

A step toward building a unified framework for managing AI bias

Saadia Afzal Rana¹, Zati Hakim Azizul¹ and Ali Afzal Awan²

¹ Department of Artificial Intelligence, Faculty of Computer Science and Information Technology, Universiti Malaya, Kuala Lumpur, Malaysia

² National University of Science and Technology, Islamabad, Pakistan

ABSTRACT

Integrating artificial intelligence (AI) has transformed living standards. However, AI's efforts are being thwarted by concerns about the rise of biases and unfairness. The problem advocates strongly for a strategy for tackling potential biases. This article thoroughly evaluates existing knowledge to enhance fairness management, which will serve as a foundation for creating a unified framework to address any bias and its subsequent mitigation method throughout the AI development pipeline. We map the software development life cycle (SDLC), machine learning life cycle (MLLC) and cross industry standard process for data mining (CRISP-DM) together to have a general understanding of how phases in these development processes are related to each other. The map should benefit researchers from multiple technical backgrounds. Biases are categorised into three distinct classes; pre-existing, technical and emergent bias, and subsequently, three mitigation strategies; conceptual, empirical and technical, along with fairness management approaches; fairness sampling, learning and certification. The recommended practices for debias and overcoming challenges encountered further set directions for successfully establishing a unified framework.

Subjects Artificial Intelligence, Data Mining and Machine Learning, Software Engineering

Keywords Algorithmic bias, Fairness management, Bias mitigation strategy, Data-driven AI system, Fairness in data mining

INTRODUCTION

Data-driven decision-making applications have been deployed in vital areas such as finance (*Baum, 2017*), judiciary (*Angwin et al., 2016; Abebe et al., 2020*), employment (*Zhao et al., 2018b*), e-commerce (*Weith & Matt, 2022*), education (*Kim, Lee & Cho, 2022*), military intelligence (*Maathuis, 2022*) and health (*Gardner et al., 2022*). On one side, the potential of AI is widely recognised and appreciated. On the other side, there is significant uncertainty in managing negative consequences and subsequent challenges (*National Institute of Standards and Technology, 2022*), and a major hindrance to progress is a bias enrooted throughout the AI pipeline's development process (*Fahse, Huber & Giffen, 2021*).

The consequences of unwanted discriminatory and unfair behaviours can be detrimental (*Pedreshi, Ruggieri & Turini, 2008*), adversely affecting human rights (*Mehrabi et al., 2021*), university admissions (*Bickel, Hammel & O'Connell, 1975*), profit and revenue (*Mikians et al., 2012*) and facing legal risks (*Pedreshi, Ruggieri & Turini, 2008; Northeastern Global News, 2020; Romei & Ruggieri, 2014*). "Bias has existed since the dawn of society" (*Ntoutsis et al., 2020*). However, AI-based decision-making is criticised for introducing different types of biases. The ever-growing worries demand AI-based systems

Submitted 7 February 2023
Accepted 11 September 2023
Published 26 October 2023

Corresponding author
Zati Hakim Azizul, zati@um.edu.my

Academic editor
Wasif Afzal

Additional Information and
Declarations can be found on
page 19

DOI 10.7717/peerj-cs.1630

© Copyright
2023 Rana et al.

Distributed under
Creative Commons CC-BY 4.0

OPEN ACCESS

to rebuild technical strategies to integrate fairness as a core component of its infrastructure. Fairness is the absence of bias or discrimination. An algorithm that makes biased decisions against specific people is considered unfair (Fenwick & Molnar, 2022). The ethical implications of systems that affect individuals' lives have sparked concerns regarding the need for fair and impartial decision-making (Pethig & Kroenung, 2022; Hildebrandt, 2021). Consequently, extensive research has been conducted to address issues of bias and unfairness, while also considering the limitations imposed by corporate policies, legal frameworks (Zuiderveen Borgesius, 2018), societal norms, and ethical responsibilities (Hobson et al., 2021; Gupta, Parra & Dennehy, 2021).

There is still a lack of an organised strategy for tackling potential biases (Richardson et al., 2021). When looking for the origin of bias in AI decision-making, it is prevalent that the issue is either data or algorithms, and the root cause is humans. Humans transmit cognitive bias while creating/generating data or designing algorithms (National Institute of Standards and Technology, 2022; Fenwick & Molnar, 2022). Thoroughly evaluating existing literature is partial to this work in learning fairness management towards proposing a unified framework to address prejudice and its subsequent mitigation method.

In order to gain a comprehensive understanding, we frame our research by mapping the software development life cycle (SDLC), machine learning life cycle (MLLC) and cross industry standard process for data mining (CRISP-DM) process model to have a general understanding of how phases in this development process are related to each other. We categorise bias into three classes; pre-existing, technical and emergent bias and, subsequently, three mitigation strategies; conceptual, empirical and technical, along with fairness management approaches; fairness sampling, learning and certification. We discuss them in light of their occurrence at a specific phase of the development cycle. The recommended practices to avoid/mitigate biases and the challenges encountered in the course of action to address them are discussed in the later part.

RATIONALE

A foundation must be established to accommodate all perplexing features, and it is necessary to study various facets of bias and fairness from the perspective of software engineering integration in AI. This study maps together the software development life cycle (SDLC), machine learning life cycle (MLLC), and cross industry standard process for data mining (CRISP-DM) processes. The mapping provides a general understanding of how phases in these software engineering (SE) and AI development processes relate and can benefit from SE's best practices within AI. The proposed framework aims to detect, identify, and localize biases on the spot and prevent them in the future by comprehending their core causes. The framework handles bias as a defect management process in software engineering.

THE AUDIENCE IT IS INTENDED FOR

We believe we are the first to present this innovative framework. The framework can help ML researchers, ML engineers, data analysts, data scientists, software engineers, and software architects develop superior versions of software applications with higher

accuracy, better defect tracking, swifter control test timings and faster time-to-market release.

SURVEY/SEARCH METHODOLOGY SECTION

Integrating bottom-up and top-down research methodologies was used to gather publications on bias in AI/ML applications. Each co-author gathered pertinent material and added it to a shared repository. The keywords for search are optimised in each search with an expectation to shortlist articles containing exact information. The three core domains of artificial intelligence, machine learning, and software engineering were the focus of the search, as shown in Fig. 1. Furthermore, 'AND' and 'OR' string operators were used along with double quotation marks to further narrow down to search fairness in AI/ML.

Interestingly, each survey introduced some new types of bias or fairness definitions. Only articles that detail bias in data, algorithms, assessment tools, fairness management approaches, bias detection, identification, mitigation strategy, fairness matrices, AI/ML/SE development cycles, datasets characteristics, AI ethics and principles issues were shortlisted during the inclusion criteria. We are further guided to our goal by sections from academic books, keynote addresses by renowned speakers (some of them are J. Buolamwini, T. Gebru (*Buolamwini & Gebru, 2018*) and C. O'Neil (*O'Neil, 2017*)), and various advisories published. Non-tabular data, domain/language-dependent technology, and the absence of experimental results are among the exclusion factors that reduce the number of articles we can find from 254 to 72. Due to a lack of mapping and tackling a variety of interlinked biases, we faced many challenges, *i.e.*, uncoordinated development efforts, inefficient use of resources, poor quality of ML models, limited scalability, amplification of biases, unfair decision-making, a lack of diversity, a loss of trust, and ethical issues. And in short found difficulties in integrating ML models into software. Addressing these issues requires a multifaceted approach that involves addressing biases at multiple levels, including in the data, the algorithms and the decision-making processes. It also requires a commitment to ethical and responsible development and use of technologies, especially when all development cycles should be synchronised together.

REVIEW OF LITERATURE

The literature review on the subject is dense, with many forms of biases and mitigation strategies, most of which are interlinked. It is essential to be aware of potential biases in the AI pipeline and the appropriate mitigation strategies to address undesirable effects (*Bailey et al., 2019*).

Realistic case study of AI bias

The previous 10 years have seen AI's widespread adoption and popularity, allowing it to permeate practically every aspect of our daily lives. However, safety and fairness concerns have prompted practitioners to prioritise them while designing and engineering AI-based applications (*Chouldechova et al., 2018; Howard & Borenstein, 2018; Osoba & Welser, 2017*). Researchers have enumerated a few applications that, because of biases, have a

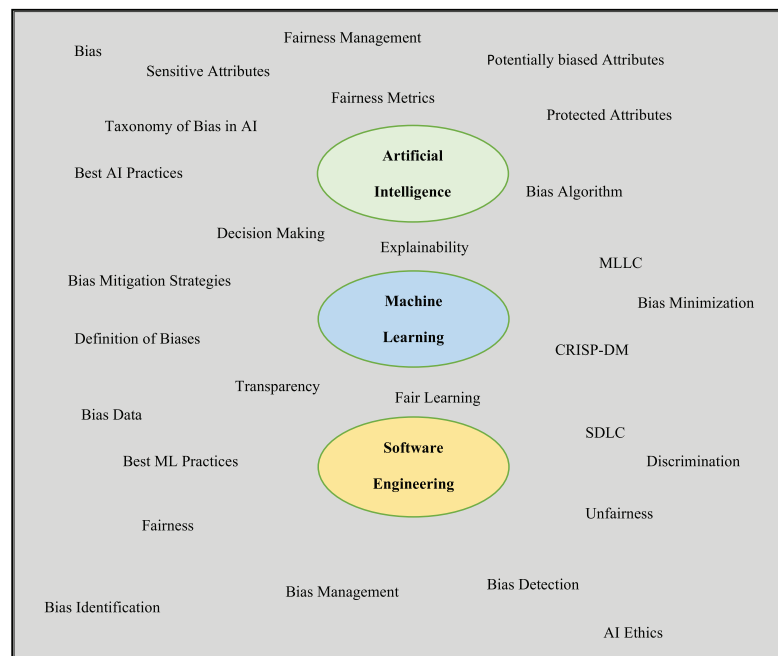


Figure 1 Search words across three predefined domains-AI, ML & SE.

Full-size  DOI: 10.7717/peerj-cs.1630/fig-1

detrimental impact on people's lives, such as biometrics apps, autonomous vehicles, AI chat-bots, robotic systems, employment matching, medical aid, and systems for children's welfare.

The US judiciary uses the correctional offender management profiling for alternative sanctions (COMPAS) (*Pagano et al., 2023*), an AI-based tool, to identify offenders more likely to commit crimes again. However, the Pulitzer Prize-winning nonprofit news organisation ProPublica (*Angwin et al., 2016*) discovered that COMPAS was racially prejudiced and that black criminals were at high risk (*Dressel & Farid, 2018*). PredPol, or predictive policing, is an AI-enabled software to predict the next crime location based on the number of arrests and frequency of calls to the police regarding different offences. PredPol was criticised for its prejudiced behaviour as it targeted racial minorities (*Benbouzid, 2019*).

An AI algorithm known as the Amazon recruiting engine was developed to evaluate the resumes of job applicants applying to Amazon and then shortlist eligible candidates for further interviews and consideration. However, it turned out that the Amazon algorithm was discriminatory regarding hiring women (*Hofeditz et al., 2022*). A labelling tool in Google Photos adds a label to a photo that corresponds to whatever is seen in the image. When it referred to images of a black software developer and his friend as gorillas, it was determined to be racist (*González Esteban & Calvo, 2022*). The well-known programme StyleGAN (*Karras, Laine & Aila, 2019*) automatically creates eerily realistic human faces and produces white faces more frequently than racial minorities.

Biases based on racial, gender, and other demographic factors restrict communities from using AI-automated technologies for regular tasks in the health and education sectors

(Egan et al., 1996; Brusseau, 2022). Organisations must know the different types of biases in their data/algorithm that can affect their machine learning (ML) models. Ultimately, identifying and mitigating biases that skewed or produced undesirable outcomes and impeded the advancements made by AI for the everyday person is more than necessary.

Assessment tools

A systematic effort has been made to provide appropriate tools for practitioners to adopt cutting-edge fairness strategies into their AI pipelines. Software toolkits and checklists are the two basic ways to ensure fairness (Richardson & Gilbert, 2021). Programming language functions that can be used to identify or lessen biases are known as toolkits. AI practitioners can employ checklists and detailed instructions by fairness specialists to ensure ethical consideration is incorporated across their pipelines (Bailey et al., 2019).

Fairness indicators and the What-If toolkit (WIT) are well-known tools that Google provides (Richardson et al., 2021). Fairness indicators are based on fairness metrics for binary and multiclass classifiers (Agarwal, Agarwal & Agarwal, 2022). The What-If tool is an interactive visual tool designed to examine, evaluate, and compare ML models (Richardson et al., 2021). UChicago's Aequitas is utilised to assess ML-based outcomes to identify various biases and make justified choices regarding the creation and implementation of such systems (Saleiro et al., 2018). IBM's AI Fairness 360 is a toolkit providing fairness detection and mitigation strategies (Bellamy et al., 2019). LinkedIn's fairness toolkit (LiFT) provides detection strategies for measuring fairness across various metrics (Brusseau, 2022). Microsoft's Fairlearn, ML Fairness Gym, Scikit's fairness tool, and PyMetrics Audit-AI are readily available tools to detect/mitigate (or both) bias and ensure fairness management. These assessment techniques expand practitioners' options for developing ethical products and maintaining stakeholder confidence (Dankwa-Mullan & Weeraratne, 2022).

Despite the availability of tools that incorporate explainable machine learning methods and fair algorithms, none of them currently offers a comprehensive set of guidelines to assist users in effectively addressing the diverse fairness concerns that arise at different stages of the machine learning decision pipeline (Pagano et al., 2023). Moreover the responsibility for identifying and mitigating bias and unfairness is often entirely placed on the developer, who may not possess sufficient expertise in addressing these challenges and cannot be solely responsible for the same (Schwartz et al., 2022). There is lack of consensus on what constitutes fairness as there is no single definition of fairness. Moreover, these tools have limited ability to detect complex forms of bias as they are designed to detect bias due to protected attributes only and casual fairness demands to investigate underlying relations between different attributes. Fairness assessment tools can themselves be biased. For example, a tool that is designed to detect bias in natural language processing models may be biased towards certain languages or dialects. Sometimes they can be time-consuming and expensive to use for organizations to adopt these tools, especially small or resource-constrained organizations. One more notable objection regarding them is that it can be difficult to interpret the results which make it difficult for organization to take

action to address any bias that is found. Majority of fairness assessment tools may not be able to identify the root cause of bias.

Fairness management research datasets

Several standard datasets are available to make research on bias and fairness easier. Each of them has sensitive or protected attributes that could be used to demonstrate unfair or biased treatment toward underprivileged groups or classes. [Table 1](#) summarises some of these datasets' characteristics. These datasets are used as a benchmark to evaluate and compare the effectiveness of bias detection and mitigation strategies.

BROAD SPECTRUM OF BIAS

Since bias has existed for as long as human civilisation, research on the topic is popular. The literature review, however, is full of various terminologies and theories that either overlap or are interlinked, and this conflict is enough to perplex researchers. In this section, we will try to throw light on a different aspect of bias in AI with a perspective to overcome confusion and increase researchers' understanding. This effort sets the groundwork for building a unified framework for fairness management.

Bias, discrimination and unfairness

The existence of bias, discrimination, and unfairness are related topics, but it is essential to distinguish them for further investigation/analysis ([Future Learn, 2013](#)). "Bias is the unfair inclination or prejudice in a judgement made by an AI system, either in favour of or against an individual or group" ([Fenwick & Molnar, 2022](#)). Discrimination can be considered a source of unfairness due to human prejudice and stereotyping based on sensitive attributes, which may happen intentionally or unintentionally. In contrast, bias can be considered a source of unfairness due to the data collection ([Luengo-Oroz et al., 2021](#)), sampling, and measurement ([Li & Chignell, 2022](#)). Discrimination is a difference in the treatment of individuals based on their membership in a group ([Ethics of AI, 2020](#)). Bias is a systematic difference in treating particular objects, people or groups compared to others. Unfairness is the presence of bias where we believe there should be no systematic difference ([Zhao et al., 2018b](#)). Bias is a statistical property, whereas fairness is generally an ethical issue.

Protected, sensitive and potentially biased attributes

Until we stop making it, machine bias will keep appearing everywhere: bias in, bias out ([Alelyani, 2021](#)). Protected attributes are characteristics that cannot be relied upon to make decisions and may be chosen based on the organisation's objectives or regulatory requirements. Similarly, sensitive attributes are characteristics of humans that may be given special consideration for social, ethical, legal or personal reasons ([Machine Learning Glossary: Fairness, \(n.d.\)](#)). In literature, these terms are used interchangeably and in place of one another. It is a potentially biased attribute if changing the value of an attribute through the alternation function has an impact on prediction, such as altering a male attribute to a female one or a black attribute to a white one ([Alelyani, 2021](#)). Any one of these attributes in the dataset necessitates special consideration. Sex, race, color, age,

Table 1 Few popular datasets, along with their characteristics.

Dataset name	Attributes	No. of records	Area
COMPAS (<i>Bellamy et al., 2019</i>)	Criminal histories, jail & prison times, demographics, COMPAS risk scores	18,610	Social
German credit (<i>Pagano et al., 2023</i>)	Housing status, personal status, amount, credit score, credit, sex	1,000	Financial
UCI adult (<i>Schwartz et al., 2022</i>)	Age, race, hours-per-week, marital status, occupation, education, sex, native country	58,842	Social
Diversity in faces (<i>National Institute of Standards and Technology, 2022</i>)	Age, pose, facial symmetry and contrast, craniofacial distances, gender, skin color, resolution along with diverse areas and ratios	1 million	Facial images
Communities and crime (<i>Dua & Graff, 2017</i>)	Crime & socio-economic data	1,994	Social
Winobias (<i>Rhue & Clark, 2020</i>)	Male or female stereotypical occupations	3,160	Coreference resolution
Recidivism in Juvenile justice (<i>Merler et al., 2019</i>)	Juvenile offenders' data and prison sentences	4,753	Social
Pilot parliaments benchmark (<i>Redmond, 2011</i>)	National parliaments data (e.g., gender and race)	1,270	Facial images

marital status, family, religion, sexual orientation, political opinion, pregnancy, physical or mental disability, career responsibilities, social origins, and national extraction are a few attributes examples (*Fair Work Ombudsman, (n.d.)*). However, it has been observed that groups/individuals still face discrimination through proxy attributes (*Chen et al., 2019*), even in the absence of some protected or sensitive traits. Proxy attributes correlate with protected or sensitive attributes, such as zip code (linked with race) (*Mehrabi et al., 2021*).

Trio-bias feedback loop

Data is the primary driving force behind most AI systems and algorithms, so they need data to be trained. As a result, the functioning of these algorithms and systems is closely tied to data availability. Underlying training data bias will manifest through an algorithm's predictions (trained on it). Furthermore, algorithms may reflect discriminatory behaviour based on specific design considerations even when the data is fair. The biased algorithm's output can then be incorporated into the existing system and impact users' choices, producing more biased data that can be used to train the new algorithm. *Figure 2* depicts the feedback loop between data biases, algorithmic biases, and user involvement. Humans are involved in data preparation and algorithm design; therefore, whether bias results from data or an algorithm, humans are the root cause.

Categorisation of biases

The human bias has more than 180 varieties, for example *Dabas (2021)*. Bias has been broadly divided into three major categories to accomplish fairness with a focus on detection and mitigation mechanisms and dispel confusion with many different types of bias (*Gan & Moussawi, 2022*). Bias categories include pre-existing, technical and emergent (*Friedman et al., 2013*).

A bias before the development of technology is referred to as pre-existing bias. This kind of bias has its roots in social structures and manifests itself in individual biases. The same is introduced into technology by people and organisations responsible for its development (*Curto et al., 2022*), whether explicitly or implicitly, consciously or unconsciously (*Friedman et al., 2013*). The AI literature identifies that the most typical biases are pre-

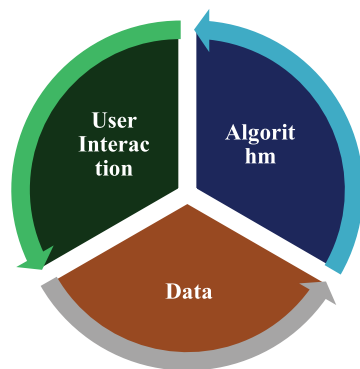


Figure 2 Trio bias feedback loop among data, algorithm and user.

Full-size DOI: 10.7717/peerj-cs.1630/fig-2

existing (*Gan & Moussawi, 2022*). A few examples of pre-existing bias are the wrong model of Microsoft bot (*Victor, 2016*), Duane Buck's murder (*O'Neil, 2017*), the stance on fairness (*O'Neil, 2017*), criminal justice models (*Bughin et al., 2018*), and the understanding of concepts and reasonability of CO₂ emissions (*Luengo-Oroz, 2019*).

Technical bias refers to concerns about a product's technological design, such as technical limitations or decisions (*Friedman et al., 2013*). Technical issues faced in prominent software like IMPACT, Tech, LSI-R, Kyle's job application and hiring algorithms are because of technical bias. Emergent bias emerges after the practical use of a design as a result of a shift in social awareness or cultural norms (*Friedman et al., 2013*). PredPol, St. George's model, COMPAS, and facial analysis tech were adversely affected due to emergent bias.

Origins of biases within the development cycle

The literature regarding the causes of biases and potential mitigation strategies is still scattered, and work on a systematic methodology for dealing with potential biases (*Nascimento et al., 2018*) and building well-established ML frameworks is in progress (*Suresh & Guttag, 2021; Ricardo, 2018; Silva & Kenney, 2019*). In this article, we attempt to overcome confusion and enhance understanding regarding the origin of different categories of bias during the development cycle by mapping SDLC, MLLC and CRISP-DM on a single scale, as shown in *Fig. 3*.

Professional designers or developers introduce pre-existing bias into the technical process, which is why it drags into the modelling phase of the CRISP-DM cycle (*Barton et al., 2019*). Technical bias, as opposed to pre-existing bias, results from how problems in the technical design are resolved. The design process and model evaluation contain many areas where technical bias can be identified (*Friedman et al., 2013; Ho & Beyan, 2020*). Pre-existing and technical biases occur prior to and within the technical development process. However, emergent bias appears during the actual use of the technical product after development (*Barton et al., 2019*). Emergent bias can be detected before deployment (during testing) and is often the most obvious category of bias (*Gan & Moussawi, 2022*). The same type of bias has several names in the literature. Therefore, biases were

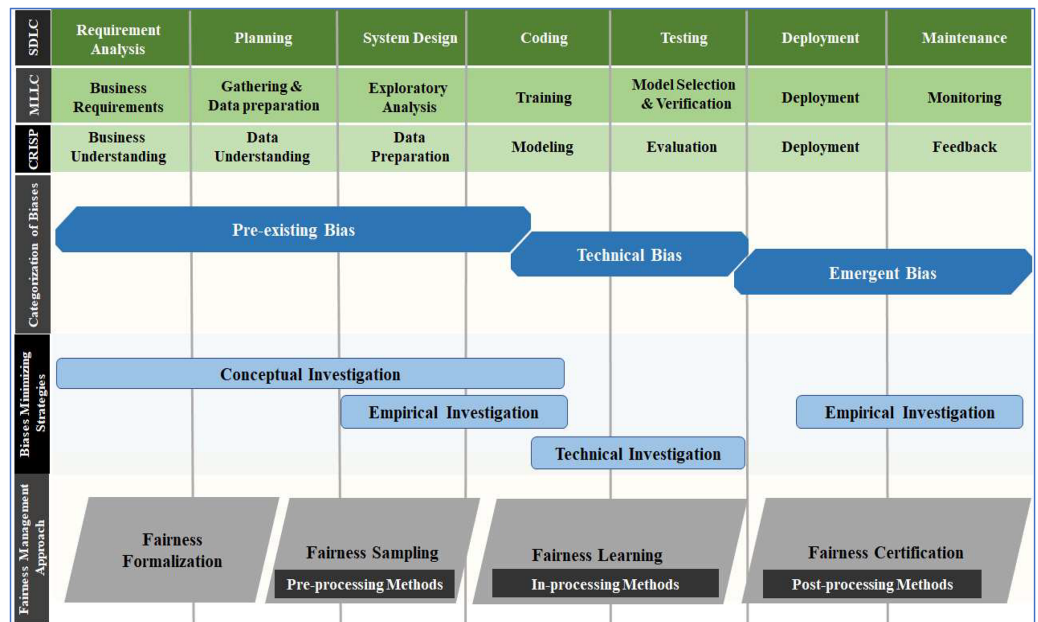


Figure 3 Mapping the SDLC, MLLC and CRISP-DM across different biases categories, minimizing strategies and fairness management. [Full-size !\[\]\(fcc3264021d438d9732560e78099f674_img.jpg\) DOI: 10.7717/peerj-cs.1630/fig-3](https://doi.org/10.7717/peerj-cs.1630/fig-3)

descriptively synthesised and characterised based on their origins (*Ho & Beyan, 2020; Fink, 2019*). Based on a thorough understanding of the development cycles, commonly encountered distinct types of biases were assigned to the phases based on their origin, as shown in [Fig. 4](#).

The bias type and subtype distribution across different phases will create an appropriate detection method (*Fahse, Huber & Giffen, 2021*). Equal opportunity, equalized odds, conditional demographic disparity, disparate impact, Euclidean distance, Mahalanobis distance, Manhattan distance, demographic disparity, and different tools and libraries are used for bias detection (*Garg & Si, 2021*).

Bias minimising strategies

Value sensitive design (VSD) is a methodology that can contribute to understanding and addressing issues of bias in AI systems (*Simon, Wong & Rieder, 2020*) and to promote transparency (*Cakir, 2020*) and ethical principles in AI systems (*Dexe et al., 2020*). It is a framework that provides recommendations for minimising or coping with various biases (*Gan & Moussawi, 2022; Friedman et al., 2013*), is adopted in this research as a potential strategy to minimise the biases associated with AI. In the VSD study, “value” is a general phrase that relates to what user values in life (*Atkinson, Bench-Capon & Bollegala, 2020*). This theoretically-based approach provides three possible solutions based on the types of investigations to prevent problems and advocate AI-specific value-oriented metrics that all stakeholders mutually agreed on.

Conceptual, empirical and technical are three different kinds of investigations. A conceptual investigation primarily focuses on analysing or prioritising different stakeholders’ values in the design and use of technology (*Barton et al., 2019; Floridi, 2010*).

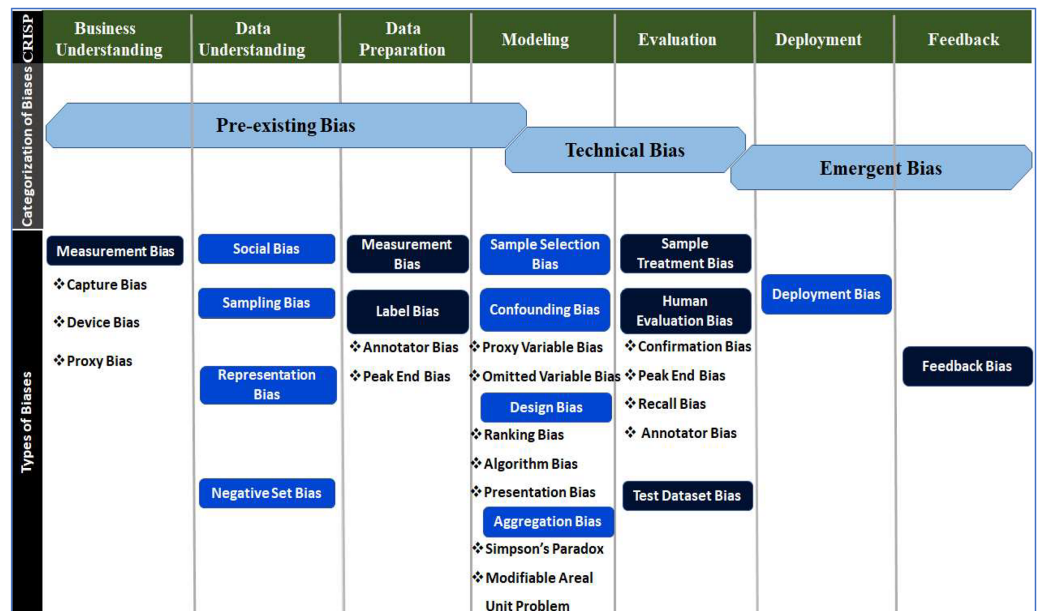



Figure 4 Distribution of biases w.r.t categories across phases of the CRISP-DM development cycle. Full-size  DOI: 10.7717/peerj-cs.1630/fig-4

Empirical investigation assesses a technical design's effectiveness using factors like how people react to technological products (Barton *et al.*, 2019; Floridi, 2010). It frequently entails observation and recording; quantitative and qualitative research methods are appropriate. Because of these characteristics, the phases are limited to data understanding (slightly), data preparation, modeling (to some extent), deployment and feedback.

The primary focus of the technical investigation is on technology while focusing on how technological characteristics and underlying mechanics promote or undermine human values and involves proactive system design to uphold values discovered in conceptual investigations (Gan & Moussawi, 2022; Friedman *et al.*, 2013; Ho & Beyan, 2020). VSD has revealed the relationships between different types of investigations and types of AI biases (Umbrello & Van de Poel, 2021). Researchers concluded by recognising the value of conceptual and empirical investigation for addressing pre-existing bias, investigations for minimising technical bias, and technical and empirical investigation for tackling emerging bias (Gan & Moussawi, 2022).

Bias mitigation methods

Several methods can mitigate a single bias, and multiple biases can be mitigated by a single method. The socio-technical approach comprising technical and nontechnical methods is widely adopted to counterattack the harmful effects of biased decisions (Fahse, Huber & Giffen, 2021). This approach mitigates bias and prevents it from recurring in the future. Table 2 depicts which method effectively mitigates different types of bias whenever it appears in different stages of the AI development cycle. The mitigation process is not executed during Feedback. From risk management's perspective, to evaluate an AI system's

Table 2 A socio-technical approach for bias mitigation across the CRISP-DM development cycle.

	Business understanding	Data understanding	Data preparation	Modelling	Evaluation	Deployment
Measurement bias	Team diversity, exchange with domain expert	Proxy estimation	Rapid prototyping			
Social bias			Learning fair representation, rapid prototyping, reweighting, optimized preprocessing, data massaging, disparate impact remover	Adversarial debiasing, multiple models, latent variable model, model interpretability equalized odds, prejudice remover		
Sampling bias				Resampling		Randomness
Representation bias	Team diversity	Data plotting, exchange with domain experts	Reweighting, data augmentation	Model interpretability		
Negative bias			Cross dataset generalization	Bag of words		
Label bias		Exchange with domain experts	Data massaging			
Sample selection bias			Reweighting			
Confounding bias						Randomness
Design Bias			Rapid prototyping	Exchange with domain experts, resampling, model interpretability, multitask learning		
Sample treatment bias				Resampling	Data augmentation	
Human evaluation bias				Resampling	Representative benchmark subgroup validity, data augmentation	
Test dataset bias			Data augmentation			
Deployment bias	Team diversity, consequences in context		Rapid prototyping			Monitoring plan, human supervision
Feedback bias						Human supervision, randomness

performance, ‘Fairness’ outclass any other vital measure, *i.e.*, dependability, efficiency, and accuracy (Barton *et al.*, 2019; Suri *et al.*, 2022).

Various strategies have been put forth by researchers and practitioners to address bias in AI. These strategies encompass data pre-processing, model selection, and post-processing decisions. However, each of these methods has its own set of limitations and difficulties, such as the scarcity of diverse and representative training data, the complexity of identifying and quantifying different forms of bias and the potential trade-offs between fairness and accuracy (Pagano *et al.*, 2023). Additionally, ethical considerations arise when determining which types of bias to prioritize and which groups should be given priority in the mitigation process (Pagano *et al.*, 2022). Bias mitigation methods can be complex, computationally expensive and can introduce new biases for example, a method that tries to balance the representation of different groups in a dataset may introduce a new bias in

favor of the majority group. They can be brittle such as that they can be sensitive to changes in the data or the model. This can make it difficult to ensure that the model remains fair over time. Another notable aspect that needs to be addressed are that there is a lack of understanding of the long-term effects of bias mitigation methods & standardized evaluation metrics. These methods can be opaque as it can be difficult to understand how they work, subsequently it can make it difficult to trust these methods, and to ensure that they are not introducing new biases. They might feel it difficult to adapt to new tasks, as result of it, they may not be effective for all machine learning models.

Despite these obstacles, the mitigation of bias in AI is of utmost importance to establish just and equitable systems that benefit everyone in society (*Balayn, Lofi & Houben, 2021*). Continuous research and development of mitigation techniques are crucial to overcome these challenges and ensure that AI systems are employed for the welfare of all individuals (*Huang et al., 2022*).

FAIRNESS MANAGEMENT APPROACH

Bias and fairness are two mutually exclusive aspects of reality. In the absence of a unified definition, an absence of bias or preference for individuals or groups based on their characteristics is generally regarded as ‘fairness’. It is necessary to do more than mitigate any bias detected to ensure that an AI system may be regarded as fair. Instead, “fairness-aware” system design should be encouraged (*Orphanou et al., 2021*). This approach incorporates “fairness” as a crucial design component from conception to deployment (sometimes extended to maintenance or upgradation). *Figure 5* shows the fairness management approach and several implementation methods across the CRISP-DM model. The following steps constitute the fairness management approach:

Fairness formalisation

Constraints, measures, specifications and criteria to ensure fairness is defined during the business understanding and data understanding phases (*Northeastern Global News, 2020*), as depicted in *Fig. 5*. In fact, at this stage, the benchmarks for auditing or evaluating fairness are stated.

Fairness sampling

Fairness sampling generally refers to preprocessing skewed data through different methods, such as oversampling (*Orphanou et al., 2021*). Issues with data, *i.e.*, inaccuracy, incompleteness, improperly labelling, too much/less, inconsistency, and silos, are addressed at this stage (*Walch, 2020*). Fair sampling is accomplished during the data understanding and data preparation phases.

Fairness learning

It is not always feasible to develop a fair model by eliminating the bias in the initial data before training. Designing a fair classifier that uses a fair algorithm is the solution in such scenarios (*Acharyya et al., 2020*). As a result, we can still use a biased dataset to train the model, and the fair algorithm still produces predictions through in-processing methods carried out during the modelling and evaluation phases (*Acharyya et al., 2020*).

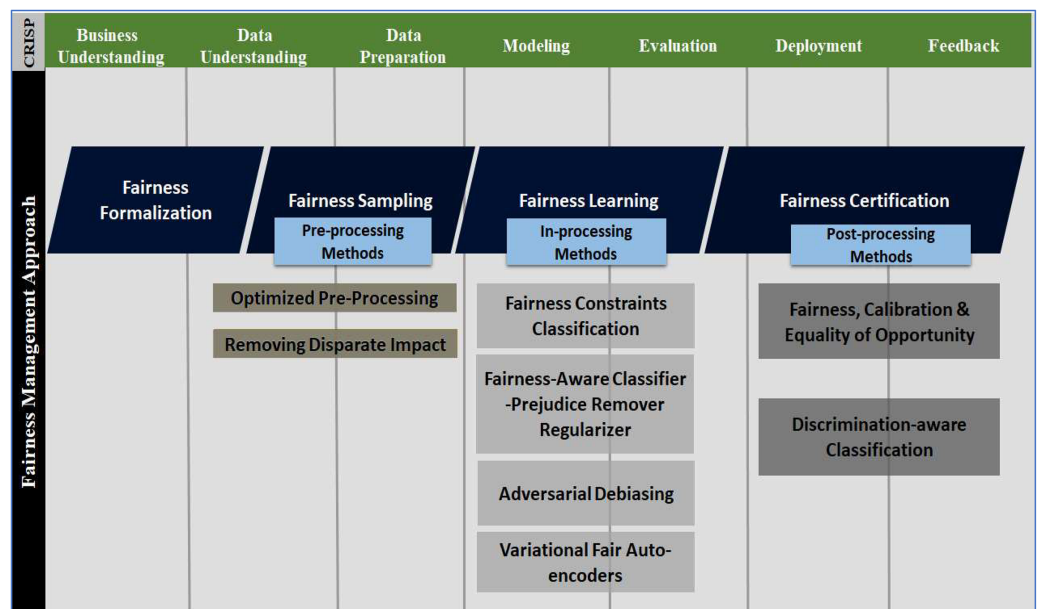


Figure 5 Fairness management approach across multiple phases of the CRISP-DM development cycle. [Full-size](#) DOI: 10.7717/peerj-cs.1630/fig-5

Fairness certification

At the final stage of the testing and deployment phases, it is evaluated whether a prediction aligns with the criteria mentioned in the fairness formalisation phase by executing post-processing methods (Floridi, 2010). Fairness certification solutions verify that unfairness does not surface during the feedback phase.

Limitations and challenges of existing approaches

After thoroughly examining various aspects of tools, methodologies, frameworks, and fairness solution spaces, it is now appropriate to consolidate and summarize the overall major limitations and challenges encountered throughout this exploration:

1. One of the shortcomings of current approaches is the lack of transparency and interpretability (Srinivasu et al., 2022). Many tools and strategies employed to mitigate biases in various domains, such as machine learning algorithms or content moderation systems, often lack clear explanations of how they address bias (Feuerriegel, Dolata & Schwabe, 2020). This lack of transparency makes it difficult for users to understand the underlying biases being addressed and the effectiveness of the applied methods (Berente et al., 2019). Additionally, without proper transparency (Society to Improve Diagnosis in Medicine, 2021), it becomes challenging to identify potential unintended consequences or biases that may arise from the bias management methods themselves (Turney, 1995).
2. Another major shortcoming lies in the limited customization and adaptability (Basereh, Caputo & Brennan, 2021). Most tools and strategies are developed with a one-size-fits-all approach, which may not adequately account for the specific biases

prevalent in different contexts or domains (Qiang, Rhim & Moon, 2023). This limits their effectiveness in managing biases that are nuanced and context-dependent. Furthermore, these methods often do not provide enough flexibility for users to customize and fine-tune the bias management mechanisms according to their specific needs and requirements.

3. Existing methods, tools, and strategies for bias management rely on manual intervention, which makes the process time-consuming and prone to human error. When bias management is carried out manually, it becomes difficult to ensure consistent and comprehensive coverage of all potential biases. Additionally, manual methods may lack scalability and efficiency, particularly when dealing with large datasets or complex models (Chhillar & Aguilera, 2022). Therefore, there is a need for more automated and robust approaches to bias management.
4. Bias can be unintentionally introduced when the dataset used to train AI models is not representative of the real-world population it is designed to serve (Schwartz et al., 2022). For example, if a facial recognition system is trained primarily on data from lighter-skinned individuals, it may exhibit higher error rates for darker-skinned individuals. This lack of diversity can perpetuate existing societal biases and lead to discriminatory outcomes (Delgado et al., 2022; Michael et al., 2022).
5. Existing methodologies tend to place a heavy emphasis on technical solutions for bias mitigation, often neglecting the importance of interdisciplinary collaboration (Madaio et al., 2022). Addressing bias in AI requires input from diverse stakeholders (Michael et al., 2022), including ethicists, social scientists, and policymakers. Without incorporating a multidisciplinary perspective, frameworks may overlook crucial ethical considerations and fail to account for the broader societal impact of AI systems.
6. Many existing strategies tend to oversimplify the concept of bias, reducing it to a binary problem. They often focus on mitigating only explicit biases while overlooking implicit biases, which are more subtle and deeply ingrained in societal structures (Peters, 2022). Addressing implicit biases requires a more nuanced understanding of the underlying social dynamics and power structures. Failure to consider these complexities can result in incomplete or ineffective bias mitigation strategies.
7. Identifying and mitigating the various interlinked biases that can arise in AI systems poses a significant challenge due to their diverse nature and complexity.
8. In some cases, there may be a potential trade-off between fairness and accuracy. For example, if an AI system is designed to be fair to all groups, it may not be as accurate as it could be.
9. There are ethical considerations around how to mitigate bias in AI systems (Straw, 2020). For example, should we prioritize fairness to individuals or to groups? Should we focus on mitigating historical bias or on preventing future bias.
10. The prevailing focus of most approaches lies in addressing bias reactively, leading to high costs associated with corrective measures. There is a pressing need to adopt a proactive stance and mitigate bias as soon as it is identified.

Lack of comprehensive evaluation frameworks is another hurdle in way to achieve fairness. While some methods may claim to address biases, there is often a lack of standardized evaluation frameworks to assess their effectiveness and potential trade-offs (Landers & Behrend, 2023). This absence of robust evaluation frameworks hinders the attainment of fairness on a global scale.

RECOMMENDED PRACTICES TO AVOID/MITIGATE BIAS

After understanding the categorisation and minimising strategies of biases and knowing how to implement a fairness management approach during the model development life cycle, it is time to look at best practices to avoid/mitigate bias and ensure fairness. A few of them are as follows:

- Considering human customs in AI use and promoting all concerned stakeholders on the board *i.e.*, developers, users/general public, policy makers *etc.* to achieve an effective strategy (Floridi, 2010).
- Domain-specific knowledge must be incorporated to detect and mitigate bias (Srinivasan & Chander, 2021).
 - When collecting data, it is vital to have expertise in extracting the most valuable data variables (Pospelov, 2022).
 - Be conscious of the data's sensitive features, including proxy features, as determined by the application (Vasudevan & Kenthapadi, 2020).
 - Datasets should, to the greatest extent possible, represent the actual population being taken into account (Celi *et al.*, 2022). Data selection by random sampling can perform effectively (Pospelov, 2022).
 - Preprocess the data to guarantee the maximum possible level of accuracy while minimising relations between results and sensitive attributes or to present data preserving privacy (Floridi, 2010).
 - For annotating the data, appropriate standards must be specified (Zhang *et al.*, 2022).
 - Consider crucial elements such as the data type, problem, desired outcome, and data size while choosing the best model for the data set (Pospelov, 2022).
- The right model choice is one of the critical components of fairness management. Compared to linear models, which give exact weights for each feature being taken into account, deep models like decision trees can more easily conceal their biases (Barba, 2021).
- Bias detection should be incorporated as a necessary component of model evaluation and focus on model accuracy and precision (Schmelzer, 2020).
- Track the models' performance in use and continually evaluate it (Schmelzer, 2020).
- Encourage all stakeholders to report bias and management should take it positively (Shestakova, 2021).

- In order to prevent perpetuating inequity, AI systems must be responsible during the design, development, evaluation, implementation, and monitoring phases (*Stoyanovich, Howe & Jagadish, 2020*).
- Process and result transparency should be defined in a way that can be interpretable without in-depth knowledge of the algorithm (*Seymour, 2018*).

CHALLENGES, OPPORTUNITIES AND FUTURE WORK

There are several obstacles to overcome before ethical AI applications and systems are successfully developed, free from bias, and wholly engineered along fairness lines. Planning to overcome these obstacles sets the direction of future research. A few serious challenges are:

- Despite numerous approaches to detect and mitigate bias/unfairness, no absolute results are yet available for the cutting edge to handle each type of biasness (*Ntoutsis et al., 2020; Pagano et al., 2023*).
- A mathematical definition cannot express all notions of fairness (*Ntoutsis et al., 2020*). From the ML perspective, literature is enriched with various definitions of fairness. One of the unsolved research issues is how to combine these definitions and propose a unique notion of fairness (*Le Quy et al., 2022*). Achieving it can evaluate AI systems in more unified and comparable manners. The ideals of fairness are incompatible with the operationalisation of the “from equality to equity” concept (*Fenwick & Molnar, 2022*), which necessitates approaching the issue from an operational war posture (*Belenguer, 2022*). Moreover, it demands integrating social and political knowledge in the primary process as critical elements (*Rajkomar et al., 2018; Elizabeth, 2017*).
- The literature review has a lot to say about the bias/fairness of the data and algorithm used by data-driven decision-making systems. However, not all areas have received the same degree of research community attention, *i.e.*, classification, clustering, word embedding, semantic role labelling, representation learning VAE, regression, PCA, named entity recognition, machine translation, language model, graph embedding, coreference resolution, and community detection (*Mehrabi et al., 2021*).
- In the context of fairness-aware ML, exploratory analysis of datasets is still not practiced widely (*Le Quy et al., 2022*). Real, synthetic and sequential decision-making datasets are not adequately exploited.
- The role of sensitive/protected attributes in measuring the performance of predictive models has been studied a lot. However, the role of proxy attributes in the same perspective demands more research work (*Le Quy et al., 2022*).
- Most research being done at the moment focuses on techniques that reduce bias in the underlying machine learning models through the algorithm’s approach. Due to this, a research gap may be filled by a data-centered approach to the subject (*Rajkomar et al., 2018*).

- The oldest benchmark dataset was gathered 48 years ago from nations with active data protection laws. However, to meet the demands of the modern day, general data quality or collection regulations still need to be researched and developed (*Ntoutsis et al., 2020*).

Establishing an agile approach to address bias in AI systems:

A framework for managing bias in AI systems should be created to encourage fairness, accountability, and transparency throughout the AI system lifetime while integrating software engineering best practices, in light of the challenges listed above. Effectively reducing bias requires a multifaceted, adaptable, transparent, scalable, accessible, interdisciplinary, and iterative approach.

Agile methods can be employed to address bias in AI systems. By adopting an agile approach, AI developers can continuously monitor and mitigate bias throughout the development lifecycle and all stakeholders will be well informed with current situation (*Benjamins, Barbado & Sierra, 2019*). Framework based on agile approach can not only mitigate bias at the spot in proactive manners but also eradicate it from appearing in future by eliminating all interlinked biases as (*Landers & Behrend, 2023; Caldwell et al., 2022*). To design a framework for fair AI several working variables play a key role. Let us explore some of these variables in detail:-

Data collection and preprocessing

The first working variable to consider is the data used to train AI models. It is essential to ensure that the data collected is representative and diverse (*Zhao et al., 2018a*), without any inherent biases (*de Bruijn, Warnier & Janssen, 2022*). Biases can emerge if the data reflects historical prejudices or imbalances (*Clarke, 2019*). Careful preprocessing is necessary to identify and address these biases to prevent unfair outcomes. For example, if a facial recognition system is trained predominantly on a specific demographic, it may exhibit racial or gender biases.

Algorithmic transparency and explainability

To design a fair AI framework, it is crucial to consider the transparency and explainability of the algorithms used (*Umbrello & Van de Poel, 2021*). Black-box algorithms that provide no insight into their decision-making processes can pose challenges in identifying and rectifying biases (*Clarke, 2019*). By promoting algorithmic transparency, stakeholders can understand how decisions are made and detect any unfairness in the system. Explainable AI techniques, such as providing interpretable explanations for decisions, can also enhance fairness and accountability (*Toreini et al., 2020*).

Evaluation metrics

Establishing appropriate evaluation metrics is another essential working variable in designing a fair AI framework. The metrics used to assess the performance of AI systems should go beyond traditional accuracy measures (*Buolamwini & Gebru, 2018*) and incorporate fairness considerations (*Clarke, 2019*). For instance, metrics like disparate impact, equal opportunity, and predictive parity can help identify and mitigate biases across different demographic groups. Evaluating AI systems on these fairness metrics ensures that fairness is a fundamental goal rather than an afterthought.

Regular audits and monitoring

A fair AI framework requires ongoing audits and monitoring to identify and rectify biases that may emerge over time. Regular assessments can help evaluate the fairness of AI algorithms and models in real-world scenarios (*Landers & Behrend, 2023*). It allows for continuous improvement and ensures that any biases are promptly addressed. Organizations should establish mechanisms to monitor the performance of AI systems, collect feedback from users, and engage in iterative improvements to enhance fairness (*Saas et al., 2022*). Moreover, encourage continuous learning within your team about bias, fairness, and ethical AI. Continuously stay informed about the most recent research and best practices in this domain and ensure ongoing updates.

Stakeholder inclusivity

Involving diverse stakeholders in the design and deployment of AI systems is a critical working variable for fair AI. Including representatives from different communities, demographic groups, and experts from various fields can help identify potential biases and ensure fairness (*Clarke, 2019*). It is crucial to consider the perspectives and experiences of those who may be disproportionately affected by AI systems. Such inclusivity can lead to a more comprehensive understanding of biases and result in fairer outcomes (*Johansen, Pedersen & Johansen, 2021*).

Ethical guidelines and governance

Ethical guidelines and governance play a vital role in shaping the design of a fair AI framework (*Hildebrandt, 2021*). Organizations should establish clear policies and guidelines that explicitly address fairness concerns. These guidelines should define what constitutes fairness and provide actionable steps to ensure it is upheld. Integrating ethical considerations into the design and decision-making processes can help prevent biases (*Jobin, Ienca & Vayena, 2019*). In nutshell, agile methods provide a framework for effectively addressing bias in AI systems. By adopting diverse teams, defining clear goals and metrics, conducting frequent reviews and iterations, ensuring transparency, curating unbiased datasets, monitoring and evaluating system performance, collaborating with stakeholders, and prioritizing ethics, organizations can develop and deploy AI systems that are fair, unbiased, and inclusive.

CONCLUSION

The presence of bias in our world is reflected in the data. It can appear at any phase of the AI model development cycle. It is not only the developers' core responsibility to ensure model fairness but AI fairness requires a collaborative effort from all stakeholders *i.e.*, policymakers, regulators (*Yavuz, 2019*), users, general public *etc.* for responsible and ethical AI solutions (*Rajkomar et al., 2018*). The future will see AI play an even more significant impact in both our personal and professional lives. Recognising the benefits and challenges of developing ethical and efficient AI models is crucial (*John-Mathews, Cardon & Balagué, 2022*). AI developers bring with them a variety of disciplines and professional experiences.

Narrowing the subject to a single profession or area of expertise would oversimplify the situation. Mapping SDLC, MLLC, and CRISP-DM on a single reference will boast an understanding of practitioners from various technical frameworks to a single point. Once the bias is identified and mitigated at each phase of the development process, the technical team will remain vigilant and unable to de-track or display negligence. The fairness management approach will further enhance the effectiveness of revealing the hidden shortcomings during the development process.

From this perspective, a firm grasp of standard practices will be the foundation for a unified AI framework for fairness management. Through the proposed framework, organisations of all sizes can manage the risk of bias (*Straw, 2020*) throughout a system's lifecycle and ensure that AI is accountable by design. In order to manage the associated risks with AI bias, the proposed framework ought to have the following key characteristics:

- Outlines a technique for conducting impact assessments
- Promotes better awareness of already-existing standards, guidelines, recommendations, best practices, methodologies, and tools and indicates the need for more effective resources.
- Integrate best practices from software engineering similar to 'defect management' to tackle bias as 'defect'. Then, detect, identify and localise bias on the spot before proceeding further and eliminating its chance of appearing in future.
- Ensure each stakeholder comprehends his or her responsibility.
- Sets out corporate governance structures, processes and safeguards that are needed to achieve desired goals.
- Law and regulation agnostic.
- In light of the field of AI's rapid growth, the framework should be updated to ensure it is up-to-date and adapted.

Briefly, AI fairness management is centered on a governance framework that encourages the prevention of bias from manifesting in a way that unjustifiably leads to less favorable or harmful outcomes and enables businesses to create more accurate and practical applications and more persuasive to customers. Overall, a framework for fair AI should provide a structured approach to manage biases and ensure that AI systems operate ethically, fairly, and transparently, while accommodating the complexities and challenges inherent in AI development and deployment.

ADDITIONAL INFORMATION AND DECLARATIONS

Funding

This work was supported by the Higher Education Commission Pakistan. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Grant Disclosures

The following grant information was disclosed by the authors:
Higher Education Commission Pakistan.

Competing Interests

The authors have no conflicts of interest to declare. All co-authors have seen and agree with the contents of the manuscript and there is no financial interest to report.

Author Contributions

- Saadia Afzal Rana conceived and designed the experiments, performed the experiments, analyzed the data, prepared figures and/or tables, authored or reviewed drafts of the article, and approved the final draft.
- Zati Hakim Azizul conceived and designed the experiments, performed the experiments, analyzed the data, prepared figures and/or tables, authored or reviewed drafts of the article, and approved the final draft.
- Ali Afzal Awan conceived and designed the experiments, performed the experiments, analyzed the data, prepared figures and/or tables, authored or reviewed drafts of the article, and approved the final draft.

Data Availability

The following information was supplied regarding data availability:

This is a literature review.

REFERENCES

- Abebe R, Barocas S, Kleinberg J, Levy K, Raghavan M, Robinson DG. 2020. Roles for computing in social change. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. New York: ACM, 252–260.
- Acharyya R, Das S, Chattoraj A, Sengupta O, Tanveer MI. 2020. Detection and mitigation of bias in ted talk ratings. *ArXiv preprint* DOI 10.48550/arXiv.2003.00683.
- Agarwal A, Agarwal H, Agarwal N. 2022. Fairness score and process standardisation: framework for fairness certification in artificial intelligence systems. *AI and Ethics* 3:1–13 DOI 10.1007/s43681-022-00147-7.
- Alelyani S. 2021. Detection and evaluation of machine learning bias. *Applied Science* 11:6271 DOI 10.3390/app11146271.
- Angwin J, Larson J, Mattu S, Kirchner L, Baker E, Goldstein AP, Azevedo IM. 2016. Machine bias. Ethics of data and analytics. *Energy and Climate Change* 2:254–264 DOI 10.1201/9781003278290.
- Atkinson K, Bench-Capon T, Bollegala D. 2020. Explanation in AI and law: past, present and future. *Artificial Intelligence* 289(4):103387 DOI 10.1016/j.artint.2020.103387.
- Bailey D, Faraj S, Hinds P, von Krogh G, Leonardi P. 2019. Special issue of organisation science: emerging technologies and organising. *Organization Science* 30(3):642–646 DOI 10.1287/orsc.2019.1299.
- Balayn A, Lofi C, Houben GJ. 2021. Managing bias and unfairness in data for decision support: a survey of machine learning and data engineering approaches to identify and mitigate bias and

- unfairness within data management and analytics systems. *The VLDB Journal* **30**(5):739–768
DOI 10.1007/s00778-021-00671-8.
- Barba P.** 2021. 6 ways to combat bias in machine learning. Available at <https://builtin.com/machine-learning/bias-machine-learning>.
- Barton C, Chettipally U, Zhou Y, Jiang Z, Lynn-Palevsky A, Le S, Calvert J, Das R.** 2019. Evaluation of a machine learning algorithm for up to 48-hour advance prediction of sepsis using six vital signs. *Computers in Biology and Medicine* **109**(8):79–84
DOI 10.1016/j.compbiomed.2019.04.027.
- Basereh M, Caputo A, Brennan R.** 2021. Fair ontologies for transparent and accountable AI: a hospital adverse incidents vocabulary case study. In: *2021 Third International Conference on Transdisciplinary AI (TransAI)*. Piscataway: IEEE, 92–97.
- Baum SD.** 2017. On the promotion of safe and socially beneficial artificial intelligence. *AI & Society* **32**(4):543–551 DOI 10.1007/s00146-016-0677-0.
- Belenguer L.** 2022. AI bias: exploring discriminatory algorithmic decision-making models and the application of possible machine-centric solutions adapted from the pharmaceutical industry. *AI and Ethics* **2**(4):1–17 DOI 10.1007/s43681-022-00138-8.
- Bellamy RKE, Dey K, Hind M, Hoffman SC, Houde S, Kannan K, Lohia P, Martino J, Mehta S, Mojsilovic A, Nagar S, Ramamurthy KN, Richards J, Saha D, Sattigeri P, Singh M, Varshney KR, Zhang Y.** 2019. AI Fairness 360: an extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development* **4**(5):4:1–4:15
DOI 10.1147/JRD.2019.2942287.
- Benbouzid B.** 2019. Values and consequences in predictive machine evaluation. A sociology of predictive policing. *Science & Technology Studies* **32**(4):119–136.
- Benjamins R, Barbado A, Sierra D.** 2019. Responsible AI by design in practice. *ArXiv preprint*
DOI 10.48550/arXiv.1909.12838.
- Berente N, Gu B, Recker J, Santhanam R.** 2019. Managing AI. *MIS Quarterly* **45**(3):1–5
DOI 10.25300/MISQ/2021/16274.
- Bickel PJ, Hammel EA, O’Connell JW.** 1975. Sex bias in graduate admissions: data from Berkeley: measuring bias is harder than is usually assumed, and the evidence is sometimes contrary to expectation. *Science* **187**(4175):398–404 DOI 10.1126/science.187.4175.398.
- Brusseau J.** 2022. Using edge cases to disentangle fairness and solidarity in AI ethics. *AI and Ethics* **2**(3):441–447 DOI 10.1007/s43681-021-00090-z.
- Bughin J, Seong J, Manyika J, Chui M, Joshi R.** 2018. Notes from the AI frontier: modeling the impact of AI on the world economy. McKinsey Global Institute, Discussion paper, September 2018, Vol. 4, 1–48.
- Buolamwini J, Gebru T.** 2018. Gender shades: intersectional accuracy disparities in commercial gender classification. In: *Conference on Fairness, Accountability and Transparency*. PMLR, 77–91.
- Cakir C.** 2020. Fairness, accountability and transparency—trust in AI and machine learning. In: *The LegalTech Book: The Legal Technology Handbook for Investors, Entrepreneurs and FinTech Visionaries*. Chichester, West Sussex, United Kingdom: John Wiley & Sons Ltd, 35–37
DOI 10.1002/9781119708063.ch9.
- Caldwell S, Sweetser P, O’Donnell N, Knight MJ, Aitchison M, Gedeon T, Johnson D, Brereton M, Gallagher M, Conroy D.** 2022. An agile new research framework for hybrid human-AI teaming: trust, transparency, and transferability. *ACM Transactions on Interactive Intelligent Systems (TiiS)* **12**(3):1–36 DOI 10.1145/3514257.

- Celi LA, Cellini J, Charpignon M-L, Dee EC, Dernoncourt F, Eber R, Mitchell WG, Moukheiber L, Schirmer J, Situ J, Paguio J, Park J, Wawira JG, Yao S. 2022. Sources of bias in artificial intelligence that perpetuate healthcare disparities—a global review. *PLoS Digital Health* **1**(3): e0000022 DOI 10.1371/journal.pdig.0000022.
- Chen J, Kallus N, Mao X, Svacha G, Udell M. 2019. Fairness under unawareness: assessing disparity when protected class is unobserved. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*. New York: ACM, 339–348.
- Chhillar D, Aguilera RV. 2022. An eye for artificial intelligence: insights into the governance of artificial intelligence and vision for future research. *Business & Society* **61**(5):1197–1241 DOI 10.1177/00076503221080959.
- Chouldechova A, Benavides-Prado D, Fialko O, Vaithianathan R. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In: *Conference on Fairness, Accountability and Transparency*. PMLR, 134–148.
- Clarke R. 2019. Regulatory alternatives for AI. *Computer Law & Security Review* **35**(4):398–409 DOI 10.1016/j.clsr.2019.04.008.
- Curto G, Jojoa Acosta MF, Comim F, Garcia-Zapirain B. 2022. Are AI systems biased against the poor? A machine learning analysis using Word2Vec and GloVe embeddings. *AI & Society* (6):1–16 DOI 10.1007/s00146-022-01494-z.
- Dabas A. 2021. Bias in artificial intelligence. Available at <https://adabhishekdabas.medium.com/bias-in-artificial-intelligence-d2ccec3abb2b>.
- Dankwa-Mullan I, Weeraratne D. 2022. Artificial intelligence and machine learning technologies in cancer care: addressing disparities, bias, and data diversity. *Cancer Discovery* **12**(6):1423–1427 DOI 10.1158/2159-8290.CD-22-0373.
- de Bruijn H, Warnier M, Janssen M. 2022. The perils and pitfalls of explainable AI: strategies for explaining algorithmic decision-making. *Government Information Quarterly* **39**(2):101666 DOI 10.1016/j.giq.2021.101666.
- Delgado J, de Manuel A, Parra I, Moyano C, Rueda J, Guersenzvaig A, Ausin T, Cruz M, Casacuberta D, Puyol A. 2022. Bias in algorithms of AI systems developed for COVID-19: a scoping review. *Journal of Bioethical Inquiry* **19**(3):1–13 DOI 10.1007/s11673-022-10200-z.
- Dexe J, Franke U, Nöu AA, Rad A. 2020. Towards increased transparency with value sensitive design. In: *International Conference on Human-Computer Interaction*. Cham: Springer International Publishing, 3–15.
- Dressel J, Farid H. 2018. The accuracy, fairness, and limits of predicting recidivism. *Science Advances* **4**(1):eaa05580 DOI 10.1126/sciadv.aao5580.
- Dua D, Graff C. 2017. UCI machine learning repository. University of California, Irvine, School of Information and Computer Sciences. Available at <https://archive.ics.uci.edu>.
- Egan KM, Trichopoulos D, Stampfer MJ, Willett WC, Newcomb PA, Trentham-Dietz A, Longnecker MP, Baron JA, Stampfer MJ, Willett WC. 1996. Jewish religion and risk of breast cancer. *The Lancet* **347**(9016):1645–1646 DOI 10.1016/S0140-6736(96)91485-3.
- Elizabeth SA. 2017. What is the point of equality? *Theories of Justice*. England, UK: Routledge, 133–183.
- Ethics of AI. 2020. Discrimination and biases. Available at <https://ethics-of-ai.mooc.fi/chapter-6/3-discrimination-and-biases/>.
- Fahse T, Huber V, Giffen BV. 2021. Managing bias in machine learning projects. In: *International Conference on Wirtschaftsinformatik*. Cham: Springer, 94–109.

- Fair Work Ombudsman.** (n.d.). Protection from discrimination at work. Available at <https://www.fairwork.gov.au/employment-conditions/protections-at-work/protection-from-discrimination-at-work>.
- Fenwick A, Molnar G.** 2022. The importance of humanising AI: using a behavioral lens to bridge the gaps between humans and machines. *Discover Artificial Intelligence* **2(1)**:1–12 DOI [10.1007/s44163-022-00030-8](https://doi.org/10.1007/s44163-022-00030-8).
- Feuerriegel S, Dolata M, Schwabe G.** 2020. Fair AI: challenges and opportunities. *Business & Information Systems Engineering* **62(4)**:379–384 DOI [10.1007/s12599-020-00650-3](https://doi.org/10.1007/s12599-020-00650-3).
- Fink A.** 2019. *Conducting research literature reviews: from the internet to paper*. Thousand Oaks: Sage Publications.
- Floridi L.** 2010. *The Cambridge handbook of information and computer ethics*. Cambridge: Cambridge University Press.
- Friedman B, Kahn PH, Borning A, Huldtgren A.** 2013. Value sensitive design and information systems. In: *Early Engagement and New Technologies: Opening Up the Laboratory*. Dordrecht: Springer, 55–95.
- Future Learn.** 2013. Bias and unfairness in data-informed decisions. Available at <https://www.futurelearn.com/info/courses/data-science-artificial-intelligence/0/steps/147783>.
- Gan I, Moussawi S.** 2022. A value sensitive design perspective on AI biases. In: *Hawaii International Conference on System Sciences*. 1–10.
- Gardner A, Smith AL, Steventon A, Coughlan E, Oldfield M.** 2022. Ethical funding for trustworthy AI: proposals to address the responsibilities of funders to ensure that projects adhere to trustworthy AI practice. *AI and Ethics* **2(2)**:277–291 DOI [10.1007/s43681-021-00069-w](https://doi.org/10.1007/s43681-021-00069-w).
- Garg A, Si R.** 2021. PCIV method for indirect bias quantification in AI and ML models. DOI [10.32628/CSEIT217251](https://doi.org/10.32628/CSEIT217251).
- González Esteban E, Calvo P.** 2022. Ethically governing artificial intelligence in the field of scientific research and innovation. *Heliyon* **8(2)**:e08946 DOI [10.1016/j.heliyon.2022.e08946](https://doi.org/10.1016/j.heliyon.2022.e08946).
- Gupta M, Parra CM, Dennehy D.** 2021. Questioning racial and gender bias in AI-based recommendations: do espoused national cultural values matter? *Information Systems Frontiers* **24(5)**:1–17 DOI [10.1007/s10796-021-10156-2](https://doi.org/10.1007/s10796-021-10156-2).
- Hildebrandt M.** 2021. The issue of bias. The framing powers of machine learning. Available at <https://mitpress.mit.edu/books/machines-we-trust>.
- Ho DA, Beyan O.** 2020. Biases in data science lifecycle. *ArXiv preprint* DOI [10.48550/arXiv.2009.09795](https://doi.org/10.48550/arXiv.2009.09795).
- Hobson Z, Yesberg JA, Bradford B, Jackson J.** 2021. Artificial fairness? Trust in algorithmic police decision-making. *Journal of Experimental Criminology* **19(1)**:1–25 DOI [10.1007/s11292-021-09484-9](https://doi.org/10.1007/s11292-021-09484-9).
- Hofeditz L, Mirbabaie M, Luther A, Mauth R, Rentemeister I.** 2022. Ethics guidelines for using AI-based algorithms in recruiting: learnings from a systematic literature review. In: *Proceedings of the 55th Hawaii International Conference on System Sciences*. 1–10.
- Howard A, Borenstein J.** 2018. The ugly truth about ourselves and our robot creations: the problem of bias and social inequity. *Science and Engineering Ethics* **24(5)**:1521–1536 DOI [10.1007/s11948-017-9975-2](https://doi.org/10.1007/s11948-017-9975-2).
- Huang J, Galal G, Etemadi M, Vaidyanathan M.** 2022. Evaluation and mitigation of racial bias in clinical machine learning models: scoping review. *JMIR Medical Informatics* **10(5)**:e36388 DOI [10.2196/36388](https://doi.org/10.2196/36388).

- Jobin A, Ienca M, Vayena E. 2019.** The global landscape of AI ethics guidelines. *Nature Machine Intelligence* **1(9)**:389–399 DOI [10.1038/s42256-019-0088-2](https://doi.org/10.1038/s42256-019-0088-2).
- Johansen J, Pedersen T, Johansen C. 2021.** Studying human-to-computer bias transference. *AI & Society* **38(4)**:1–25 DOI [10.1007/s00146-021-01328-4](https://doi.org/10.1007/s00146-021-01328-4).
- John-Mathews JM, Cardon D, Balagué C. 2022.** From reality to world. A critical perspective on AI fairness. *Journal of Business Ethics* **178(4)**:1–15 DOI [10.1007/s10551-022-05055-8](https://doi.org/10.1007/s10551-022-05055-8).
- Karras T, Laine S, Aila T. 2019.** A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Piscataway: IEEE, 4401–4410.
- Kim J, Lee H, Cho YH. 2022.** Learning design to support student-AI collaboration: perspectives of leading teachers for AI in education. *Education and Information Technologies* **27(1)**:1–36 DOI [10.1007/s10639-021-10831-6](https://doi.org/10.1007/s10639-021-10831-6).
- Landers RN, Behrend TS. 2023.** Auditing the AI auditors: a framework for evaluating fairness and bias in high stakes AI predictive models. *American Psychologist* **78(1)**:36 DOI [10.1037/amp0000972](https://doi.org/10.1037/amp0000972).
- Le Quy T, Roy A, Iosifidis V, Zhang W, Ntoutsis E. 2022.** A survey on datasets for fairness-aware machine learning. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **12(3)**: e1452 DOI [10.1002/widm.1452](https://doi.org/10.1002/widm.1452).
- Li J, Chignell M. 2022.** FMEA-AI: AI fairness impact assessment using failure mode and effects analysis. *AI and Ethics* **2(4)**:1–14 DOI [10.1007/s43681-022-00145-9](https://doi.org/10.1007/s43681-022-00145-9).
- Luengo-Oroz M. 2019.** Solidarity should be a core ethical principle of AI. *Nature Machine Intelligence* **1(11)**:494 DOI [10.1038/s42256-019-0115-3](https://doi.org/10.1038/s42256-019-0115-3).
- Luengo-Oroz M, Bullock J, Pham KH, Lam CSN, Luccioni A. 2021.** From artificial intelligence bias to inequality in the time of COVID-19. *IEEE Technology and Society Magazine* **40(1)**:71–79 DOI [10.1109/MTS.2021.3056282](https://doi.org/10.1109/MTS.2021.3056282).
- Maathuis C. 2022.** On explainable AI solutions for targeting in cyber military operations. *International Conference on Cyber Warfare and Security* **17(1)**:166–175 DOI [10.34190/iccws.17.1.38](https://doi.org/10.34190/iccws.17.1.38).
- Machine Learning Glossary: Fairness.** (n.d.). Google developers. Available at <https://developers.google.com/machine-learning/glossary/fairness> (accessed 22 March 2022).
- Madaio M, Egede L, Subramonyam H, Wortman Vaughan J, Wallach H. 2022.** Assessing the fairness of AI systems: AI practitioners' processes, challenges, and needs for support. *Proceedings of the ACM on Human-Computer Interaction* **6(CSCW1)**:1–26 DOI [10.1145/3512899](https://doi.org/10.1145/3512899).
- Mehrabi N, Morstatter F, Saxena N, Lerman K, Galstyan A. 2021.** A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* **54(6)**:1–35 DOI [10.1145/3457607](https://doi.org/10.1145/3457607).
- Merler M, Ratha N, Feris RS, Smith JR. 2019.** Diversity in faces. *ArXiv preprint* DOI [10.48550/arXiv.1901.10436](https://doi.org/10.48550/arXiv.1901.10436).
- Michael K, Abbas R, Jayashree P, Bandara RJ, Aloudat A. 2022.** Biometrics and AI bias. *IEEE Transactions on Technology and Society* **3(1)**:2–8 DOI [10.1109/TTS.2022.3156405](https://doi.org/10.1109/TTS.2022.3156405).
- Mikians J, Gyarmati L, Erramilli V, Laoutaris N. 2012.** Detecting price and search discrimination on the Internet. In: *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*. New York: ACM, 79–84.
- Nascimento AM, da Cunha MAVC, de Souza Meirelles F, Scornavacca E Jr, De Melo VV. 2018.** A literature analysis of research on artificial intelligence in management information system (MIS). In: *AMCIS*.

- National Institute of Standards and Technology.** 2022. There's more to AI bias than biased data, NIST report highlights. Available at <https://www.nist.gov/news-events/news/2022/03/theres-more-ai-bias-biased-data-nist-report-highlights>.
- Northeastern Global News.** 2020. Here's what happened when Boston tried to assign students good schools close to home. Available at <https://news.northeastern.edu/2018/07/16/heres-what-happened-when-boston-tried-to-assign-students-good-schools-close-to-home/>.
- Ntoutsis E, Fafalios P, Gadiraju U, Iosifidis V, Nejdil W, Vidal M-E, Ruggieri S, Turini F, Papadopoulos S, Krasanakis E, Kompatsiaris I, Kinder-Kurlanda K, Wagner C, Karimi F, Fernandez M, Alani H, Berendt B, Kruegel T, Heinze C, Broelemann K, Kasneci G, Tiropanis T, Staab S.** 2020. Bias in data-driven artificial intelligence systems—an introductory survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **10**(3):e1356 DOI [10.1002/widm.1356](https://doi.org/10.1002/widm.1356).
- Orphanou K, Otterbacher J, Kleanthous S, Batsuren K, Giunchiglia F, Bogina V, Tal AS, Kuflik T.** 2021. Mitigating bias in algorithmic systems—a fish-eye view. *ACM Computing Surveys (CSUR)* **55**(5):1–37 DOI [10.1145/3527152](https://doi.org/10.1145/3527152).
- Osoba OA, Welser W IV.** 2017. *An intelligence in our image: the risks of bias and errors in artificial intelligence*. Santa Monica: Rand Corporation.
- O'Neil C.** 2017. *Weapons of math destruction: how big data increases inequality and threatens democracy*. Belvoir Castle: Crown.
- Pagano TP, Loureiro RB, Lisboa FVN, Cruz GOR, Peixoto RM, de Sousa Guimarães GA, dos Santos LL, Araujo MM, Cruz M, de Oliveira ELS, Winkler I, Nascimento EGS.** 2022. Bias and unfairness in machine learning models: a systematic literature review. *ArXiv preprint* DOI [10.48550/arXiv.2202.08176](https://doi.org/10.48550/arXiv.2202.08176).
- Pagano TP, Loureiro RB, Lisboa FVN, Peixoto RM, Guimarães GAS, Cruz GOR, Araujo MM, Santos LL, Cruz MAS, Oliveira ELS, Winkler I, Nascimento EGS.** 2023. Bias and unfairness in machine learning models: a systematic review on datasets, tools, fairness metrics, and identification and mitigation methods. *Big Data and Cognitive Computing* **7**(1):15 DOI [10.3390/bdcc7010015](https://doi.org/10.3390/bdcc7010015).
- Pedreshi D, Ruggieri S, Turini F.** 2008. Discrimination-aware data mining. In: *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York: ACM, 560–568.
- Peters U.** 2022. Algorithmic political bias in artificial intelligence systems. *Philosophy & Technology* **35**(2):25 DOI [10.1007/s13347-022-00512-8](https://doi.org/10.1007/s13347-022-00512-8).
- Pethig F, Kroenung J.** 2022. Biased humans, (un)biased algorithms? *Journal of Business Ethics* **183**:1–16 DOI [10.1007/s10551-022-05071-8](https://doi.org/10.1007/s10551-022-05071-8).
- Pospelov S.** 2022. How to reduce bias in machine learning. Available at <https://www.spiceworks.com/tech/artificial-intelligence/guest-article/how-to-reduce-bias-in-machine-learning/>.
- Qiang V, Rhim J, Moon A.** 2023. No such thing as one-size-fits-all in AI ethics frameworks: a comparative case study. *AI & Society* (**38**):1–20 DOI [10.1007/s00146-023-01653-w](https://doi.org/10.1007/s00146-023-01653-w).
- Rajkomar A, Hardt M, Howell MD, Corrado G, Chin MH.** 2018. Ensuring fairness in machine learning to advance health equity. *Annals of Internal Medicine* **169**(12):866–872 DOI [10.7326/M18-1990](https://doi.org/10.7326/M18-1990).
- Redmond M.** 2011. Communities and crime unnormalised data set. UCI Machine Learning Repository. Available at <http://www.ics.uci.edu/mllearn/MLRepository.html>.
- Rhue L, Clark J.** 2020. Automatically signaling quality? A study of the fairness-economic tradeoffs in reducing bias through AI/ML on digital platforms. A Study of the Fairness-Economic

- Tradeoffs in Reducing Bias through AI/ML on Digital Platforms (January 10, 2020). NYU Stern School of Business DOI 10.2139/ssrn.3547502.
- Ricardo BY. 2018.** Bias on the web. *Communications of the ACM* **61(6)**:54–61 DOI 10.1145/3209581.
- Richardson B, Garcia-Gathright J, Way SF, Thom J, Cramer H. 2021.** Towards fairness in practice: a practitioner-oriented rubric for evaluating fair ML toolkits. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. New York: ACM, 1–13.
- Richardson B, Gilbert JE. 2021.** A framework for fairness: a systematic review of existing fair AI solutions. *ArXiv preprint* DOI 10.48550/arXiv.2112.05700.
- Romei A, Ruggieri S. 2014.** A multidisciplinary survey on discrimination analysis. *The Knowledge Engineering Review* **29(5)**:582–638 DOI 10.1017/S0269888913000039.
- Saas H, Zlatkin-Troitschanskala O, Reichert-Schlx J, Brückner S, Kuhn C, Dormann C. 2022.** Aera (iPosterSessions—an aMuze! Interactive system). Available at <https://aera22-aera.ipostersessions.com/Default.aspx?s=0F-E4-C1-0C-7F-64-F5-D8-26-42-67-7C-A6-EB-B7-DC>.
- Saleiro P, Kuester B, Hinkson L, London J, Stevens A, Anisfeld A, Rodolfa KT, Ghani R. 2018.** Aequitas: a bias and fairness audit toolkit. *ArXiv preprint* DOI 10.48550/arXiv.1811.05577.
- Schmelzer R. 2020.** 6 ways to reduce different types of bias in machine learning. Available at <https://www.techtarget.com/searchenterpriseai/feature/6-ways-to-reduce-different-types-of-bias-in-machine-learning>.
- Schwartz R, Vassilev A, Greene K, Perine L, Burt A, Hall P. 2022.** Towards a standard for identifying and managing bias in artificial intelligence. *NIST Special Publication* **1270(10)**:6028 DOI 10.6028/NIST.SP.1270.
- Seymour W. 2018.** Detecting bias: Does an algorithm have to be transparent in order to be fair? In: Bates J, Clough PD, Jäschke R, Otterbacher J, eds. *Proceedings of the International Workshop on Bias in Information, Algorithms, and Systems (BIAS 2018)*, Vol. 2103, 1–7 DOI 10.1145/3183880.3183906.
- Shestakova V. 2021.** Best practices to mitigate bias and discrimination in artificial intelligence. *Performance Improvement* **60(6)**:6–11 DOI 10.1002/pfi.21987.
- Silva S, Kenney M. 2019.** Algorithms, platforms, and ethnic bias. *Communications of the ACM* **62(11)**:37–39 DOI 10.1145/3318157.
- Simon J, Wong PH, Rieder G. 2020.** Algorithmic bias and the value sensitive design approach. *Internet Policy Review* **9(4)**:1–16 DOI 10.14763/2020.4.1534.
- Society to Improve Diagnosis in Medicine. 2021.** Bias in artificial intelligence. Available at <https://www.improvediagnosis.org/publications/improvedx-may-2021/bias-in-artificial-intelligence/>.
- Srinivasan R, Chander A. 2021.** Biases in AI systems. *Communications of the ACM* **64(8)**:44–49 DOI 10.1145/3464903.
- Srinivasu PN, Sandhya N, Jhaveri RH, Raut R. 2022.** From blackbox to explainable AI in healthcare: existing tools and case studies. *Mobile Information Systems* **2022(1)**:1–20 DOI 10.1155/2022/8167821.
- Stoyanovich J, Howe B, Jagadish HV. 2020.** Responsible data management. *Proceedings of the VLDB Endowment* **13(12)**:3474–3488 DOI 10.14778/3415478.3415570.
- Straw I. 2020.** The automation of bias in medical artificial intelligence (AI): decoding the past to create a better future. *Artificial Intelligence in Medicine* **110(3)**:101965 DOI 10.1016/j.artmed.2020.101965.

- Suresh H, Gutttag J. 2021.** A framework for understanding sources of harm throughout the machine learning life cycle. In: *Equity and Access in Algorithms, Mechanisms, and Optimisation*. New York: ACM, 1–9.
- Suri JS, Agarwal S, Jena B, Saxena S, El-Baz A, Agarwal V, Kalra MK, Saba L, Viskovic K, Fatemi M, Naidu S. 2022.** Five strategies for bias estimation in artificial intelligence-based hybrid deep learning for acute respiratory distress syndrome COVID-19 lung infected patients using AP (AI) bias 2.0: a systematic review. *Transactions and Instrument Measures* (27):140–148 DOI 10.1109/TIM.2022.3174270.
- Toreini E, Aitken M, Coopamootoo K, Elliott K, Zelaya CG, Van Moorsel A. 2020.** The relationship between trust in AI and trustworthy machine learning technologies. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. New York: ACM, 272–283.
- Turney P. 1995.** Bias and the quantification of stability. *Machine Learning* 20(1):23–33 DOI 10.1007/BF00993473.
- Umbrello S, Van de Poel I. 2021.** Mapping value sensitive design onto AI for social good principles. *AI and Ethics* 1(3):283–296 DOI 10.1007/s43681-021-00038-3.
- Vasudevan S, Kenthapadi K. 2020.** Lift: a scalable framework for measuring fairness in ML applications. In: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. New York: ACM, 2773–2780.
- Victor D. 2016.** Microsoft created a Twitter bot to learn from users. It quickly became a racist Jerk. The New York Times, 24. Available at <https://www.nytimes.com/2016/03/25/technology/microsoft-created-a-twitter-bot-to-learn-from-users-it-quickly-became-a-racist-jerk.html>.
- Walch K. 2020.** 9 data quality issues that can sideline AI projects. Available at <https://www.techtarget.com/searchenterpriseai/feature/9-data-quality-issues-that-can-sideline-AI-projects>.
- Weith HVK, Matt C. 2022.** When do customers perceive artificial intelligence as fair? An assessment of AI-based B2C E-Commerce. Available at <https://boris.unibe.ch/id/eprint/159654>.
- Yavuz C. 2019.** Machine bias: artificial intelligence and discrimination. Available at <https://lup.lub.lu.se/luur/download?func=downloadFile&recordId=8987035&fileId=8987040>.
- Zhang Z, Li J, Stork DG, Mansfield E, Russell J, Adams C, Wang JZ. 2022.** Reducing bias in AI-based analysis of visual artworks. *IEEE BITS the Information Theory Magazine* 34:1–16 DOI 10.1109/MBITS.2022.3197102.
- Zhao J, Wang T, Yatskar M, Ordonez V, Chang KW. 2018a.** Gender bias in coreference resolution: evaluation and debiasing methods. *ArXiv preprint* DOI 10.48550/arXiv.1804.06876.
- Zhao J, Zhou Y, Li Z, Wang W, Chang KW. 2018b.** Learning gender-neutral word embeddings. *ArXiv preprint* DOI 10.48550/arXiv.1809.01496.
- Zuiderveen Borgesius F. 2018.** Discrimination, artificial intelligence, and algorithmic decision-making. Council of Europe, Directorate General of Democracy. Available at <https://rm.coe.int/discrimination-artificial-intelligence-and-algorithmic-decisionmaking/1680925d73>.