because of human input while providing an effective experience between human and robot” ² .

Table 1: Definitions of concepts similar to Socially Responsible AI.

humanitarian projects or purposes, they do not regard the AI systems as unethical if they do not provide such services. We explore the nature of Social Responsibility of AI by focusing on its components to help AI technologists to reconcile these obligations. Though these four components are depicted as separate concepts, they are not mutually exclusive. It is necessary for AI technologists and researchers to recognize that these obligations are in a constant but dynamic tension with one another.

2.3 Comparisons of Similar Concepts

Based on Definition 1 and the pyramid of socially responsibility of AI, we compare Socially Responsible AI with other similar concepts, as illustrated in Table 1. The results show that Socially Responsible AI holds a systematic view that subsumes existing concepts and further considers the fundamental responsibilities of AI systems – to be functional and legal, as well as their philanthropic responsibilities – to be able to improve life quality of well beings and address challenging societal issues. In the rest of this survey, we focus our discussions on the *ethical* (Section 3, essentials of SRAs) and *philanthropic* (Section 4, roles of SRAs) responsibilities of AI given that both the functional and legal responsibilities are the usual focuses in AI research and development. An overview of SRAs research is illustrated in Figure 2, which we will refer back to throughout the remainder of the survey. Importantly, in our view, the essentials of SRAs work toward ethical responsibilities, and their roles in society encompass both ethical and philanthropic responsibilities.

2. <https://www.cognizant.com/glossary/human-centered-ai>

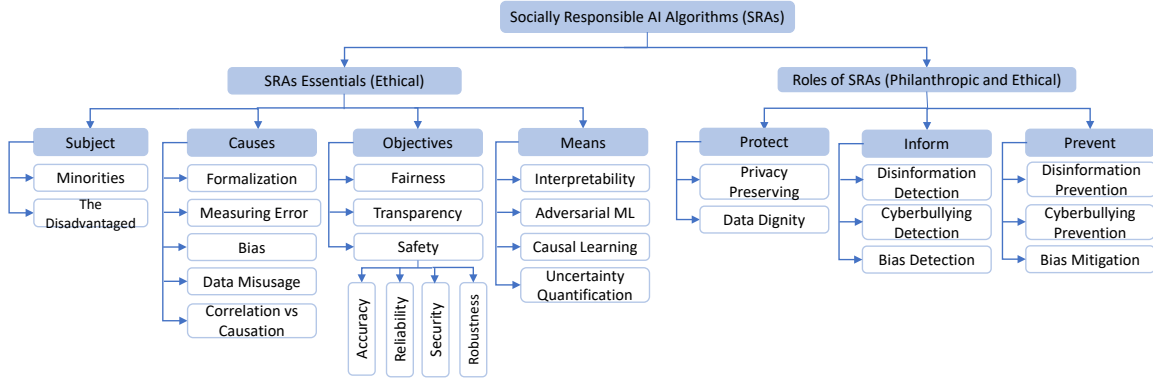


Figure 2: An overview of SRAs Research.

3. Socially Responsible AI Algorithms (SRAs)

The role of AI technologists and researchers carries a number of responsibilities. The most obvious is developing accurate, reliable, and trustworthy algorithms that can be depended on by their users. Yet, this has never been a trivial task. For example, due to the various types of human biases, e.g., confirmation bias, gender bias, and anchoring bias, AI technologists and researchers often inadvertently inject these same kinds of bias into the developed algorithms, especially when using machine learning techniques. For example, supervised machine learning is a common technique for learning and validating algorithms through manually annotated data, loss functions, and related evaluation metrics. Numerous uncertainties – e.g., imbalanced data, ill-defined criteria for data annotation, over-simplified loss functions, and unexplainable results – potentially lurk in this “beautiful” pipeline and will eventually lead to negative outcomes such as biases and discrimination. With the growing reliance on AI in almost any field in our society, we must bring upfront the vital question about *how to develop Socially Responsible AI Algorithms*. While conclusive answers are yet to be found, we attempt to provide a systematic framework of SRAs (illustrated in Figure 3) to discuss the components of AI’s ethical responsibilities, the roles of SRAs in terms of AI’s philanthropic and ethical responsibilities, and the feedback from users routed back as inputs to SRAs. We hope to broaden future discussions on this subject. In this regard, we define SRAs as follows:

Definition 2 (Socially Responsible AI Algorithms) *Socially Responsible AI Algorithms are the intelligent algorithms that prioritize the needs of all stakeholders as the highest priority, especially the minoritized and disadvantaged users, in order to make just and trustworthy decisions. These obligations include protecting and informing users; preventing and mitigating negative impact; and maximizing the long-term beneficial impact. Socially Responsible AI Algorithms constantly receive feedback from users to continually accomplish the expected social values.*

In this definition, we highlight that the functional (e.g., maximizing profits) and societal (e.g., transparency) objectives are integral parts of AI algorithms. SRAs aim to be socially responsible while still meeting and exceeding business objectives.

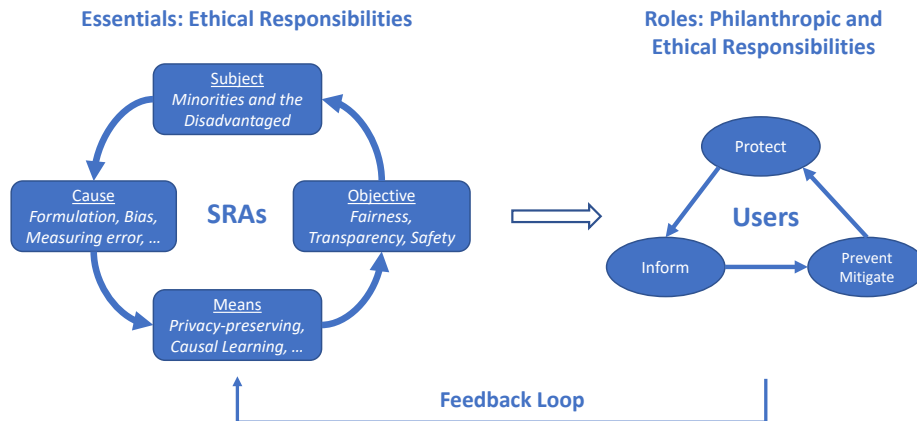


Figure 3: The framework of Socially Responsible AI Algorithms (SRAs). It consists of the essentials (i.e., the internal mechanisms) of SRAs (left), their roles (right), and feedback received from end users for helping SRAs gradually achieve the expected social values (bottom). The essentials of SRAs center on the ethical responsibilities of AI and the roles of SRAs require philanthropic responsibilities and ethical responsibilities.

3.1 Subjects of Socially Indifferent AI Algorithms

Every human being can be a potential victim of socially indifferent AI algorithms. Mirroring society, the ones who suffer the most, both in frequency and severity, are minorities and disadvantaged groups such as black, indigenous and people of color (BIPOC), and females. For example, Google mislabeled an image of two black people as “gorillas” (Guynn, 2015) and more frequently showed ads of high-paying jobs to males than females (Carpenter, 2015). Similar gender bias was also found in Facebook algorithms behind the job ads (Horwitz, 2021). In domains with high-stakes decisions, e.g., financial services, healthcare, and criminal justice, it is not uncommon to identify instances where socially indifferent AI algorithms favor privileged groups. For example, the algorithm used in Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) was found almost twice as likely to mislabel a black defendant as a future risk than a white defendant (Angwin, Larson, Mattu, & Kirchner, 2016). Identifying the subjects of socially indifferent AI algorithms depends on the context. In another study, the journalistic organization ProPublica³ investigated algorithms that determine online prices for Princeton Review’s tutoring classes. The results showed that people who lived in higher income areas were charged twice as much as the general public and than people living in a zip code with high population density. Asians were 1.8 times more likely to pay higher price, regardless of their income (Angwin & Larson, 2015). Analogously, these AI algorithms might put poor people who cannot afford internet service at disadvantage because they simply have never seen such data samples in the training process.

3. <https://www.propublica.org/>

When it comes to purpose-driven collection and use of data, each individual can be the subject of socially indifferent AI algorithms. Users’ personal data are frequently collected and used without their consent. Such data includes granular details such as contact information, online browsing and session record, social media consumption, location and so on. While most of us are aware of our data being used, few have controls to where and how the data is used, and by whom. The misuse of data and lack of knowledge causes users to become the victims of privacy-leakage and distrust.

3.2 Causes of Socially Indifferent AI Algorithms

There are many potential factors that can cause AI algorithms to be socially indifferent. Here, we list several causes that have been frequently discussed in literature (Mehrabi et al., 2019; Getoor, 2019). They are formalization, measuring errors, bias, privacy, and correlation versus causation.

3.2.1 FORMALIZATION

AI algorithms encompass data formalization, label formalization, formalization of loss function and evaluation metrics. We unconsciously make some frame of reference commitment to each of these formalizations. Firstly, the social and historical context are often left out when transforming raw data into numerical feature vectors. Therefore, AI algorithms are trained on pre-processed data with important contextual information missing. Secondly, data annotation can be problematic for a number of reasons. For example, what are the criteria? Who defines the criteria? Who are the annotators? How can it be ensured that they all follow the criteria? What we have for model training are only proxies of the true labels (Getoor, 2019). Ill-formulated loss functions can also result in socially indifferent AI algorithms. Many loss functions in the tasks are over-simplified to solely focus on maximizing profits and minimizing losses. The concerns of unethical optimization are recently discussed by Beale, Battey, Davison, and MacKay (2019). Unknown to AI systems, certain strategies in the optimization space that are considered as unethical by stakeholder may be selected to satisfy the simplified task requirements. Lastly, the use of inappropriate benchmarks for evaluation may push algorithms away from the overarching goal of the task and fuel injustice.

3.2.2 MEASURING ERRORS

Another cause of socially indifferent AI algorithms is the errors when measuring algorithm performance. When reporting results, researchers typically proclaim the proposed algorithms can achieve certain accuracy or F1 scores. However, this is based on assumptions that the training and test samples are representative of the target population and their distributions are similar enough. Yet, how often does the assumption hold in practice? As illustrated in Figure 4, with non-representative samples, the learned model can achieve zero training error and perform well on the testing data at the initial stage. However, with more data being tested later, the model performance deteriorates because the learned model does not represent the true model.




Figure 4: An example of measuring errors. The green line denotes the learned model and the blue one is the true model. ‘+’ and ‘-’ represent training data belonging to different classes; ‘X’ represents testing data. Image taken from Getoor’s slides for 2019 IEEE Big Data keynote (Getoor, 2019) with permission.

3.2.3 BIAS

Bias is one of the most discussed topics regarding responsible AI. We here focus on the data bias, automation bias, and algorithmic bias (Getoor, 2019).

Data Bias. Data, especially big data, is often heterogeneous – data with high variability of types and formats, e.g., text, image, and video. The availability of multiple data sources brings unprecedented opportunities as well as unequivocally presented challenges (Li, Cheng, Wang, Morstatter, Trevino, Tang, & Liu, 2017). For instance, high-dimensional data such as text is infamous for the danger of overfitting and the curse of dimensionality. Additionally, it is rather challenging to find subset of features that are predictive but uncorrelated. The required number samples for generalization also grows proportionally with feature dimension. One example is how the U.S. National Security Agency tried to use AI algorithms to identify potential terrorists. The Skynet project collected cellular network traffic in Pakistan and extracted 80 features for each cell phone user with only 7 known terrorists (Gershgorin, 2019). The algorithm ended up identifying an Al Jazeera reporter covering Al Qaeda as a potential terrorist. Data heterogeneity is also against the well known *i.i.d.* assumption in most learning algorithms (Li et al., 2017). Therefore, training these algorithms on heterogeneous data can result in undesired results. Imbalanced subgroups is another source of data bias. As illustrated in (Mehrabi et al., 2019), regression analysis based on the subgroups with balanced fitness level suggests positive correlation between BMI and daily pasta calorie intake whereas that based on less balanced data shows almost no relationship.

Automation Bias. This type of bias refers to our preference to results suggested by automated decision-making systems while ignoring the contradictory information. Too much reliance on the automated systems without sparing additional thoughts in making final decisions, we might end up abdicating decision responsibility to AI algorithms.

Algorithmic Bias. Algorithmic bias regards biases added purely by the algorithm itself (Baeza-Yates, 2018). Some algorithms are inadvertently taught prejudices and unethical biases by societal patterns hidden in the data. Typically, models fit better to features that frequently appear in the data. For example, an automatic AI recruiting tool will




Figure 5: Confounders are common reasons for spurious correlation between two variables that are not causally connected.

learn to make decisions for a given applicant of a software engineer position using observed patterns such as “experience”, “programming skills”, “degree”, and “past projects”. For a position where gender disparity is large, the algorithms mistakenly interpret this collective imbalance as a useful pattern in the data rather than undesirable noise that should have been discarded. Algorithmic bias is systematic and repeatable error in an AI system that creates discriminated outcome, e.g., privileging wealthy users over others. It can amplify, operationalize, and even legitimize institutional bias (Getoor, 2019).

3.2.4 DATA MISUSE

Data is the fuel and new currency that has empowered tremendous progress in AI research. Search engines have to rely on data to craft precisely personalized recommendation that improves the online experience of consumers, including online shopping, book recommendation, entertainment, and so on. However, users’ data are frequently misused without the consent and awareness of users. One example is the Facebook-Cambridge Analytical scandal (Wikipedia, 2021b) where millions of Facebook users’ personal data was collected by Cambridge Analytica (Wikipedia, 2021c), without their consent. In a more recent study (Cabañas, Cuevas, Arrate, & Cuevas, 2020), researchers show that Facebook allows advertisers to exploit its users’ sensitive information for tailored ad campaigns. To make things worse, users often have no clue about where, how, and why their data is being used, and by whom. The lack of knowledge and choice over their data causes users to undervalue their personal data, and further creates issues such as privacy and distrust.

3.2.5 CORRELATION VS CAUSATION

AI Algorithms can become socially indifferent when correlation is misinterpreted as causation. For example, in the diagram in Figure 5, we observe a strong correlation between the electric bill of an ice cream shop and ice cream sales. Apparently, high electric bill cannot *cause* the ice cream sales to increase. Rather, weather is the common cause of electric bill and the sale, i.e., high temperature causes high electric bill and the increased ice cream sales. Weather – the *confounder* – creates a spurious correlation between electric bill and ice cream sales. Causality is a generic relationship between the cause and the outcome (Guo, Cheng, Li, Hahn, & Liu, 2020). While correlation helps with prediction, causation is important for decision making. One typical example is Simpson’s Paradox (Blyth, 1972). It describes a phenomenon where a trend or association observed in subgroups maybe opposite to that observed when these subgroups are aggregated. For instance, in the study of analyzing the




Figure 6: The objectives of Socially Responsible AI Algorithms.

sex bias in graduation admissions at UC Berkeley (Bickel, Hammel, & O’Connell, 1975), the admission rate was found higher in male applicants when using the entire data. However, when the admission data were separated and analyzed over the departments, female candidates had equal or even higher admission rate over male candidates.

3.3 Objectives of Socially Responsible AI Algorithms

Essentially, the goal is to (re)build *trust* in AI. By definition, trust is the “firm belief in the reliability, truth or ability of someone or something”⁴. It is a high-level concept that needs to be specified by more concrete objectives. We here discuss the SRAs objectives that have been discussed comparatively frequently in literature. They are fairness, transparency, and safety as illustrated in Figure 6.

3.3.1 FAIRNESS

Fairness in AI has gained substantial attentions in both research and industry since 2010. For decades, researchers found it rather challenging to present a unified definition of fairness in part because fairness is a societal and ethical concept. This concept is mostly subjective, changes over social context, and evolves over time, making fairness a rather challenging goal to achieve in practice. Because SRAs is a decision-making process commensurate with social values, we here adopt a fairness definition in the context of decision-making:

Definition 3 (Fairness) *“Fairness is the absence of any prejudice or favoritism toward an individual or a group based on their inherent or acquired characteristics” (Mehrabi et al., 2019).*

Note that even an ideally “fair” AI system defined in a specific context might still lead to biased decisions as the entire decision making process involves numerous elements such as policy makers and environment. While the concept of fairness is difficult to pin down, unfairness/bias/discrimination might be easier to identify. There are six types of discrimination (Mehrabi et al., 2019). Direct discrimination results from protected attributes of individuals while indirect discrimination from seemingly neural and non-protected attributes. Systemic

4. Definition from Oxford Languages.

discrimination relates to policies that may show discrimination against subgroups of population. Statistical discrimination occurs when decision makers use average statistics to represent individuals. Depending whether the differences amongst different groups can be justified or not, we further have explainable and unexplainable discrimination.

3.3.2 TRANSPARENCY

Transparency is another important but quite ambiguous concept. This is partly because AI alone can be defined in more than 70 ways (Legg, Hutter, et al., 2007). When we seek a transparent algorithm, we are asking for an *understandable explanation* of how it works (Yeo, 2020): What does the training set look like? Who collected the data? What is the algorithm doing? There are mainly three types of transparency with regard to human interpretability of AI algorithms (Weller, 2017): For a *developer*, the goal of transparency is to understand how the algorithm works and get a sense of why; for a *deployer* who owns and releases the algorithm to the public, the goal of transparency is to make the consumers to feel safe and comfortable to use the system; and what transparency means to a *user* is understanding what the AI system is doing and why. We may further differentiate global transparency from local transparency, the former aims to explain the entire system whereas the latter explains a decision within a particular context.

Yet, at the same time, disclosures about AI can pose potential risks: explanations can be hacked and releasing additional information may make AI more vulnerable to attacks. It is becoming clear that transparency is often beneficial but not universally good (Weller, 2017). The AI “transparency paradox” encourages different parties of AI systems to think more carefully about how to balance the transparency and the risks it poses. We can also see related discussions in recent work such as (Slack, Hilgard, Jia, Singh, & Lakkaraju, 2020). The paper studied how the widely recognized interpretable algorithms LIME (Ribeiro, Singh, & Guestrin, 2016) and SHAP (Lundberg & Lee, 2017) could be hacked. As the authors illustrated, explanations can be purposefully manipulated, leading to a loss of trust not only in the model but also in its explanations (Slack et al., 2020). Consequently, while working towards the goal of transparency, we must also recognize that privacy and security are the indispensable conditions we need to satisfy.

3.3.3 SAFETY

Because AI systems operate in a world with much uncertainty, volatility, and flux, another objective of SRAs is to be safe, accurate, and reliable (Varshney & Alemzadeh, 2017). There are four operational objectives relevant to Safety: *accuracy*, *reliability*, *security*, and *robustness* (Leslie, 2019). In machine learning, *accuracy* is typically measured by error rate or the fraction of instances for which the algorithm produces an incorrect output. As a standard performance metric, accuracy should be the fundamental component to establishing the approach to safe AI. It is necessary to specify a proper performance measure for evaluating any AI systems. For instance, when data for classification tasks is extremely imbalanced, precision, recall, and F1 scores are more appropriate than accuracy. The objective of *reliability* is to ensure that AI systems behave as we anticipate. It is a measure of consistency and is important to establish confidence in the safety of AI systems. *Security* encompasses the protection of information integrity, confidentiality, and continuous functionality to its

users. Under harsh conditions (e.g., adversarial attack, perturbations, and implementation error), AI systems are expected to functions reliably and accurately, i.e., *Robustness*.

3.4 Means Towards Socially Responsible AI Algorithms

In this section, we review four primary machine learning techniques and statistical methods for achieving the goals of SRAs – interpretability and explainability, adversarial machine learning, causal learning, and uncertainty quantification. Existing surveys have conducted comprehensive reviews on each of these techniques: e.g., interpretability (Carvalho et al., 2019; Tjoa & Guan, 2020), causal learning (Guo et al., 2020; Yao, Chu, Li, Li, Gao, & Zhang, 2020), adversarial machine learning (Chakraborty, Alam, Dey, Chattopadhyay, & Mukhopadhyay, 2018; Akhtar & Dasgupta, 2019), and uncertainty quantification (Kabir, Khosravi, Hosen, & Nahavandi, 2018). We thereby focus on the basics and the most frequently discussed methods in each means.

3.4.1 INTERPRETABILITY AND EXPLAINABILITY

Interpretability and explainability are the keys to increasing transparency of AI algorithms. This is extremely important when we leverage these algorithms for high-stakes prediction applications, which deeply impact people’s lives (Rudin, 2019). Existing work in machine learning interpretability can be categorized according to different criteria. Depending on when the interpretability methods are applicable (before, during, or after building the machine learning model), we have pre-model (before), in-model (during), and post-model (after) interpretability. Pre-model techniques are only applicable to the data itself. It requires an in-depth understanding of the data before building the model, e.g., sparsity and dimensionality. Therefore, it is closely related to data interpretability (Carvalho et al., 2019), in which classic descriptive statistics and data visualization methods are often used, including Principal Component Analysis (Wold, Esbensen, & Geladi, 1987) and t-SNE (Maaten & Hinton, 2008), and clustering methods such as k -means (Hartigan & Wong, 1979). In-model interpretability asks for intrinsically interpretable AI algorithms (e.g., Yang, Yang, Dyer, He, Smola, & Hovy, 2016), we can also refer to it as intrinsic interpretability. It can be achieved through imposition of constraints on the model such as causality, sparsity, or physical conditions from domain knowledge (Rudin, 2019). In-model interpretability answers question *how the model works* (Lipton, 2018). Decision trees, rule-based models, linear regression, attention network, and disentangled representation learning are in-model interpretability techniques. Post-model interpretability, or post-hoc interpretability (e.g., Mordvintsev, Olah, & Tyka, 2015; Ribeiro et al., 2016), is applied after model training. It answers the question *what else can the model tell us* (Lipton, 2018). Post-model interpretability include local explanations (Ribeiro et al., 2016), saliency maps (Simonyan, Vedaldi, & Zisserman, 2013), example-based explanations (Kim, Khanna, & Koyejo, 2016), influence functions (Koh & Liang, 2017), feature visualization (Erhan, Bengio, Courville, & Vincent, 2009), and explaining by base interpretable models (Craven & Shavlik, 1996).

Another criterion to group current interpretability techniques is model-specific vs model-agnostic. Model-specific interpretation is based on internals of a specific model (Molnar, 2020). To illustrate, the coefficients of a linear regression model belong to model-specific interpretation. Model-agnostic methods do not have access to the model inner workings,

rather, they are applied to any machine learning model after it has been trained. Essentially, the goal of interpretability is to help the user understand the decisions made by the machine learning models through the tool *explanation*. There are pragmatic and non-pragmatic theories of explanation. The former indicates that explanation should be a *good* answer that can be easily understood by the audience. The non-pragmatic theory emphasizes the *correctness* of the answer to the why-question. Both need to have the following properties (Robnik-Šikonja & Bohanec, 2018): expressive power, translucency, portability, and algorithmic complexity.

3.4.2 ADVERSARIAL MACHINE LEARNING

Machine learning models, especially deep learning models, are vulnerable to crafted adversarial examples, which are imperceptible to human eyes but can easily fool deep neural networks (NN) in the testing/deploying stage (Yuan, He, Zhu, & Li, 2019). Adversarial examples have posed great concerns in the security and integrity of various applications. Adversarial machine learning, therefore, closely relates to the robustness of SRAs.

The security of any machine learning model is measured with regard to the adversarial goals and capabilities (Chakraborty et al., 2018). Identifying the threat surface (Papernot, McDaniel, Sinha, & Wellman, 2016a) of an AI system built on machine learning models is critical to understand where and how an adversary may subvert the system under attack. For example, the attack surface in a standard automated vehicle system can be defined with regard to the data processing pipeline. Typically, there are three types of attacks the attack surface can identify: evasion attack – the adversary attempts to evade the system by manipulating malicious samples during testing phase, poisoning attack – the adversary attempts to poison the training data by injecting carefully designed samples into the learning process, and exploratory attack – it tries to collect as much information as possible about the learning algorithm of the underlying system and pattern in training data. Depending on the amount of information available to an adversary about the system, we can define different types of adversarial capabilities. In the training phase (i.e., training phase capabilities), there are three broad attack strategies: (1) data injection. The adversary can only augment new data to the training set; (2) data modification. The adversary has full access to the training data; and (3) logic corruption. The adversary can modify the learning algorithm. In the testing phase (i.e., testing phase capabilities), adversarial attacks focus on producing incorrect outputs. For white-box attack, an adversary has full knowledge about the model used for prediction: algorithm used in training, training data distribution, and the parameters of the fully trained model. The other type of attack is black-box attack, which, on the contrary, assumes no knowledge about the model and only uses historical information or information about the settings. The primary goal of black-box attack is to train a local model with the data distribution, i.e., non-adaptive attack, and with carefully selected dataset by querying the target model, i.e., adaptive attack.

Exploratory attacks do not have access to the training data but aim to learn the current state by probing the learner. Commonly used techniques include model inversion attack (Fredrikson, Lantz, Jha, Lin, Page, & Ristenpart, 2014; Fredrikson, Jha, & Ristenpart, 2015), model extraction using APIs (Tramèr, Zhang, Juels, Reiter, & Ristenpart, 2016), and inference attack (Ateniese, Mancini, Spognardi, Villani, Vitali, & Felici, 2015; Shokri,

Stronati, Song, & Shmatikov, 2017). The popular attacks are evasion attacks where malicious inputs are craftily manipulated so as to fool the model to make false predictions. Poisoning attacks, however, modify the input during the training phase to obtain the desired results. Some of the well-known techniques are generative adversarial network (GAN) (Goodfellow, Pouget-Abadie, Mirza, Xu, Warde-Farley, Ozair, Courville, & Bengio, 2014), adversarial examples generation (including training phase modification, e.g., Barreno, Nelson, Sears, Joseph, & Tygar, 2006, and testing phase modification, e.g., Papernot, McDaniel, Wu, Jha, & Swami, 2016b), GAN-based attack in collaborative deep learning (Hitaj, Ate-niese, & Perez-Cruz, 2017), and adversarial classification (Dalvi, Domingos, Sanghai, & Verma, 2004).

3.4.3 CAUSAL LEARNING

Causal inference and reasoning is a critical ingredient for AI to achieve human-level intelligence, an overarching goal of Socially Responsible AI. The momentum of integrating causality into responsible AI is growing, as witnessed by a number of works (e.g., Kusner, Loftus, Russell, & Silva, 2017; Xu, Wu, Yuan, Zhang, & Wu, 2019; Holzinger, Langs, Denk, Zatloukal, & Müller, 2019) studying SRAs through causal learning methods.

Basics of Causal Learning. The two fundamental frameworks in causal learning are *structural causal models* (Pearl, 2009) and *potential outcome* (Rubin, 1974). Structural causal models rely on the causal graph, which is a special class of Bayesian network with edges denoting causal relationships. A more structured format is referred to as structural equations. One of the fundamental notions in structural causal models is the *do*-calculus (Pearl, 2009), an operation for *intervention*. The difficulty to conduct causal study is the difference between the observational and interventional distribution, the latter describes what the distribution of outcome Y is if we were to set covariates $X = x$. Potential outcome framework interprets causality as given the treatment and outcome, we can only observe one potential outcome. The counterfactuals – potential outcome that would have been observed if the individual had received a different treatment – however, can never be observed in reality. These two frameworks are the foundations of causal effect estimation (estimating effect of a treatment) and causal discovery (learning causal relations amongst different variables).

Many important concepts in causal inference have been adapted to AI such as intervention and counterfactual reasoning. Here, we introduce the causal concept most frequently used in SRAs – propensity score, defined as “conditional probability of assignment to a particular treatment given a vector of observed covariates” (Rosenbaum & Rubin, 1983). A popular propensity-based approach is Inverse Probability of Treatment Weighting (Hirano, Imbens, & Ridder, 2003). To synthesize a randomized control trial (Rubin, 1974), it uses covariate balancing to weigh instances based on their propensity scores and the probability of an instance to receive the treatment. Let t_i and x_i be the treatment assignment and covariate of instance i , the weight w_i is typically computed by the following formula:

$$w_i = \frac{t_i}{P(t_i|x_i)} + \frac{1 - t_i}{1 - P(t_i|x_i)}, \quad (1)$$

where $P(t_i|x_i)$ quantifies the propensity score. The weighted average of the observed outcomes for the treatment and control groups are defined as

$$\hat{\tau} = \frac{1}{n_1} \sum_{i:t_i=1} w_i y_i - \frac{1}{n_0} \sum_{i:t_i=0} w_i y_i, \quad (2)$$

where n_1 and n_0 denote the sizes of the treated and controlled groups.

Causal Learning for SRAs. Firstly, it is becoming increasingly popular to use causal models to solve fairness-related issues. For example, the subject of causality and its importance to address fairness issue was discussed in (Loftus, Russell, Kusner, & Silva, 2018). Causal models can also be used to discover and eliminate discrimination to make decisions that are irrespective of sensitive attributes, on individual-, group-, and system-level, see, e.g., (Zhang, Wu, & Wu, 2016; Zhang & Wu, 2017; Nabi & Shpitser, 2018). Secondly, bias alleviation is another field where causal learning methods are frequently discussed and affect many machine learning applications at large. The emerging research on debiasing recommender system (Wang, Golbandi, Bendersky, Metzler, & Najork, 2018b; Wang, Liang, Charlin, & Blei, 2018c; Joachims, Swaminathan, & Schnabel, 2017) can serve as one example. Due to the biased nature of user behavior data, recommender systems inevitably involve with various discrimination-related issues: recommending less career coaching services and high-paying jobs to women (Lambrecht & Tucker, 2019; Datta, Tschantz, & Datta, 2015), recommending more male-authored books (Ekstrand, Tian, Kazi, Mehrpouyan, & Kluver, 2018), and minorities are less likely to become social influencers (Karimi, Génois, Wagner, Singer, & Strohmaier, 2018; Stoica, Riederer, & Chaintreau, 2018). Gender and ethnic biases were even found in a broader context, e.g., word embeddings trained on 100 years of text data (Garg, Schiebinger, Jurafsky, & Zou, 2018). Causal approaches such as (Yang & Feng, 2020) aim to mitigate such bias in word embedding relations.

Thirdly, causal learning methods also have had discernible achievements in transparency, especially the interpretability of black-box algorithms. Causality is particularly desired since these algorithms only capture correlations not real causes (Moraffah, Karami, Guo, Raglin, & Liu, 2020). Further, it has been suggested that counterfactual explanations are the highest level of interpretability (Pearl, 2018). For model-based interpretations, causal interpretability aims to explain the causal effect of a model component on the final decision (Chattopadhyay, Manupriya, Sarkar, & Balasubramanian, 2019; Parafita & Vitrià, 2019; Narendra, Sankaran, Vijaykeerthy, & Mani, 2018). One example to differentiate it from traditional interpretability is only causal interpretability is able to answer question such as “What is the effect of the n -th filter of the m -th layer of a neural network on the prediction of the model?”. Counterfactual explanations is a type of example-based explanations, in which we look for data instances that can explain the underlying data distributions. Counterfactual explanations are human friendly, however, it is possible to have different true versions of explanations for the predicted results, i.e., the Rashomon effect (Molnar, 2020). Studies such as (Wachter, Mittelstadt, & Russell, 2017; Grath, Costabello, Van, Sweeney, Kamiab, Shen, & Lecue, 2018; Liu, Kailkhura, Loveland, & Han, 2019) are proposed to address this issue. For detailed discussion on causal interpretability, please refer to (Moraffah et al., 2020). Lastly, causal learning is inherently related to the robustness or adaptability of AI systems, which have been noted to lack the capability of reacting to new circumstances they are not trained for. Causal relationship, however, is expected to be invariant

and robust across environments (Pearl, 2019, 2009). This complements intensive earlier efforts toward “transfer learning”, “domain adaptation”, and “lifelong learning” (Chen & Liu, 2018). Some current work seeking to extrapolate the relationship between AI robustness and causality includes the independent causal mechanism principle (Peters, Janzing, & Schölkopf, 2017; Schölkopf, 2019), invariant prediction (Arjovsky, Bottou, Gulrajani, & Lopez-Paz, 2019), and disentangled causal mechanism (Suter, Miladinovic, Schölkopf, & Bauer, 2019; Bengio, Deleu, Rahaman, Ke, Lachapelle, Bilaniuk, Goyal, & Pal, 2019).

3.4.4 UNCERTAINTY QUANTIFICATION

AI research continues to develop new state-of-the-art algorithms with superior performance and large-scaled datasets with high quality. Even using the best models and training data, it is still infeasible for AI systems to cover all the potential situations when deployed into real-world applications. As a matter of fact, AI systems always encounter new samples that are different from those used for training. The core question is how to leverage the strengths of these uncertainties. Recent research, e.g., (Bhatt, Antorán, Zhang, Liao, Sattigeri, Fogliato, Melancon, Krishnan, Stanley, Tickoo, et al., 2020), has advocated to measure, communicate, and use uncertainty as a form of transparency. There are also tools such as IBM’s Uncertainty Quantification 360⁵ to provide AI practitioners access to related resources as common practices for AI transparency. Consequently, uncertainty quantification plays a crucial role in the optimization and decision-making process in SRAs. There are typically two kinds of uncertainties in risk analysis processes: first, the aleatory uncertainty describes the inherent randomness of systems. For example, an AI system can present different results even with the same set of inputs. The uncertainty arises from underlying random variations within the data. Second, the epistemic uncertainty represents the effect of an unknown phenomenon or an internal parameter. The primary reason leading to this type of uncertainty is the lack of observed data. As the variation among the data in aleatory uncertainty is often observable, we can well quantify the uncertainty and assess the risks. Quantification of epistemic uncertainty is more challenging because AI systems are forced to extrapolate over unseen situations (Ståhl, Falkman, Karlsson, & Mathiason, 2020). In the literature of uncertainty quantification, one of the most widely recognized techniques are prediction intervals (PI). For neural-network-based models, PI can be categorized into multi-step PI construction methods (e.g., Bayesian method) and direct PI construction methods (e.g., lower upper bound estimation). Here, we briefly discuss several methods in each category. Please refer to the survey (Kabir et al., 2018) for more details.

Multi-Step Prediction Intervals Construction Methods. Delta method, Bayesian method, Mean-Variance Estimation method, and Bootstrap method are the four conventional multi-step methods reported in literature. Delta method constructs PIs through nonlinear regression using Tylor series expansion. Particularly, we linearize neural network models through optimization by minimizing the error-based loss function, sum square error. Under the assumption that uncertainty is from normal and homogeneous distribution, we then employ standard asymptotic theory to construct PIs. Delta method has been used in numerous case studies, e.g., (Lu & Viljanen, 2009; Ho, Xie, Tang, Xu, & Goh, 2001). Bayesian learning provides a natural framework for constructing PIs (Kasiviswanathan &

5. <http://uq360.mybluemix.net/overview>




Figure 7: Illustration of what Socially Responsible AI Algorithms (SRAs) can do. It requires philanthropic responsibilities and ethical responsibilities.

Sudheer, 2016; Ungar, De Veaux, & Rosengarten, 1996) as it optimizes the posterior distribution of parameters from the assumed prior distribution. Despite its high generalization power, Bayesian techniques are limited by large computational complexity due to the calculation of Hessian matrix. Bootstrap method is the most popular among the four conventional multi-step PI construction methods (Errouissi, Cardenas-Barrera, Meng, Castillo-Guerra, Gong, & Chang, 2015; Zio, 2006; Dybowski & Roberts, 2001). It includes smooth, parametric, wild, pairs, residual, Gaussian process, and other types of bootstrap techniques. In NN-based pairs bootstrap algorithm, for example, the key is to generate bootstrapped pairs by uniform sampling with replacement from the original training data. The estimation is then conducted for a single bootstrapped dataset (Kabir et al., 2018).

Direct Prediction Intervals Construction Methods. This category of methods can tackle some of the limitations in previous methods, such as high demanding in computational power and stringent assumptions. When NN models are constructed through direct training without any assumptions, they can provide more adaptive and smarter PIs for any distribution of targets (Chu, Li, Pedro, & Coimbra, 2015). Lower Upper Bound estimation method is such a technique that can be applied to arbitrary distribution of targets with more than one order reduced computation time. It directly calculates the lower and the upper bounds through trained NNs. Initially, Lower Upper Bound estimation NNs are optimized with the coverage width-based criterion, which presents several limitations. With all the benefits of the original Lower Upper Bound estimation method, the NN-based Direct Interval Forecasting method (Wan, Xu, Pinson, Dong, & Wong, 2013) has much shorter computation time and narrower PIs credited to the improved cost function and the reduced average coverage error. Other approaches for improving the cost function of Lower Upper Bound estimation include the normalized root-mean-square width and particle swarm optimization (Quan, Srinivasan, & Khosravi, 2014), optimal system by (Hosen, Khosravi, Nahavandi, & Creighton, 2014), the independent width and penalty factors (Khosravi, Nahavandi, Srinivasan, & Khosravi, 2014), the deviation from mid-interval consideration (Marín, Valencia, & Sáez, 2016), and the deviation information-based criterion (Zhang, Wu, Wong, Xu, Dong, & Iu, 2014).

4. Roles of SRAs

So far, we have introduced the essentials of SRAs to achieve the expected ethical responsibilities. But pragmatic questions regarding their intended use remain: How to operationalize SRAs? What can SRAs eventually do for societal well-being to address societal challenges? Both ethical and philanthropic responsibilities are indispensable ingredients of the answers. While the ultimate goal of SRAs is to *do good* and *be a good AI citizen*, their ethical responsibilities should be ensured first. When AI fails to fulfill its ethical responsibilities, its philanthropic benefits can be insignificant. For instance, despite the immense public good of COVID-19 vaccines, there has been great controversy about algorithms for their distribution, which have been shown to be inequitable (Branswell, 2021). Some argue that distribution algorithms should prioritize saving more lives and bringing the economy back more rapidly (Cowen, 2020); they support such an ‘unfair’ allocation, but we would argue that that is not unfairness, but simply a difference of values and ethics. In our view, roles of SRAs are expected to encompass both ethical and philanthropic responsibilities. In this survey, we describe three dimensions that SRAs can help with to improve the quality of human life as illustrated in Figure 7: Protect (e.g., protect users’ personal information), Inform (e.g., fake news early detection), and Prevent/Mitigate (e.g., cyberbullying mitigation). We illustrate each dimension with research findings in several emerging societal issues. Particularly, for *protecting* dimension, we focus on privacy preserving and data dignity; for *informing* and *preventing/mitigating* dimensions, we discuss three societal issues that raise growing concerns recently: disinformation, abusive language, and unwanted bias. Because there are many various forms of abusive language such as hate speech and profanity, and the body of work related to each form is vast and diverse, spanning multiple interconnected disciplines, this survey uses the form of cyberbullying as a representative for the illustrations.

4.1 Protecting

The protecting dimension aims to cover or shield humans (especially the most vulnerable or at-risk) from harm, injury, and negative impact of AI systems, in order to intervene. This can be the protection of users’ personal data and their interactions with AI systems. Two typical examples are privacy preserving and data dignity.

4.1.1 PRIVACY-PRESERVING

The capability of deep learning models has been greatly improved by the emerging powerful infrastructures such as clouds and collaborative learning for model training. The fuel of this power, however, comes from data, particularly sensitive data. This has raised growing privacy concerns such as illegitimate use of private data and the disclosure of sensitive data (Boulemtafes, Derhab, & Challal, 2020; Beigi & Liu, 2020). Existing threats against privacy are typically from attacks such as the adversarial examples we discussed in Sec. 3.4.2. Specifically, there are direct information exposure (e.g., untrusted clouds), which is caused by direct intentional or unintentional data breaches, and indirect (inferred) information exposure (e.g., parameter inference), which is caused by direct access to the model or output. Existing privacy-preserving mechanisms can be classified into three cat-

egories, namely, private data aggregation methods, private training, and private inference (Mirshghallah, Taram, Vepakomma, Singh, Raskar, & Esmailzadeh, 2020).

Data aggregation methods are either context-free or context-aware. A context-free approach such as differential privacy (Dwork, 2008), is unaware of the context or what the data will be used for. Context-aware approach such as information-theoretic privacy (Schaefer, Boche, Khisti, & Poor, 2017), on the other hand, is aware of the context in which the data will be used. A naïve technique for privacy protection is to remove identifiers from data, such as name, address, and zip code. It has been used for protecting patients’ information while processing their medical records, but the results are unsatisfying (Sweeney, 2002; Narayanan & Shmatikov, 2008; Homer, Szeling, Redman, Duggan, Tembe, Muehling, Pearson, Stephan, Nelson, & Craig, 2008). The k -Anonymity method can prevent information from re-identification by showing at least k samples with exact same set of attributes for given combination of attributes that the adversary has access to (Sweeney, 2002). The most commonly used data aggregation method is differential privacy, which aims to estimate the effect of removing an individual from the dataset and keep the effect of the inclusion of one’s data small. Some notable work includes the Laplace mechanism (Dwork, McSherry, Nissim, & Smith, 2006), differential privacy with Advanced Composition (Dwork, Roth, et al., 2014), and local differential privacy (Kairouz, Oh, & Viswanath, 2014; Erlingsson, Pihur, & Korolova, 2014).

Information-theoretic privacy is a context-aware approach that explicitly models the dataset statistics. By contrast, context-free methods assume worst-case dataset statistics and adversaries. This line of research was studied by Diaz, Wang, Calmon, and Sankar (2019), Pinceti, Kosut, and Sankar (2019), and Varodayan and Khisti (2011). The second type of privacy-preserving mechanism works during the training phase. Established work in private training is mostly used to guarantee differential privacy or semantic security and encryption (Goldwasser & Micali, 1984). The two most common methods for encryption are homomorphic encryption (Gentry, 2009) and secure multi-party computation (Makri, Rotaru, Smart, & Vercauteren, 2019). The third type of privacy-preserving mechanism works during the inference phase. It aims at the trained systems that are deployed to offer inference-as-a-service (Mirshghallah et al., 2020). Most methods in private inference are similar to those in private training, except for the information-theoretic privacy (Malekzadeh, Clegg, Cavallaro, & Haddadi, 2019, 2018, 2020). It is typically used to offer information-theoretic mathematical or empirical evidence of how these methods operate to improve privacy. There is also work using differential privacy (Wang, Zhang, Bao, Zhu, Cao, & Yu, 2018a), homomorphic encryption (Gilad-Bachrach, Dowlin, Laine, Lauter, Naehrig, & Wernsing, 2016; Chabanne, de Wargny, Milgram, Morel, & Prouff, 2017), and secure multi-party computation (Liu, Juuti, Lu, & Asokan, 2017).

4.1.2 DATA DIGNITY

Beyond privacy preserving, what is more urgent to accomplish is data dignity. It allows users to have absolute control to how their data is being used and they are paid accordingly (Getoor, 2019). Data dignity encompasses the following aspects (Grover, 2020):

- To help users objectively determine the benefits and risks associated with their digital presence and personal data.

- To let users control how their data will be used and the purpose of using the data.
- To allow users to negotiate the terms of using their data.
- To give users complete right and autonomy to be found, analyzed, or forgotten, apart from the fundamental right over their data.

There are business models such as the Microsoft Data Bank designed to give users the control of their data and those shared by the Art of Research (Hart, 2019) about how people can buy and sell their personal data.

4.2 Informing

The informing dimension aims to deliver the facts or information to users, particularly the potential negative results, in a timely way. We illustrate it with a focus on the discussions of detecting disinformation, cyberbullying, and bias.

4.2.1 DISINFORMATION DETECTION

Disinformation is false information that is deliberately created and spread to deceive people, a social group, organization, or country (Pacepa & Rychlak, 2013). The online information ecosystem is never short of disinformation and misinformation, and the growing concerns have been raised recently. Tackling disinformation is rather challenging mainly because (1) disinformation exists almost in all domains; (2) it is ever-changing with new problems, challenges, and threats emerging every day; (3) it entails the joint efforts of interdisciplinary research – computer science, social science, politics, policy making, and psychology, cognitive science (Bhattacharjee, Shu, Gao, & Liu, 2020). Accurate and efficient identification of disinformation is the core to combat disinformation. Existing prominent approaches for disinformation detection primarily rely on news content, social context, user comments, fact-checking tools, and explainable and cross-domain detection.

Early work on disinformation detection has been focused on hand-crafted features extracted from text, such as lexical and syntactic features (Feng, Banerjee, & Choi, 2012; Ott, Choi, Cardie, & Hancock, 2011). Apart from text, online platforms also provide abundant social information that can be leveraged to enrich the textual features, e.g., number of retweets and likes on Twitter. Informed by theories in social science and network science, another line of work exploits social network information to improve the detection performance. Common features are social context (Shu, Wang, & Liu, 2019), user profile (Shu, Wang, & Liu, 2018), user engagement (Shu, Mahudeswaran, Wang, Lee, & Liu, 2020), and relationships among news articles, readers, and publishers (Della Vedova, Tacchini, Moret, Ballarin, DiPierro, & de Alfaro, 2018). A unique function of online platforms is that they allow users to interact through comments. Recent work has shown that user comments can provide weak supervision signal for identifying the authenticity of news articles, which enables early detection of disinformation (Shu, Zheng, Li, Mukherjee, Awadallah, Ruston, & Liu, 2020). When the user comments are unavailable, it is possible to learn users’ response to news articles and then generate user responses (Qian, Gong, Sharma, & Liu, 2018). Fact-checking can be achieved manually or automatically. Manual fact-checking relies on domain experts or crowdsourced knowledge from users. Automatic fact-checking uses structure knowledge bases such as knowledge graph to verify the authenticity of news articles,

see, e.g., (Ciampaglia, Shiralkar, Rocha, Bollen, Menczer, & Flammini, 2015). Beyond within-domain detection, other tasks such as cross-domain detection (Janicka, Pszona, & Wawer, 2019), explanation (Shu, Cui, Wang, Lee, & Liu, 2019), and causal understanding of fake news dissemination (Cheng, Guo, Shu, & Liu, 2021) have also been discussed in literature.

4.2.2 CYBERBULLYING DETECTION

Cyberbullying differs from other forms of abusive language in that it is not an one-off incident but “aggressively intentional acts carried out by a group or an individual using electronic forms of contact, *repeatedly* or *over time* against victims who cannot easily defend themselves” (Smith, Mahdavi, Carvalho, Fisher, Russell, & Tippett, 2008). The increasingly reported number of cyberbullying cases on social media and the resulting detrimental impact have raised great concerns in society. Cyberbullying detection is regularly figured as a binary classification problem. While it shares some similarities with document classification, it should be noted that cyberbullying identification is inherently more complicated than simply identifying oppressive content (Salawu, He, & Lumsden, 2017).

Distinct characteristics of cyberbullying such as *power imbalance* and *repetition* of aggressive acts are central to marking a message or a social media session (Cheng, Silva, Hall, & Liu, 2020b) as cyberbullying. Several major challenges in cyberbullying detection have been discussed in literature such as the formulation of the unique bullying characteristics, e.g., repetition, data annotation, and severe class imbalance. Depending on the employed features, established work can be classified into four categories: content-based, sentiment-based, user-based, and network-based methods. Features extracted from social media content are lexical items such as keywords, Bag of Words, pronoun and punctuation. Empirical evaluations have shown that textual features are the most informative predictors for cyberbullying detection (Cheng, Shu, Wu, Silva, Hall, & Liu, 2020a). For instance, using number of offensive terms as content features is effective in detecting offensive and cursing behavior (Dinakar, Jones, Havasi, Lieberman, & Picard, 2012; Dadvar, Trieschnigg, Ordelman, & de Jong, 2013; Kontostathis, Reynolds, Garron, & Edwards, 2013); Computing content similarity between tweets from different users can help capture users’ personality traits and peer influence, two important factors of cyberbullying occurrences (Cheng, Li, Silva, Hall, & Liu, 2019a). Sentiment-based features typically include key-words, phrases and emojis, and they are often combined with content-based features (Dani, Li, & Liu, 2017). A notable work (Xu, Jun, Zhu, & Bellmore, 2012) identified seven types of emotions in tweets such as anger, empathy, and fear. User-based features are typical characteristics of users, e.g., personality (e.g., hostility), demographics (e.g., age), and user activity (e.g., active users (Balakrishnan, 2015)). Hostility and neuroticism are found to be strongly related to cyberbullying behavior (Biel, Aran, & Gatica-Perez, 2011; Mishna, Khoury-Kassabri, Gadalla, & Daciuk, 2012). Further, gender and age are indicative of cyberbullying in certain cases (Al-garadi, Varathan, & Ravana, 2016). Network-based features measure the sociability of online users, e.g., number of friends, followers, and network embeddedness (Cheng, Li, Silva, Hall, & Liu, 2019b; Huang, Singh, & Atrey, 2014). In addition, a number of methods seek to capture the temporal dynamics to characterize the repetition of cyberbullying, such

as (Cheng, Guo, Silva, Hall, & Liu, 2019; Chen & Li, 2020; Ge, Cheng, & Liu, 2021; Cheng, Guo, Silva, Hall, & Liu, 2021).

4.2.3 BIAS DETECTION

Compared to the well-defined notions of fairness, bias detection is much less studied and the solution is not as straightforward as it may seem (Fu, Huang, & Singh, 2020). The challenges arise from various perspectives. First, the data and algorithms used to make a decision are often not available to policy makers or enforcement agents. Second, algorithms are becoming increasingly complex and the uninterpretability limits an investigator’s ability to identify systematic discrimination through analysis of algorithms. Rather, they have to examine the output from algorithms to check for anomalous results, increasing the difficulty and uncertainty of the task.

Data exploratory analysis is a simple but effective tool to detect data bias. In this initial step of data analysis, we can use basic data statistics and visual exploration to understand what is in a dataset and the characteristics of the data. For algorithmic bias, one of the earliest methods is to compare the selection rate of different groups. Discrimination is highly possible if the selection rate for one group is sufficiently lower than that for other groups. For example, the US Equal Employment Opportunity Commission (EEOC) advocates the “four-fifths rule” or “80% rule” (Feldman, Friedler, Moeller, Scheidegger, & Venkatasubramanian, 2015) to identify a disparate impact. Suppose Y denotes a binary class (e.g., hire or not), A is the protected attribute (e.g., gender), a dataset presents disparate impact if

$$\frac{Pr(Y = 1|A = 0)}{Pr(Y = 1|A = 1)} \leq \tau = 0.8. \tag{3}$$

However, statistical disparity does not necessarily indicate discrimination. If one group has disproportionately more qualified members, we may expect the differences between groups in the results.

A more frequently used approach is regression analysis (Ayres, 2010), which is performed to examine the likelihood of favorable (or adverse) decisions across groups based on sensitive attributes. A significant, non-zero coefficient of the sensitive attributes given a correctly specified regression signals the presence of discrimination. However, we cannot guarantee to observe all the factors the decision maker considers. Therefore, instead of using rate at which decisions are made (e.g., the loan approval rates), bias detection can be based on the success rate of the decisions (e.g., the payback rate of the approved applicants (Becker, 1993)), i.e., the *outcome test*. Another less popular statistical approach for bias detection is benchmarking. The major challenge of benchmarking analysis is identifying the distribution of the sensitive attributes of the benchmark population where sensitive attributes are unlikely to influence the identification of being at-risk. Some solutions can be seen in (McConnell & Scheidegger, 2001; Lange, Blackman, & Johnson, 2001). Recently, AI researchers have developed tools to automatically detect bias. For instance, drawing on techniques in natural language processing and moral foundation theories, the tool by Mokherian, Abeliuk, Cummings, and Lerman (2020) can understand structure and nuances of content consistently showing up on left-leaning and right-leaning news sites, aiming to help consumers better prepare for unfamiliar news source. In earlier efforts, an

international research group launched a non-profit organization Project Implicit⁶ in 1998 aimed at detecting implicit social bias.

4.3 Preventing/Mitigating

If both of the first two dimensions fail, we may rely on the last dimension to prevent/mitigate the negative impact of socially indifferent AI algorithms on the end-users. We continue the discussions about disinformation, cyberbullying, and bias, with a focus on the prevention and mitigation strategies.

4.3.1 DISINFORMATION PREVENTION/MITIGATION

Preventing the generation/spread of disinformation and mitigating its negative impact is an urgent task because disinformation typically spread faster (Vosoughi, Roy, & Aral, 2018) than normal information due to the catchy news content and the ranking algorithms operating behind the online news platforms. To increase user engagement, social recommender systems are designed to recommend popular posts and trending content. Therefore, disinformation often gains more visibility. An effective approach for disinformation mitigation is to govern this visibility of news, e.g., recommendation and ranking based algorithms. Mitigation also relates to early detection.

Network intervention can slow down the spread of disinformation by influencing the exposed users in a social network. For example, we can launch a counter-cascade that consists of fact-checked version of false news articles. This is commonly referred to as the influence limitation or minimization problem (Bhattacharjee et al., 2020). Given a network with accessible counter-cascade, the goal is to find a (minimum) set of nodes in this network such that the effect of the original cascade can be minimized. A variety of approximation algorithms (Budak, Agrawal, & El Abbadi, 2011; Nguyen, Yan, Thai, & Eidenbenz, 2012) have been proposed to solve the NP-hard problem and the variants. When applied to disinformation mitigation, they seek to inoculate as many nodes as possible in a short period of time. It is possible to extend the two cascades into tasks with multiple cascades, where we can further consider the different priorities of these cascades, i.e., each cascade influences the node in the network differently (Tong, Du, & Wu, 2018). The second method for disinformation mitigation is content flagging: social media platforms allow users to ‘flag’ or ‘report’ a news content if they find it offensive, harmful, and/or false. Big social media companies such as Facebook hired professional moderators to manually investigate and/or remove these content. However, considering the millions of news generated/spread every minute (Marr, 2021), it is impractical for these moderators to manually review all the news. The solution comes to the crowd wisdom – users can choose to ‘flag’ the content if it violates the community guidelines of the platform. Some platforms can further provide feedback for these users about if their fact-check is correct or not. User behavior is an effective predictor for disinformation detection (Cheng et al., 2021), therefore, the third prevention method leverages the differences between user behaviors to identify susceptible or gullible users. For example, it is shown in (Rajabi, Shehu, & Purohit, 2019) that groups of vulnerable Twitter users can be identified in fake news consumption. Other studies (Guess, Nagler, & Tucker, 2019) also suggest that older people are more likely to spread disinformation.

6. <https://implicit.harvard.edu/implicit>

4.3.2 CYBERBULLYING PREVENTION/MITIGATION

In contrast to the large amount of work in cyberbullying detection, efforts for its prevention and mitigation have been a few. Some research suggests that prevention/mitigation strategies are defined at different levels (Smith et al., 2008). At technological level, we can consider providing parental control service, firewall blocking service, online services rules, text-message control, and mobile parental control, e.g., KnowBullying and BullyBlocker (Silva, Rich, & Hall, 2016). Another effective tool is psychological approach, such as talking and listening to cyber-victims, providing counseling services, encouraging victims to make new relations and join social clubs. At education level, we are responsible to educate end-users, help improve their technical and cognitive skills. At administrative level, it is important for organizations and government to develop policies to regulate using free service and enhance workplace environment. Therefore, the goal of cyberbullying prevention/mitigation can only be accomplished with interdisciplinary collaborations, e.g., psychology, public health, computer science, and other behavioral and social sciences (Kraft & Wang, 2009). One example is that computer and social scientists attempted to understand behavior of users in realistic environments by designing social media site for experimentation such as controlled study and post-study survey (DiFranzo, Taylor, Kazerooni, Wherry, & Bazarova, 2018; Ashktorab, 2017).

Existing solutions to preventing cyberbullying can report/control/warn about message content (e.g., Dinakar et al., 2012; Vishwamitra, Zhang, Tong, Hu, Luo, Kowalski, & Mazer, 2017), provide support for victims (e.g., Vishwamitra et al., 2017), and educate both victims and bullies (e.g., Dinakar et al., 2012). A variety of anti-bully apps are also available to promote well-being of users. For example, NoMoreBullyingMe App provides online meditation techniques to support victims; “Honestly” App (Shaul, 2015) encourages users to share positive responses with each other (e.g., sing a song). However, current cyberbullying prevention strategies often do not work as desired because of the complexity and nuance with which adolescents bully others online (Ashktorab & Vitak, 2016).

4.3.3 BIAS MITIGATION

Prior approaches for bias mitigation focus on either designing fair machine learning algorithms or theorizing on the social and ethical aspects of machine learning discrimination (Caton & Haas, 2020). From the technical aspect, approaches to fairness typically can be categorized into pre-processing (prior to modelling), in-processing (at the point of modelling), and post-processing (after modelling). One condition to use pre-processing approaches is that the algorithm is allowed to modify the training data (Bellamy, Dey, Hind, Hoffman, Houde, Kannan, Lohia, Martino, Mehta, Mojsilović, et al., 2019). We can then transform the data to remove the discrimination (d’Alessandro, O’Neil, & LaGatta, 2017). In-processing approaches eliminate bias by modifying algorithms during the training process (d’Alessandro et al., 2017). We can either incorporate fairness notion into the objective function or impose fairness constraint (Bellamy et al., 2019; Berk, Heidari, Jabbari, Joseph, Kearns, Morgenstern, Neel, & Roth, 2017). When neither training data nor model can be modified, we can use post-processing approaches to reassign the predicted labels based on a defined function and a holdout set which was not used in the model training phase (Bellamy et al., 2019; Berk et al., 2017). Most of these approaches are built on the notion of protected

or sensitive variables that define the (un)privileged groups. Commonly used protected variables are age, gender, marital status, race, and disabilities. A shared characteristic of these groups is they are disproportionately (less) more likely to be positively classified. Fairness measures are important to quantify fairness in the development of fairness approaches. However, creating generalized notions of fairness quantification is a challenging task (Caton & Haas, 2020). Depending on the protected target, fairness metrics are usually designed for individual fairness (e.g., every one is treated equally), group fairness (e.g., different groups such as women vs men are treated equally), or subgroup fairness. Drawing on theories in causal inference, individual fairness also includes counterfactual fairness which describes that a decision is fair towards an individual if the result was same when s/he had taken a different sensitive attribute (Kusner et al., 2017).

Recent years have witnessed immense progress of fair machine learning – a variety of methods have been proposed to address bias and discrimination over different applications. We focus on two mainstream methods: fair classification and fair regression. A review of machine learning fairness can be referred to (Mehrabi et al., 2019; Caton & Haas, 2020).

(1) *Fair Classification*. For a (binary) classifier with sensitive variable S , the target variable Y , and the classification score R , general fairness desiderata have three “non-discrimination” criteria: *Independence*, i.e., $R \perp\!\!\!\perp S$; *Separation*, i.e., $R \perp\!\!\!\perp S | Y$; and *Sufficiency*, i.e., $Y \perp\!\!\!\perp S | R$. Fair machine learning algorithms need to adopt/create specific fairness definitions that fit into context (Goel, Yaghini, & Faltings, 2018; Kamishima, Akaho, Asoh, & Sakuma, 2012; Krasanakis, Spyromitros-Xioufis, Papadopoulos, & Kompatsiaris, 2018; Menon & Williamson, 2018; Calders & Verwer, 2010). Common methods in fair classification include blinding (Calmon, Wei, Vinzamuri, Ramamurthy, & Varshney, 2017; Hardt, Price, & Srebro, 2016), causal methods (Galhotra, Brun, & Meliou, 2017; Johnson, Foster, & Stine, 2016), transformation (Feldman et al., 2015; Dwork, Hardt, Pitassi, Reingold, & Zemel, 2012; Wei, Ramamurthy, & Calmon, 2020), sampling and subgroup analysis (Celis, Deshpande, Kathuria, & Vishnoi, 2016; Dwork, Immorlica, Kalai, & Leiserson, 2018), adversarial learning (Feng, Yang, Lyu, Tan, Sun, & Wang, 2019; Xu et al., 2019; Sattigeri, Hoffman, Chenthamarakshan, & Varshney, 2019), reweighing (Jiang & Nachum, 2020; Calders & Verwer, 2010), and regularization and constraint optimization (Kamishima et al., 2012; Bechavod & Ligett, 2017; Cheng, Mosallanezhad, Silva, Hall, & Liu, 2021).

(2) *Fair Regression*. The goal of fair regression is to jointly minimize the difference between true and predicted values and ensure fairness. It follows the general formulation of fair classification but with continuous rather than binary/categorical target variable. Accordingly, the fairness definition, metrics, and the basic algorithms are adapted from classification to regression. For example, it is suggested using statistical parity and bounded-group-loss metrics to measure fairness in regression (Agarwal, Dudík, & Wu, 2019). Bias in linear regression is considered as the effects of a sensitive attribute on the target variable through the mean difference between groups and AUC metrics (Calders, Karim, Kamiran, Ali, & Zhang, 2013). One commonly used approach in fair regression is regularization, e.g., (Kamishima et al., 2012; Berk et al., 2017).

Apart from fair machine learning, algorithm operators are encouraged to share enough details about how research is carried out to allow others to replicate it. This is a leap for mitigating bias as it helps end-users with different technical background to understand how the algorithm works before making any decision. It is also suggested that AI technologists




Figure 8: Primary challenges and open problems we confront in developing SRAs. Some challenges relate to SRAs’ internal mechanisms that fulfill AI’s ethical responsibilities whilst others relate to SRAs’ roles to which both ethical and philanthropic responsibilities are the keys.

and researchers develop a bias impact statement as a self-regulatory practice. It can help probe and avert any potential biases that are injected into or resultant from algorithmic decision (Lee, Resnick, & Barton, 2019). Some example questions in the statement are “What will the automated decision do?”, “How will potential bias be detected?”, and “What are the operator incentives”. In algorithm design, researchers are also responsible to encourage the role of diversity within the team, training data, and the level of cultural sensitivity. The “diversity-in-design” mechanism aims to take deliberate and transparent actions to address the upfront cultural biases and stereotypes. Furthermore, we might also consider updating nondiscrimination and other civil rights laws to interpret and redress online disparate impacts (Lee et al., 2019). An example of such consideration is to unambiguously define the thresholds and parameters for the disparate treatments of protected groups before the algorithm design.

5. Open Problems and Challenges

This survey reveals that the current understanding of SRAs is insufficient and future efforts are in great need. Here, we describe several primary challenges, as summarized in Figure 8, in an attempt to broaden the discussions on future directions and potential solutions.

Causal Learning. The correlation fallacy causes AI algorithms to meet with fundamental obstacles in order to commit social responsibility. These obstacles are robustness, explainability, and cause-effect connections (Pearl, 2019). The era of big data has changed the ways of learning causality, and meanwhile, causal learning becomes an indispensable ingredient for AI systems to achieve human-level intelligence. There are a number of benefits to incorporate causality in the next-generation of AI. For example, teaching AI algorithms to understand “why” can help them transfer their knowledge to different but similar domains. Early efforts in SRAs attempted to employ causal learning concept and methods such as intervention, counterfactual, *do*-calculus, propensity scoring to address fairness (e.g., counterfactual fairness) and interpretability (causal interpretability) issues. They have shown prominent results in these tasks.

Context Matters. Context is the core to SRAs due to its inherently elaborate nature, e.g., the “Transparency Paradox”. Understanding and quantifying the relationships among

the various principles (some are tradeoffs and some are not), e.g., fairness, transparency, and safety, have to be placed in specific context. One such context is the *social context*. Existing SRAs (e.g., fair machine learning), once introduced into a new social context, may render current technical interventions ineffective, inaccurate, and even dangerously misguided (Selbst, Boyd, Friedler, Venkatasubramanian, & Vertesi, 2019). A recent study (Sühr, Hilgard, & Lakkaraju, 2020) found that while fair ranking algorithms such as Det-Greedy (Geyik, Ambler, & Kenthapadi, 2019) help increase the exposure of minority candidates, their effectiveness is limited by the job contexts in which employers have a preference to particular genders. How to properly integrate social context into SRAs is still an open problem. *Algorithmic context* (e.g., supervised learning, unsupervised learning, and reinforcement learning) is also extremely important when designing SRAs for the given data. A typical example is the feedback loop problem in predictive policing (Ensign, Friedler, Neville, Scheidegger, & Venkatasubramanian, 2018). A subtle algorithmic choice can have huge ramifications on the results. Consequently, we need to understand the algorithmic context to make the right algorithmic choices when designing socially responsible AI systems. Designing context-aware SRAs is the key to achieving Social Responsibility of AI.

Responsible Model Release and Governance. Nontransparent model reporting is one of the main causes of AI indifferent behaviors. As a critical step to clarify the intended use cases of AI systems and the contexts for which they are well suited, responsible model release and governance has been receiving growing attentions from both industry and academia. One role of SRAs is to bring together the tools, solutions, practices, and people to govern the built AI systems across its life cycle (Dobrin, 2020). At this early stage, some research results suggested that released models be accompanied by documentation detailing various characteristics of the systems, e.g., what it does, how it works, and why it matters. For example, the AI FactSheets (Arnold, Bellamy, Hind, Houde, Mehta, Mojsilović, Nair, Ramamurthy, Olteanu, Piorkowski, et al., 2019) advocates to use a factsheet completed and voluntarily released by AI developers to increase the transparency of their services. A similar concept is model cards (Mitchell, Wu, Zaldivar, Barnes, Vasserman, Hutchinson, Spitzer, Raji, & Gebru, 2019), short documents that provide benchmarked evaluation for the trained AI models in a variety of conditions, e.g., different cultural or demographic groups. Typically, a model card should include the model details, intended use, evaluation metrics, training/evaluation data, ethical considerations, and caveats and recommendations. To help increase transparency, manage risk, and build trust in AI, AI technologists and researchers are responsible to address various challenges faced in creating useful AI release documentation (Hind, Houde, Martino, Mojsilovic, Piorkowski, Richards, & Varshney, 2020) and develop effective AI governance tools.

AI Defenses. Developing AI systems that outwit malicious AI is still at an early stage (Knight, 2020). Since we have not fully understood how AI systems work, they are not only vulnerable to attack but also likely to fail in surprising ways (Yuan et al., 2019; Chakraborty et al., 2018). As a result, it is critical and urgent to work on designing systems that are provably robust to help ensure that the AI systems are not vulnerable to adversaries. At least two capabilities an “AI firewall” needs to be equipped with: one capability is to probe an AI algorithm for weaknesses (e.g., perturb the input of an AI system to make it misbehave) and the other one is to automatically intercept potentially problematic inputs. Some big tech companies have started building their own AI defenses to identify the weak

spots, e.g., the “red team” in Facebook, the software framework released by Microsoft, Nvidia, IBM, and 9 other companies. AI defenses reflect the fundamental weakness in modern AI and make AI systems more robust and intelligent.

AI Ethics Principles and Policies. Current AI principles and policies for ethical practice have at least two common criticisms: (1) they are too vaguely formulated to prove to be helpful in guiding practice; and (2) they are defined primarily by AI researchers and powerful people with mainstream populations in mind (Young, Magassa, & Friedman, 2019). For the first criticism, to help operationalize AI principles in practice and organizations confront inevitable value trade-offs, it has been suggested to redefine AI principles based on philosophical theories in applied ethics (Canca, 2020). Particularly, it categorizes published AI principles (e.g., fairness, accountability, and transparency) into three widely used *core principles* in applied ethics: autonomy, beneficence (avoiding harm and doing good), and justice. The core principles “invoke those values that theories in moral and political philosophy argue to be intrinsically valuable, meaning their value is not derived from something else” (Canca, 2020). Existing AI principles are *instrumental principles* that “build on concepts whose values are derived from their instrumental effect in protecting and promoting intrinsic values” (Canca, 2020). Operationable AI principles help effectively put ethical AI in practice and reduce the responsible AI Gap in companies. To address the second criticism, we need to best elicit the inputs and values of diverse voices from the Subjects of SRAs, i.e., the minority and disadvantaged groups, and incorporate their perspectives into the tech policy document design process. If we align values of AI systems through a panel of people (who are compensated for doing this), they too can influence the system behavior, and not just the powerful people or AI researchers.

Understanding *Why*. Many AI systems are designed and developed without fully understanding *why*: What do we wish the AI system do? This is often the reason that these systems fail to represent the goals of the real tasks, a primary source of AI risks. The problem can become more challenging when the AI system is animated through a number of lines of code that lack nuance, creating a machine that does not align with our true intentions. As the first step, understanding *why* clearly defines our social expectation of AI systems and paves way for more specific questions such as “What is the problem? Who will define it? and what are the right people to include?”. Answering *why* helps us effectively abolish the development of socially indifferent AI systems in the first place and also helps understand the kinds of deception an AI system may learn by itself.

Long-Term Effect. SRAs include social concepts such as fairness that can evolve over time along with the constant changes of human values and social dynamics. This raises the concerns about the commitment SRAs need to fulfill in the long term. For example, despite the various types of fairness definitions, once introduced into the dimension of time, the number of fairness definitions may be explosive. In addition, current fairness criteria may be considered as unfair in the future. Fairness criteria are essentially designed to promote long-term well-being. However, even a static fairness notion can fail to protect the target groups when there is a feedback loop in the overall system (Liu, Dean, Rolf, Simchowitz, & Hardt, 2018). How to build AI systems that can commit long-term responsibility is extremely challenging and rarely studied thus far. Initial results of long-term fairness (Liu et al., 2018; Hu & Chen, 2018) highlight the importance of measurement and temporal modeling in the evaluation of fairness criteria.

Humans in the Loop. While existing techniques in SRAs have indeed made significant progress towards responsible AI systems, their usefulness can be limited in some settings where the decisions made are actually poorer for every individual. For issues of fairness in prediction, for example, many findings (e.g., Pfohl, Foryciarz, & Shah, 2020) have shown the concerns about the fairness-performance trade-off: the imposition of fairness comes at a cost to model performance. Predictions are less reliable and moreover, different notions of fairness can make approaches to fairness conflict with one another. Having human in the loop matters when it comes to contextualizing the objectives of SRAs, especially for high-stake decisions. For instance, there are situations where the cut-off values of fairness for two subgroups are different, and humans can help calibrate the differences.

Responsible AI Gap in Industry. The far-reaching effect of reputational damage and employee disengagement result from AI misbehavior has forced company executives to begin understanding the risks of poorly designed AI systems and the importance of SRAs. While seeing many potential benefits of developing responsible AI systems such as increasing market share and long-term profitability, companies lack the knowledge of how to cross the “Responsible AI Gap” between principles and tangible actions (Mills, Baltassis, Santinelli, Carlisi, Duranton, & Gallego, 2020). This is partly because companies view responsible AI solely as risk-avoidance mechanism and overlook its financial rewards. To capture the benefits of responsible AI in companies’ day-to-day business, companies need to go far beyond SRAs and examine every aspect of the end-to-end AI systems. A recent article (Mills et al., 2020) suggested six basic steps to bridge the gulf between responsible AI and the reality: *Empower responsible AI leadership, Develop principles, policies, and training, Establish human and AI governance, Conduct Responsible AI reviews, Integrate tools and methods, and Build and test a response plan.* Even though the gap might be huge, small efforts built over time can let SRAs achieve a transformational impact on the businesses.

Interdisciplinary Research. Current public dialog on SRAs has been focused on a narrow subset of fields, blinding us to the opportunities presented by interdisciplinary research. It is necessary to work with researchers from different disciplines whose contributions are sorely needed, e.g., psychologist, social scientist, educators, and humanities. Non-profit organizations are both the beneficiaries and benefactors of SRAs. In partnering with non-profits and social enterprises will not only unleash AI’s potential for benefiting societal well-being, but also let AI technologists and researchers have the opportunity to encounter the real problems we are currently facing. A better understanding of what problems need to be solved helps identify SRAs that need to be created. Moreover, as big tech companies bankroll more work of academic researchers, much of ethics-based research gets concentrated in the hands of a few companies that can afford it (Rivero, 2020). This is problematic because we are over reliant on the same companies that are producing socially indifferent AI systems. We need interdisciplinary and decentralized research to create SRAs and simultaneously achieve the four levels in the pyramid of Social Responsibility of AI.

SRAs for Social Good. The last challenge regards the intended use of SRAs. When SRAs are leveraged to uplift humanity, a trust in AI is further enhanced. There has been a burgeoning of AI-for-social-good movement that produces AI algorithms to help reduce poverty, hunger, inequality, injustice, climate change, ill health, and other causes of human suffering (Varshney, 2019). Compared to deploying cutting-edge AI systems to solve these critical issues, a more urgent question to examine is “What makes an AI project good” in or-

der to prevent the detrimental consequences of AI. In addition to Protecting, Informing, and Preventing, social good applications also relate closely to *Fundraise* and *Greenlight* (Siegel, 2020). Applying SRAs to target solicitations for donations largely helps with fundraising for non-profits, charitable organizations, and universities. Greenlight describes how SRAs can help allocate grants and other types of resources by predicting the success rates of project proposals. It plays an important role in improving execution effectiveness of organizations. Developing social good applications that leverage power of SRAs to benefit society is an equally endeavor for AI technologists and researchers.

6. Conclusion

This survey examines multiple dimensions of research in Social Responsibility of AI, seeking to broaden the current discussions primarily focused on decision-making algorithms that perform scoring and classification tasks. We argue that having a full scope of AI to capture the connections among all the major dimensions is the key to Socially Responsible AI Algorithms (SRAs). This work starts with an inclusive definition of Social Responsibility of AI, highlighting the principles (e.g., Fairness, Inclusiveness), means (e.g., SRAs), and objective (e.g., improving humanity). To better frame the Social Responsibility of AI, we also introduce the pyramid with four-level responsibilities of AI systems: functional responsibilities, legal responsibilities, ethical responsibilities, and philanthropic responsibilities. We then focus our discussions on how to achieve Social Responsibility of AI via the proposed framework SRAs. In the definition of SRAs, we emphasize that the functional and societal aspects are integral parts of AI algorithms. Given that both the functional and legal responsibilities are the usual focuses in AI research and development, we particularly investigate the essentials to achieve AI’s ethical responsibilities: the subjects, causes, objectives, and means. For the intended use (i.e., roles) of SRAs, we discuss the need of philanthropic and ethical responsibilities for AI systems to protect and inform users, and prevent/mitigate the negative impact. We conclude with several open problems and major challenges in SRAs. At this pivotal moment in the development of AI, it is of vital importance to discuss AI ethics and specify Social Responsibility of AI. Drawing from the theory of moral license (List & Momeni, 2018) – when humans are good, we give ourselves moral license to be bad – we argue that simply asking AI to do good is insufficient and inefficient, and more can be done for AI technologists and researchers to develop socially responsible AI systems. We hope this work can propel future research in various fields to tackle together the challenges and steer a course towards a beneficial AI future.

Acknowledgements

This material is based upon work supported by, or in part by, the U.S. Army Research Laboratory (ARL), the U.S. Army Research Office (ARO), the Office of Naval Research (ONR) under contract/grant numbers W911NF2110030, W911NF2020124, and N00014-21-1-4002, as well as by the National Science Foundation (NSF) grants 1909555 and 2036127. We thank Dr. Lise Getoor and Dr. Hosagrahar V. Jagadish for their invaluable suggestions.

References

- Abdalla, M., & Abdalla, M. (2020). The grey hoodie project: Big tobacco, big tech, and the threat on academic integrity..
- Agarwal, A., Dudík, M., & Wu, Z. S. (2019). Fair regression: Quantitative definitions and reduction-based algorithms..
- Akhtar, Z., & Dasgupta, D. (2019). A brief survey of adversarial machine learning and defense strategies..
- Al-garadi, M. A., Varathan, K. D., & Ravana, S. D. (2016). Cybercrime detection in on-line communications: The experimental case of cyberbullying detection in the twitter network. *Computers in Human Behavior*, *63*, 433–443.
- Angwin, J., & Larson, J. (2015). The tiger mom tax: Asians are nearly twice as likely to get a higher price from princeton review. *Retrieved September, 1, 2015*.
- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). Machine bias. *ProPublica*, *May*, *23*, 2016.
- Arjovsky, M., Bottou, L., Gulrajani, I., & Lopez-Paz, D. (2019). Invariant risk minimization..
- Arnold, M., Bellamy, R. K., Hind, M., Houde, S., Mehta, S., Mojsilović, A., Nair, R., Ramamurthy, K. N., Olteanu, A., Piorkowski, D., et al. (2019). Factsheets: Increasing trust in AI services through supplier’s declarations of conformity. *IBM Journal of Research and Development*, *63*(4/5), 6–1.
- Ashktorab, Z. (2017). *Designing Cyberbullying Prevention and Mitigation Tools*. Ph.D. thesis.
- Ashktorab, Z., & Vitak, J. (2016). Designing cyberbullying mitigation and prevention solutions through participatory design with teenagers. In *CHI*, pp. 3895–3905.
- Asimov, N. (2020) <https://www.govtech.com/em/safety/Stanfords-Vaccine-Algorithm-Left-Frontline-Workers-at-Back-of-Line-.html>.
- Ateniese, G., Mancini, L. V., Spognardi, A., Villani, A., Vitali, D., & Felici, G. (2015). Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers. *International Journal of Security and Networks*, *10*(3), 137–150.
- Ayres, I. (2010). Testing for discrimination and the problem of “included variable bias”,”. Tech. rep., mimeo, Yale Law School.
- Baeza-Yates, R. (2018). Bias on the web. *Communications of the ACM*, *61*(6), 54–61.
- Balakrishnan, V. (2015). Cyberbullying among young adults in malaysia: The roles of gender, age and internet frequency. *Computers in Human Behavior*, *46*, 149–157.
- Barreno, M., Nelson, B., Sears, R., Joseph, A. D., & Tygar, J. D. (2006). Can machine learning be secure?. In *ASIACCS*, pp. 16–25.
- Beale, N., Batty, H., Davison, A. C., & MacKay, R. S. (2019). An unethical optimization principle..

- Bechavod, Y., & Ligett, K. (2017). Penalizing unfairness in binary classification..
- Becker, G. S. (1993). Nobel lecture: The economic way of looking at behavior. *Journal of Political Economy*, 101(3), 385–409.
- Beigi, G., & Liu, H. (2020). A survey on privacy in social media: Identification, mitigation, and applications. *ACM Transactions on Data Science*, 1(1), 1–38.
- Bellamy, R. K., Dey, K., Hind, M., Hoffman, S. C., Houde, S., Kannan, K., Lohia, P., Martino, J., Mehta, S., Mojsilović, A., et al. (2019). AI fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development*, 63(4/5), 4–1.
- Bengio, Y., Deleu, T., Rahaman, N., Ke, R., Lachapelle, S., Bilaniuk, O., Goyal, A., & Pal, C. (2019). A meta-transfer objective for learning to disentangle causal mechanisms..
- Berk, R., Heidari, H., Jabbari, S., Joseph, M., Kearns, M., Morgenstern, J., Neel, S., & Roth, A. (2017). A convex framework for fair regression..
- Bhatt, U., Antorán, J., Zhang, Y., Liao, Q. V., Sattigeri, P., Fogliato, R., Melancon, G. G., Krishnan, R., Stanley, J., Tickoo, O., et al. (2020). Uncertainty as a form of transparency: Measuring, communicating, and using uncertainty..
- Bhattacharjee, A., Shu, K., Gao, M., & Liu, H. (2020). Disinformation in the online information ecosystem: Detection, mitigation and challenges..
- Bickel, P. J., Hammel, E. A., & O’Connell, J. W. (1975). Sex bias in graduate admissions: Data from Berkeley. *Science*, 187(4175), 398–404.
- Biel, J.-I., Aran, O., & Gatica-Perez, D. (2011). You are known by how you vlog: Personality impressions and nonverbal behavior in youtube.. In *ICWSM*, pp. 446–449. Citeseer.
- Blyth, C. R. (1972). On Simpson’s paradox and the sure-thing principle. *Journal of the American Statistical Association*, 67(338), 364–366.
- Boulemtafes, A., Derhab, A., & Challal, Y. (2020). A review of privacy-preserving techniques for deep learning. *Neurocomputing*, 384, 21–45.
- Branswell, H. (2021) <https://www.statnews.com/2021/03/08/a-pandemic-expert-weighs-in-on-the-long-road-ahead-for-covid-19-vaccine-distribution/>.
- Budak, C., Agrawal, D., & El Abbadi, A. (2011). Limiting the spread of misinformation in social networks. In *WWW*, pp. 665–674.
- Cabañas, J. G., Cuevas, Á., Arrate, A., & Cuevas, R. (2020). Does facebook use sensitive data for advertising purposes?. *Communications of the ACM*, 64(1), 62–69.
- Calders, T., Karim, A., Kamiran, F., Ali, W., & Zhang, X. (2013). Controlling attribute effect in linear regression. In *ICDM*, pp. 71–80. IEEE.
- Calders, T., & Verwer, S. (2010). Three naive Bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2), 277–292.
- Calmon, F., Wei, D., Vinzamuri, B., Ramamurthy, K. N., & Varshney, K. R. (2017). Optimized pre-processing for discrimination prevention. In *NeurIPS*, pp. 3992–4001.
- Canca, C. (2020). Operationalizing AI ethics principles. *Communications of the ACM*, 63(12), 18–21.

- Carpenter, J. (2015) <https://www.washingtonpost.com/news/the-intersect/wp/2015/07/06/googles-algorithm-shows-prestigious-job-ads-to-men-but-not-to-women-heres-why-that-should-worry-you/>.
- Carroll, A. B., et al. (1991). The pyramid of corporate social responsibility: Toward the moral management of organizational stakeholders. *Business Horizons*, 34(4), 39–48.
- Carvalho, D. V., Pereira, E. M., & Cardoso, J. S. (2019). Machine learning interpretability: A survey on methods and metrics. *Electronics*, 8(8), 832.
- Caton, S., & Haas, C. (2020). Fairness in machine learning: A survey..
- Celis, L. E., Deshpande, A., Kathuria, T., & Vishnoi, N. K. (2016). How to be fair and diverse?..
- Chabanne, H., de Wargny, A., Milgram, J., Morel, C., & Prouff, E. (2017). Privacy-preserving classification on deep neural network.. *IACR Cryptol. ePrint Arch.*, 2017, 35.
- Chakraborty, A., Alam, M., Dey, V., Chattopadhyay, A., & Mukhopadhyay, D. (2018). Adversarial attacks and defences: A survey..
- Chattopadhyay, A., Manupriya, P., Sarkar, A., & Balasubramanian, V. N. (2019). Neural network attributions: A causal perspective..
- Chen, H.-Y., & Li, C.-T. (2020). Henin: Learning heterogeneous neural interaction networks for explainable cyberbullying detection on social media..
- Chen, Z., & Liu, B. (2018). Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3), 1–207.
- Cheng, L., Guo, R., Shu, K., & Liu, H. (2021). Causal understanding of fake news dissemination on social media. In *KDD*.
- Cheng, L., Guo, R., Silva, Y., Hall, D., & Liu, H. (2019). Hierarchical attention networks for cyberbullying detection on the instagram social network. In *SDM*, pp. 235–243. SIAM.
- Cheng, L., Guo, R., Silva, Y. N., Hall, D., & Liu, H. (2021). Modeling temporal patterns of cyberbullying detection with hierarchical attention networks. *ACM/IMS Transactions on Data Science*, 2(2), 1–23.
- Cheng, L., Li, J., Silva, Y. N., Hall, D. L., & Liu, H. (2019a). PI-bully: Personalized cyberbullying detection with peer influence.. In *IJCAI*, pp. 5829–5835.
- Cheng, L., Li, J., Silva, Y. N., Hall, D. L., & Liu, H. (2019b). Xbully: Cyberbullying detection within a multi-modal context. In *WSDM*, pp. 339–347.
- Cheng, L., Mosallanezhad, A., Silva, Y. N., Hall, D. L., & Liu, H. (2021). Mitigating bias in session-based cyberbullying detection: A non-compromising approach. In *Proceedings of ACL*.
- Cheng, L., Shu, K., Wu, S., Silva, Y. N., Hall, D. L., & Liu, H. (2020a). Unsupervised cyberbullying detection via time-informed gaussian mixture model. In *CIKM*.
- Cheng, L., Silva, Y., Hall, D., & Liu, H. (2020b). Session-based cyberbullying detection: Problems and challenges..

- Chu, Y., Li, M., Pedro, H. T., & Coimbra, C. F. (2015). Real-time prediction intervals for intra-hour dni forecasts. *Renewable Energy*, *83*, 234–244.
- Ciampaglia, G. L., Shiralkar, P., Rocha, L. M., Bollen, J., Menczer, F., & Flammini, A. (2015). Computational fact checking from knowledge networks. *PloS One*, *10*(6), e0128193.
- Cowen, T. (2020) <https://www.bloomberg.com/opinion/articles/2020-11-23/vaccine-distribution-shouldn-t-be-fair>.
- Craven, M., & Shavlik, J. W. (1996). Extracting tree-structured representations of trained networks. In *NeurIPS*, pp. 24–30.
- Dadvar, M., Trieschnigg, D., Ordelman, R., & de Jong, F. (2013). Improving cyberbullying detection with user context. In *ECIR*, pp. 693–696. Springer.
- d’Alessandro, B., O’Neil, C., & LaGatta, T. (2017). Conscientious classification: A data scientist’s guide to discrimination-aware classification. *Big Data*, *5*(2), 120–134.
- Dalvi, N., Domingos, P., Sanghai, S., & Verma, D. (2004). Adversarial classification. In *KDD*, pp. 99–108.
- Dani, H., Li, J., & Liu, H. (2017). Sentiment informed cyberbullying detection in social media. In *ECML PKDD*, pp. 52–67. Springer.
- Datta, A., Tschantz, M. C., & Datta, A. (2015). Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. *PoPETs*, *2015*(1), 92–112.
- Della Vedova, M. L., Tacchini, E., Moret, S., Ballarin, G., DiPierro, M., & de Alfaro, L. (2018). Automatic online fake news detection combining content and social signals. In *2018 FRUCT*, pp. 272–279. IEEE.
- Diaz, M., Wang, H., Calmon, F. P., & Sankar, L. (2019). On the robustness of information-theoretic privacy measures and mechanisms. *IEEE Transactions on Information Theory*, *66*(4), 1949–1978.
- DiFranzo, D., Taylor, S. H., Kazerooni, F., Wherry, O. D., & Bazarova, N. N. (2018). Upstanding by design: Bystander intervention in cyberbullying. In *CHI*, pp. 1–12.
- Dinakar, K., Jones, B., Havasi, C., Lieberman, H., & Picard, R. (2012). Common sense reasoning for detection, prevention, and mitigation of cyberbullying. *ACM Tüs*, *2*(3), 1–30.
- Dobrin, S. (2020) <https://www.ibm.com/blogs/watson/2020/12/how-ibm-is-advancing-ai-governance-to-help-clients-build-trust-and-transparency/>.
- Dwork, C. (2008). Differential privacy: A survey of results. In *TAMC*, pp. 1–19. Springer.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. In *ITCS*, pp. 214–226.
- Dwork, C., Immorlica, N., Kalai, A. T., & Leiserson, M. (2018). Decoupled classifiers for group-fair and efficient machine learning. In *FAT**, pp. 119–133.
- Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pp. 265–284. Springer.

- Dwork, C., Roth, A., et al. (2014). The algorithmic foundations of differential privacy.. *Fnt-TCS*, 9(3-4), 211–407.
- Dybowski, R., & Roberts, S. J. (2001). Confidence intervals and prediction intervals for feed-forward neural networks..
- Ekstrand, M. D., Tian, M., Kazi, M. R. I., Mehrpouyan, H., & Kluver, D. (2018). Exploring author gender in book rating and recommendation. In *RecSys*, pp. 242–250.
- Ensign, D., Friedler, S. A., Neville, S., Scheidegger, C., & Venkatasubramanian, S. (2018). Runaway feedback loops in predictive policing. In *FAT**, pp. 160–171.
- Erhan, D., Bengio, Y., Courville, A., & Vincent, P. (2009). Visualizing higher-layer features of a deep network. *University of Montreal*, 1341(3), 1.
- Erlingsson, Ú., Pihur, V., & Korolova, A. (2014). Rappor: Randomized aggregatable privacy-preserving ordinal response. In *CCS*, pp. 1054–1067.
- Errouissi, R., Cardenas-Barrera, J., Meng, J., Castillo-Guerra, E., Gong, X., & Chang, L. (2015). Bootstrap prediction interval estimation for wind speed forecasting. In *2015 IEEE ECCE*, pp. 1919–1924. IEEE.
- Feige, I. (2019) <https://faculty.ai/blog/what-is-ai-safety/>.
- Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., & Venkatasubramanian, S. (2015). Certifying and removing disparate impact. In *KDD*, pp. 259–268.
- Feng, R., Yang, Y., Lyu, Y., Tan, C., Sun, Y., & Wang, C. (2019). Learning fair representations via an adversarial framework..
- Feng, S., Banerjee, R., & Choi, Y. (2012). Syntactic stylometry for deception detection. In *ACL*, pp. 171–175.
- Fredrikson, M., Jha, S., & Ristenpart, T. (2015). Model inversion attacks that exploit confidence information and basic countermeasures. In *CCS*, pp. 1322–1333.
- Fredrikson, M., Lantz, E., Jha, S., Lin, S., Page, D., & Ristenpart, T. (2014). Privacy in pharmacogenetics: An end-to-end case study of personalized warfarin dosing. In *USENIX Security*, pp. 17–32.
- Fu, R., Huang, Y., & Singh, P. V. (2020). AI and algorithmic bias: Source, detection, mitigation and implications..
- Galhotra, S., Brun, Y., & Meliou, A. (2017). Fairness testing: testing software for discrimination. In *ESEC/FSE*, pp. 498–510.
- Garg, N., Schiebinger, L., Jurafsky, D., & Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *PNAS*, 115(16), E3635–E3644.
- Ge, S., Cheng, L., & Liu, H. (2021). Improving cyberbullying detection with user interaction. In *Proceedings of the Web Conference 2021*, pp. 496–506.
- Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. In *STOC*, pp. 169–178.
- Gershgorn, D. (2019) <https://www.popsci.com/nsas-skynet-might-not-be-able-to-tell-what-makes-terrorist/>.

- Getoor, L. (2019). Responsible data science. In *Big Data*, pp. 1–1. IEEE.
- Geyik, S. C., Ambler, S., & Kenthapadi, K. (2019). Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In *KDD*, pp. 2221–2231.
- Gilad-Bachrach, R., Dowlin, N., Laine, K., Lauter, K., Naehrig, M., & Wernsing, J. (2016). Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *ICML*, pp. 201–210.
- Goel, N., Yaghini, M., & Faltings, B. (2018). Non-discriminatory machine learning through convex fairness criteria. In *AIES*, pp. 116–116.
- Goldwasser, S., & Micali, S. (1984). Probabilistic encryption. *JCSS*, 28(2), 270–299.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *NeurIPS*, pp. 2672–2680.
- Grath, R. M., Costabello, L., Van, C. L., Sweeney, P., Kamiab, F., Shen, Z., & Lecue, F. (2018). Interpretable credit application predictions with counterfactual explanations..
- Grover, V. (2020) <https://www.martechadvisor.com/articles/data-governance/data-dignity-for-better-data-privacy/>.
- Guess, A., Nagler, J., & Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on facebook. *Science Advances*, 5(1), eaau4586.
- Guo, R., Cheng, L., Li, J., Hahn, P. R., & Liu, H. (2020). A survey of learning causality with data: Problems and methods. *CSUR*, 53(4), 1–37.
- Guynn, J. (2015) <https://www.usatoday.com/story/tech/2015/07/01/google-apologizes-after-photos-identify-black-people-as-gorillas/29567465/>.
- Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. In *NeurIPS*, pp. 3315–3323.
- Hart, V. (2019). Data dignity at radicalxchange - the art of research. <https://theartofresearch.org/data-dignity-at-radicalxchange/>.
- Hartigan, J. A., & Wong, M. A. (1979). Algorithm as 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (applied statistics)*, 28(1), 100–108.
- Hind, M., Houde, S., Martino, J., Mojsilovic, A., Piorkowski, D., Richards, J., & Varshney, K. R. (2020). Experiences with improving the transparency of AI models and services. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–8.
- Hirano, K., Imbens, G. W., & Ridder, G. (2003). Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4), 1161–1189.
- Hitaj, B., Ateniese, G., & Perez-Cruz, F. (2017). Deep models under the gan: information leakage from collaborative deep learning. In *CCS*, pp. 603–618.
- Ho, S., Xie, M., Tang, L., Xu, K., & Goh, T. (2001). Neural network modeling with confidence bounds: a case study on the solder paste deposition process. *IEEE Transactions on Electronics Packaging Manufacturing*, 24(4), 323–332.

- Holzinger, A., Langs, G., Denk, H., Zatloukal, K., & Müller, H. (2019). Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(4), e1312.
- Homer, N., Szelling, S., Redman, M., Duggan, D., Tembe, W., Muehling, J., Pearson, J. V., Stephan, D. A., Nelson, S. F., & Craig, D. W. (2008). Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays. *PLoS Genet*, 4(8), e1000167.
- Horwitz, J. (2021). Facebook algorithm shows gender bias in job ads, study finds. <https://www.wsj.com/articles/facebook-shows-men-and-women-different-job-ads-study-finds-11617969600>.
- Hosen, M. A., Khosravi, A., Nahavandi, S., & Creighton, D. (2014). Improving the quality of prediction intervals through optimal aggregation. *IEEE Transactions on Industrial Electronics*, 62(7), 4420–4429.
- Hu, L., & Chen, Y. (2018). A short-term intervention for long-term fairness in the labor market. In *The Web Conference*, pp. 1389–1398.
- Huang, Q., Singh, V. K., & Atrey, P. K. (2014). Cyberbullying detection using social and textual analysis. In *SAM*, pp. 3–6.
- Jagadish, V. H. (2019). Responsible data science. In *2019 WSDM*, pp. 1–1. ACM.
- Janicka, M., Pszona, M., & Wawer, A. (2019). Cross-domain failures of fake news detection. *Computación y Sistemas*, 23(3).
- Jiang, H., & Nachum, O. (2020). Identifying and correcting label bias in machine learning. In *AISTATS*, pp. 702–712.
- Joachims, T., Swaminathan, A., & Schnabel, T. (2017). Unbiased learning-to-rank with biased feedback. In *WSDM*, pp. 781–789.
- Johnson, K. D., Foster, D. P., & Stine, R. A. (2016). Impartial predictive modeling: Ensuring fairness in arbitrary models..
- Kabir, H. D., Khosravi, A., Hosen, M. A., & Nahavandi, S. (2018). Neural network-based uncertainty quantification: A survey of methodologies and applications. *IEEE Access*, 6, 36218–36234.
- Kairouz, P., Oh, S., & Viswanath, P. (2014). Extremal mechanisms for local differential privacy. In *NeurIPS*, pp. 2879–2887.
- Kamishima, T., Akaho, S., Asoh, H., & Sakuma, J. (2012). Fairness-aware classifier with prejudice remover regularizer. In *ECML PKDD*, pp. 35–50. Springer.
- Karimi, F., Génois, M., Wagner, C., Singer, P., & Strohmaier, M. (2018). Homophily influences ranking of minorities in social networks. *Scientific Reports*, 8(1), 1–12.
- Kasiviswanathan, K., & Sudheer, K. (2016). Comparison of methods used for quantifying prediction interval in artificial neural network hydrologic models. *Modeling Earth Systems and Environment*, 2(1), 22.
- Khosravi, A., Nahavandi, S., Srinivasan, D., & Khosravi, R. (2014). Constructing optimal prediction intervals by using neural networks and bootstrap method. *IEEE Transactions on Neural Networks and Learning Systems*, 26(8), 1810–1815.

- Kim, B., Khanna, R., & Koyejo, O. O. (2016). Examples are not enough, learn to criticize! criticism for interpretability. In *NeurIPS*, pp. 2280–2288.
- Knight, W. (2020). This company uses AI to outwit malicious AI. <https://www.wired.com/story/company-uses-ai-outwit-malicious-ai/>.
- Koh, P. W., & Liang, P. (2017). Understanding black-box predictions via influence functions..
- Kontostathis, A., Reynolds, K., Garron, A., & Edwards, L. (2013). Detecting cyberbullying: query terms and techniques. In *Web Science*, pp. 195–204.
- Kraft, E. M., & Wang, J. (2009). Effectiveness of cyber bullying prevention strategies: A study on students’ perspectives.. *International Journal of Cyber Criminology*, 3(2).
- Krasanakis, E., Spyromitros-Xioufis, E., Papadopoulos, S., & Kompatsiaris, Y. (2018). Adaptive sensitive reweighting to mitigate bias in fairness-aware classification. In *The Web Conference*, pp. 853–862.
- Kusner, M. J., Loftus, J., Russell, C., & Silva, R. (2017). Counterfactual fairness. In *NeurIPS*, pp. 4066–4076.
- Lambrecht, A., & Tucker, C. (2019). Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads. *Management Science*, 65(7), 2966–2981.
- Lange, J. E., Blackman, K. O., & Johnson, M. B. (2001). Speed violation survey of the new jersey turnpike: Final report..
- Lee, N. T., Resnick, P., & Barton, G. (2019). Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms..
- Legg, S., Hutter, M., et al. (2007). A collection of definitions of intelligence. *FAIA*, 157, 17.
- Leslie, D. (2019). Understanding artificial intelligence ethics and safety..
- Li, J., Cheng, K., Wang, S., Morstatter, F., Trevino, R. P., Tang, J., & Liu, H. (2017). Feature selection: A data perspective. *CSUR*, 50(6), 1–45.
- Lipton, Z. C. (2018). The mythos of model interpretability. *Queue*, 16(3), 31–57.
- List, J. A., & Momeni, F. (2018). When corporate social responsibility backfires: Theory and evidence from a natural field experiment. Tech. rep., National Bureau of Economic Research.
- Liu, J., Juuti, M., Lu, Y., & Asokan, N. (2017). Oblivious neural network predictions via minionn transformations. In *CCS*, pp. 619–631.
- Liu, L. T., Dean, S., Rolf, E., Simchowitz, M., & Hardt, M. (2018). Delayed impact of fair machine learning..
- Liu, S., Kailkhura, B., Loveland, D., & Han, Y. (2019). Generative counterfactual introspection for explainable deep learning..
- Loftus, J. R., Russell, C., Kusner, M. J., & Silva, R. (2018). Causal reasoning for algorithmic fairness..

- Lu, T., & Viljanen, M. (2009). Prediction of indoor temperature and relative humidity using neural network models: model comparison. *Neural Computing and Applications*, 18(4), 345.
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *NeurIPS*, pp. 4765–4774.
- Maaten, L. v. d., & Hinton, G. (2008). Visualizing data using t-sne. *JMLR*, 9(Nov), 2579–2605.
- Makri, E., Rotaru, D., Smart, N. P., & Vercauteren, F. (2019). EPIC: efficient private image classification (or: Learning from the masters). In *CT-RSA*, pp. 473–492. Springer.
- Malekzadeh, M., Clegg, R. G., Cavallaro, A., & Haddadi, H. (2018). Protecting sensory data against sensitive inferences. In *Proceedings of the 1st Workshop on Privacy by Design in Distributed Systems*, pp. 1–6.
- Malekzadeh, M., Clegg, R. G., Cavallaro, A., & Haddadi, H. (2019). Mobile sensor data anonymization..
- Malekzadeh, M., Clegg, R. G., Cavallaro, A., & Haddadi, H. (2020). Privacy and utility preserving sensor-data transformations..
- Marín, L. G., Valencia, F., & Sáez, D. (2016). Prediction interval based on type-2 fuzzy systems for wind power generation and loads in microgrid control design. In *2016 FUZZ-IEEE*, pp. 328–335. IEEE.
- Marr, B. (2021) <https://bernardmarr.com/how-much-data-do-we-create-every-day-the-mind-blowing-stats-everyone-should-read/>.
- McConnell, E. H., & Scheidegger, A. R. (2001). Race and speeding citations: Comparing speeding citations issued by air traffic officers with those issued by ground traffic officers. In *ACJS, Washington, DC*.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2019). A survey on bias and fairness in machine learning..
- Menon, A. K., & Williamson, R. C. (2018). The cost of fairness in binary classification. In *FAT**, pp. 107–118.
- Miller C, C. R. (2019). People, power and technology: The tech workers’ view. https://doteveryone.org.uk/wp-content/uploads/2019/04/PeoplePowerTech_Doteveryone_May2019.pdf.
- Mills, S., Baltassis, E., Santinelli, M., Carlisi, C., Durantoni, S., & Gallego, A. (2020). Six steps to bridge the responsible AI gap. <https://www.bcg.com/publications/2020/six-steps-for-socially-responsible-artificial-intelligence>.
- Mirshghallah, F., Taram, M., Vepakomma, P., Singh, A., Raskar, R., & Esmaeilzadeh, H. (2020). Privacy in deep learning: A survey..
- Mishna, F., Khoury-Kassabri, M., Gadalla, T., & Daciuk, J. (2012). Risk factors for involvement in cyber bullying: Victims, bullies and bully-victims. *Children and Youth Services Review*, 34(1), 63–70.

- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. In *FAT**, pp. 220–229.
- Mokhberian, N., Abeliuk, A., Cummings, P., & Lerman, K. (2020). Moral framing and ideological bias of news. In *SocInfo*, pp. 206–219. Springer.
- Molnar, C. (2020). *Interpretable Machine Learning*. Lulu. com.
- Moraffah, R., Karami, M., Guo, R., Raglin, A., & Liu, H. (2020). Causal interpretability for machine learning-problems, methods and evaluation. *SIGKDD Explorations*, 22(1), 18–33.
- Mordvintsev, A., Olah, C., & Tyka, M. (2015). Inceptionism: Going deeper into neural networks..
- Nabi, R., & Shpitser, I. (2018). Fair inference on outcomes. In *AAAI*, Vol. 2018, p. 1931. NIH Public Access.
- Narayanan, A., & Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. In *2008 IEEE sp*, pp. 111–125. IEEE.
- Narendra, T., Sankaran, A., Vijaykeerthy, D., & Mani, S. (2018). Explaining deep learning models using causal inference..
- Nguyen, N. P., Yan, G., Thai, M. T., & Eidenbenz, S. (2012). Containment of misinformation spread in online social networks. In *Web Science*, pp. 213–222.
- Ott, M., Choi, Y., Cardie, C., & Hancock, J. T. (2011). Finding deceptive opinion spam by any stretch of the imagination..
- Pacepa, I. M., & Rychlak, R. J. (2013). *Disinformation: Former Spy Chief Reveals Secret Strategy for Undermining Freedom, Attacking Religion, and Promoting Terrorism*. Wnd Books.
- Papernot, N., McDaniel, P., Sinha, A., & Wellman, M. (2016a). Towards the science of security and privacy in machine learning..
- Papernot, N., McDaniel, P., Wu, X., Jha, S., & Swami, A. (2016b). Distillation as a defense to adversarial perturbations against deep neural networks. In *IEEE Symposium on SP*, pp. 582–597. IEEE.
- Parafita, Á., & Vitrià, J. (2019). Explaining visual models by causal attribution..
- Pearl, J. (2009). *Causality*. Cambridge university press.
- Pearl, J. (2018). Theoretical impediments to machine learning with seven sparks from the causal revolution..
- Pearl, J. (2019). The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM*, 62(3), 54–60.
- Peters, J., Janzing, D., & Schölkopf, B. (2017). *Elements of causal inference*. The MIT Press.
- Pfohl, S. R., Foryciarz, A., & Shah, N. H. (2020). An empirical characterization of fair machine learning for clinical risk prediction..

- Pinceti, A., Kosut, O., & Sankar, L. (2019). Data-driven generation of synthetic load datasets preserving spatio-temporal features. In *2019 IEEE PESGM*, pp. 1–5. IEEE.
- Qian, F., Gong, C., Sharma, K., & Liu, Y. (2018). Neural user response generator: Fake news detection with collective user intelligence.. In *IJCAI*, Vol. 18, pp. 3834–3840.
- Quan, H., Srinivasan, D., & Khosravi, A. (2014). Particle swarm optimization for construction of neural network-based prediction intervals. *Neurocomputing*, *127*, 172–180.
- Rajabi, Z., Shehu, A., & Purohit, H. (2019). User behavior modelling for fake information mitigation on social web. In *SBP-BRiMS*, pp. 234–244. Springer.
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?” explaining the predictions of any classifier. In *KDD*, pp. 1135–1144.
- Rivero, N. (2020). Google showed us the danger of letting corporations lead AI research. <https://qz.com/1945293/the-dangers-of-letting-google-lead-ai-research/>.
- Robnik-Šikonja, M., & Bohanec, M. (2018). Perturbation-based explanations of prediction models. In *Human and Machine Learning*, pp. 159–175. Springer.
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, *70*(1), 41–55.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, *66*(5), 688.
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, *1*(5), 206–215.
- Salawu, S., He, Y., & Lumsden, J. (2017). Approaches to automated detection of cyberbullying: A survey..
- Sattigeri, P., Hoffman, S. C., Chenthamarakshan, V., & Varshney, K. R. (2019). Fairness gan: Generating datasets with fairness properties using a generative adversarial network. *IBM Journal of Research and Development*, *63*(4/5), 3–1.
- Schaefer, R. F., Boche, H., Khisti, A., & Poor, H. V. (2017). *Information Theoretic Security and Privacy of Information Systems*. Cambridge University Press.
- Schölkopf, B. (2019). Causality for machine learning..
- Schwab, K. (2021). ‘this is bigger than just timnit’: How google tried to silence a critic and ignited a movement. <https://www.fastcompany.com/90608471/timnit-gebru-google-ai-ethics-equitable-tech-movement>.
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *FAT**, pp. 59–68.
- Shaul, B. (2015). Honestly looks to combat cyberbullying on ios, android. <https://www.adweek.com/performance-marketing/honestly-looks-to-combat-cyberbullying-on-ios-android/>.
- Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In *2017 IEEE Symposium on SP*, pp. 3–18. IEEE.

- Shu, K., Cui, L., Wang, S., Lee, D., & Liu, H. (2019). dEFEND: Explainable fake news detection. In *KDD*, pp. 395–405.
- Shu, K., Mahudeswaran, D., Wang, S., Lee, D., & Liu, H. (2020). Fakenewsnet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big Data*, 8(3), 171–188.
- Shu, K., Wang, S., & Liu, H. (2018). Understanding user profiles on social media for fake news detection. In *2018 IEEE MIPR*, pp. 430–435. IEEE.
- Shu, K., Wang, S., & Liu, H. (2019). Beyond news contents: The role of social context for fake news detection. In *WSDM*, pp. 312–320.
- Shu, K., Zheng, G., Li, Y., Mukherjee, S., Awadallah, A. H., Ruston, S., & Liu, H. (2020). Leveraging multi-source weak social supervision for early detection of fake news..
- Siegel, E. (2020) <https://www.kdnuggets.com/how-machine-learning-works-for-social-good.html/>.
- Silva, Y. N., Rich, C., & Hall, D. (2016). Bullyblocker: Towards the identification of cyberbullying in social networking sites. In *ASONAM*, pp. 1377–1379. IEEE.
- Simonyan, K., Vedaldi, A., & Zisserman, A. (2013). Deep inside convolutional networks: Visualising image classification models and saliency maps..
- Singh, R., Vatsa, M., & Ratha, N. (2021). Trustworthy AI. In *8th ACM IKDD CODS and 26th COMAD*, pp. 449–453.
- Slack, D., Hilgard, S., Jia, E., Singh, S., & Lakkaraju, H. (2020). Fooling lime and shap: Adversarial attacks on post hoc explanation methods. In *AIES*, pp. 180–186.
- Smith, P. K., Mahdavi, J., Carvalho, M., Fisher, S., Russell, S., & Tippett, N. (2008). Cyberbullying: Its nature and impact in secondary school pupils. *Journal of Child Psychology and Psychiatry*, 49(4), 376–385.
- Ståhl, N., Falkman, G., Karlsson, A., & Mathiason, G. (2020). Evaluation of uncertainty quantification in deep learning. In *IPMU*, pp. 556–568. Springer.
- Stoica, A.-A., Riederer, C., & Chaintreau, A. (2018). Algorithmic glass ceiling in social networks: The effects of social recommendations on network diversity. In *The Web Conference*, pp. 923–932.
- Sühr, T., Hilgard, S., & Lakkaraju, H. (2020). Does fair ranking improve minority outcomes? understanding the interplay of human and algorithmic biases in online hiring..
- Suter, R., Miladinovic, D., Schölkopf, B., & Bauer, S. (2019). Robustly disentangled causal mechanisms: Validating deep representations for interventional robustness. In *ICML*, pp. 6056–6065. PMLR.
- Sweeney, L. (2002). k-anonymity: A model for protecting privacy. *IJUFKS*, 10(05), 557–570.
- Thiebes, S., Lins, S., & Sunyaev, A. (2020). Trustworthy artificial intelligence..
- Tjoa, E., & Guan, C. (2020). A survey on explainable artificial intelligence (xai): Toward medical xai..
- Tong, A., Du, D.-Z., & Wu, W. (2018). On misinformation containment in online social networks. In *NeurIPS*, pp. 341–351.

- Tramèr, F., Zhang, F., Juels, A., Reiter, M. K., & Ristenpart, T. (2016). Stealing machine learning models via prediction apis. In *USENIX Security*, pp. 601–618.
- Ungar, L. H., De Veaux, R. D., & Rosengarten, E. (1996). Estimating prediction intervals for artificial neural networks. In *Proc. of the 9th Yale WALIS*.
- Varodayan, D., & Khisti, A. (2011). Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage. In *ICASSP*, pp. 1932–1935. IEEE.
- Varshney, K. R. (2019). Trustworthy machine learning and artificial intelligence. *XRDS: Crossroads, The ACM Magazine for Students*, 25(3), 26–29.
- Varshney, K. R., & Alemzadeh, H. (2017). On the safety of machine learning: Cyber-physical systems, decision sciences, and data products. *Big Data*, 5(3), 246–255.
- Vishwamitra, N., Zhang, X., Tong, J., Hu, H., Luo, F., Kowalski, R., & Mazer, J. (2017). Mcdefender: Toward effective cyberbullying defense in mobile online social networks. In *CODASPY*, pp. 37–42.
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151.
- Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the gdpr. *Harv. JL & Tech.*, 31, 841.
- Wan, C., Xu, Z., Pinson, P., Dong, Z. Y., & Wong, K. P. (2013). Probabilistic forecasting of wind power generation using extreme learning machine. *IEEE Transactions on Power Systems*, 29(3), 1033–1044.
- Wang, J., Zhang, J., Bao, W., Zhu, X., Cao, B., & Yu, P. S. (2018a). Not just privacy: Improving performance of private deep learning in mobile cloud. In *KDD*, pp. 2407–2416.
- Wang, X., Golbandi, N., Bendersky, M., Metzler, D., & Najork, M. (2018b). Position bias estimation for unbiased learning to rank in personal search. In *WSDM*, pp. 610–618.
- Wang, Y., Liang, D., Charlin, L., & Blei, D. M. (2018c). The deconfounded recommender: A causal inference approach to recommendation..
- Wei, D., Ramamurthy, K. N., & Calmon, F. d. P. (2020). Optimized score transformation for fair classification. In *AISTATS*.
- Weller, A. (2017). Challenges for transparency..
- Wikipedia (2021a) [https://en.wikipedia.org/w/index.php?title=Robustness_\(computer_science\)&oldid=1009774103](https://en.wikipedia.org/w/index.php?title=Robustness_(computer_science)&oldid=1009774103). Page Version ID: 1009774103.
- Wikipedia (2021b) https://en.wikipedia.org/w/index.php?title=Facebook%E2%80%93Cambridge_Analytica_data_scandal&oldid=1035933869. Page Version ID: 1035933869.
- Wikipedia (2021c) https://en.wikipedia.org/w/index.php?title=Cambridge_Analytica&oldid=1036071985. Page Version ID: 1036071985.
- Wold, S., Esbensen, K., & Geladi, P. (1987). Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3), 37–52.

- Xu, D., Wu, Y., Yuan, S., Zhang, L., & Wu, X. (2019). Achieving causal fairness through generative adversarial networks.. In *IJCAI*, pp. 1452–1458.
- Xu, J.-M., Jun, K.-S., Zhu, X., & Bellmore, A. (2012). Learning from bullying traces in social media. In *NAACL HLT*, pp. 656–666. ACL.
- Yang, Z., & Feng, J. (2020). A causal inference method for reducing gender bias in word embedding relations.. In *AAAI*, pp. 9434–9441.
- Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., & Hovy, E. (2016). Hierarchical attention networks for document classification. In *NAACL HLT*, pp. 1480–1489.
- Yao, L., Chu, Z., Li, S., Li, Y., Gao, J., & Zhang, A. (2020). A survey on causal inference..
- Yeo, C. (2020). What is transparency in AI?. <https://medium.com/fair-bytes/what-is-transparency-in-ai-bd08b2e901ac>.
- Young, M., Magassa, L., & Friedman, B. (2019). Toward inclusive tech policy design: a method for underrepresented voices to strengthen tech policy documents. *Ethics and Information Technology*, *21*(2), 89–103.
- Yuan, X., He, P., Zhu, Q., & Li, X. (2019). Adversarial examples: Attacks and defenses for deep learning. *IEEE Transactions on Neural Networks and Learning Systems*, *30*(9), 2805–2824.
- Zhang, G., Wu, Y., Wong, K. P., Xu, Z., Dong, Z. Y., & Iu, H. H.-C. (2014). An advanced approach for construction of optimal wind power prediction intervals. *IEEE Transactions on Power Systems*, *30*(5), 2706–2715.
- Zhang, L., & Wu, X. (2017). Anti-discrimination learning: a causal modeling-based framework. *JDSA*, *4*(1), 1–16.
- Zhang, L., Wu, Y., & Wu, X. (2016). A causal framework for discovering and removing direct and indirect discrimination..
- Zio, E. (2006). A study of the bootstrap method for estimating the accuracy of artificial neural networks in predicting nuclear transient processes. *IEEE Transactions on Nuclear Science*, *53*(3), 1460–1478.