



## OPEN ACCESS

## EDITED BY

Jose Santamaria Lopez,  
University of Jaén, Spain

## REVIEWED BY

Luisa Varriale,  
University of Naples Parthenope, Italy  
Reza Shahbazian,  
University of Calabria, Italy  
Tim Hulsen,  
Philips, Netherlands

## \*CORRESPONDENCE

Dominik Kowald  
✉ dominik.kowald@tugraz.at  
Simone Kopeinik  
✉ skopeinik@know-center.at

RECEIVED 19 July 2024

ACCEPTED 11 November 2024

PUBLISHED 29 November 2024

## CITATION

Kowald D, Scher S, Pammer-Schindler V,  
Müllner P, Waxnegger K, Demelius L, Fessler A,  
Toller M, Mendoza Estrada IG, Šimić I, Sabol V,  
Trügler A, Veas E, Kern R, Nad T and Kopeinik S  
(2024) Establishing and evaluating trustworthy  
AI: overview and research challenges.  
*Front. Big Data* 7:1467222.  
doi: 10.3389/fdata.2024.1467222

## COPYRIGHT

© 2024 Kowald, Scher, Pammer-Schindler,  
Müllner, Waxnegger, Demelius, Fessler, Toller,  
Mendoza Estrada, Šimić, Sabol, Trügler, Veas,  
Kern, Nad and Kopeinik. This is an  
open-access article distributed under the  
terms of the [Creative Commons Attribution  
License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or  
reproduction in other forums is permitted,  
provided the original author(s) and the  
copyright owner(s) are credited and that the  
original publication in this journal is cited, in  
accordance with accepted academic practice.  
No use, distribution or reproduction is  
permitted which does not comply with these  
terms.

# Establishing and evaluating trustworthy AI: overview and research challenges

Dominik Kowald<sup>1,2\*</sup>, Sebastian Scher<sup>1,3</sup>,  
Viktoria Pammer-Schindler<sup>1,2</sup>, Peter Müllner<sup>1</sup>,  
Kerstin Waxnegger<sup>1</sup>, Lea Demelius<sup>1,2</sup>, Angela Fessler<sup>1,2</sup>,  
Maximilian Toller<sup>1</sup>, Inti Gabriel Mendoza Estrada<sup>1</sup>, Ilija Šimić<sup>1</sup>,  
Vedran Sabol<sup>1</sup>, Andreas Trügler<sup>1,2,3</sup>, Eduardo Veas<sup>1,2</sup>,  
Roman Kern<sup>1,2</sup>, Tomislav Nad<sup>4</sup> and Simone Kopeinik<sup>1\*</sup>

<sup>1</sup>Know Center Research GmbH, Graz, Austria, <sup>2</sup>Institute of Interactive Systems and Data Science, Graz University of Technology, Graz, Austria, <sup>3</sup>Department of Geography and Regional Science, Wegener Center for Climate and Global Change, University of Graz, Graz, Austria, <sup>4</sup>SGS Digital Trusts Services GmbH, Graz, Austria

Artificial intelligence (AI) technologies (re-)shape modern life, driving innovation in a wide range of sectors. However, some AI systems have yielded unexpected or undesirable outcomes or have been used in questionable manners. As a result, there has been a surge in public and academic discussions about aspects that AI systems must fulfill to be considered trustworthy. In this paper, we synthesize existing conceptualizations of trustworthy AI along six requirements: (1) human agency and oversight, (2) fairness and non-discrimination, (3) transparency and explainability, (4) robustness and accuracy, (5) privacy and security, and (6) accountability. For each one, we provide a definition, describe how it can be established and evaluated, and discuss requirement-specific research challenges. Finally, we conclude this analysis by identifying overarching research challenges across the requirements with respect to (1) interdisciplinary research, (2) conceptual clarity, (3) context-dependency, (4) dynamics in evolving systems, and (5) investigations in real-world contexts. Thus, this paper synthesizes and consolidates a wide-ranging and active discussion currently taking place in various academic sub-communities and public forums. It aims to serve as a reference for a broad audience and as a basis for future research directions.

## KEYWORDS

trustworthy AI, artificial intelligence, fairness, human agency, robustness, privacy, accountability, transparency

## 1 Introduction

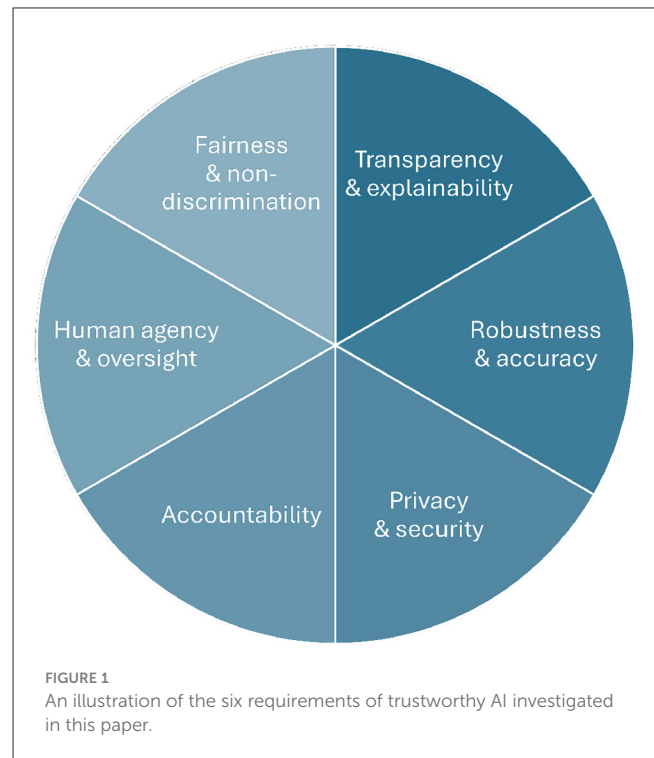
From sophisticated chatbots like Chat-GPT to AI-driven recommender systems enhancing our entertainment experiences on platforms like Netflix and Spotify (Anderson et al., 2020), the impact of AI on our lives is significant. AI-based decision support systems are proving invaluable in critical fields such as life science and healthcare (Rajpurkar et al., 2022). Similarly, AI is reshaping hiring and human resources practice (Van den Broek et al., 2021) and transforming the banking and finance landscape with innovative solutions (Cao, 2022). However, in the past, some AI systems have been used in questionable manners, which has led to unexpected or undesirable results. Examples include biased algorithms perpetuating discrimination in recruitment processes (Chen, 2023) or AI-driven recommender systems favoring popular content and, with this, users

interested in popular content (Kowald et al., 2020; Kowald and Lacic, 2022). Alongside biases in algorithms, AI systems rely on training data, including personal and private user information, which raises concerns for potential privacy and security breaches. One example is the Equifax data breach, in which private data records of millions of users were compromised (Zou and Schaub, 2018). Additionally, when thinking of self-driving cars, unreliable AI-based systems could even cause physical harm, as demonstrated by the unfortunate Uber car crash in 2018, in which a malfunctioning algorithm did not detect and, as a consequence, killed a pedestrian on the road (Kohli and Chadha, 2020).

As a consequence, there has been an increase in public and academic discussions about the essential requirements AI systems must fulfill to be considered trustworthy. There is also a growing consensus on the necessity of setting up standards and regulations to ensure and validate the trustworthiness of AI. In this respect, the European Commission (EC) has proposed the AI Act (Madiaga, 2021), a comprehensive regulatory framework for supporting the responsible development and deployment of AI technologies within the European Union. The AI Act seeks to establish clear rules governing the development and deployment of AI systems while imposing strict requirements for high-risk AI applications. The various interpretations of trustworthy AI add further complexity to this discourse by encompassing not just technical requirements but also human-centered and legal considerations. Another important framework proposed by the European Commission has been the “Assessment List for Trustworthy AI (ALTAI)” (Ala-Pietilä et al., 2020; Radclyffe et al., 2023), which enables organizations to self-assess the trustworthiness of AI solutions based on a checklist.

This paper contributes insights into this discourse by analyzing the state-of-the-art regarding six aspects of AI systems that are typically understood as requirements for systems to be viewed as trustworthy. These requirements are: (1) human agency and oversight, (2) fairness and non-discrimination, (3) transparency and explainability, (4) robustness and accuracy, (5) privacy and security, and (6) accountability (see Figure 1). We define each of these six requirements, introduce methods to establish and implement these requirements in AI systems, and discuss corresponding validation methods and evaluation metrics. Such validation efforts are crucial from scientific and practical perspectives and might serve as a prerequisite for certifying AI systems and models (Winter et al., 2021). Finally, for each of these requirements, we outline ongoing research challenges and future research perspectives.

The contributions of our work are two-fold: firstly, we give a comprehensive overview of the requirements of trustworthy AI, in which we cover different viewpoints on trustworthy AI, including technical and also human-centered and legal considerations. Secondly, we discuss open issues and challenges in defining, establishing, and evaluating these requirements of trustworthy AI. Therefore, the guiding research question of this work is defined as follows: *What is the current state of research regarding the establishment and evaluation of comprehensive - technical, human-centered, and legal - requirements of trustworthy AI?* To address



this research question, we follow the methodology described in Section 2.3.

Our work complements existing surveys and articles on trustworthy AI in two main ways. Firstly, existing overview articles such as Chatila et al. (2021); Thiebes et al. (2021); Akbar et al. (2024); Díaz-Rodríguez et al. (2023) tend to focus on definitions of trustworthy AI and neglect evaluation aspects, which is one key aspect of our article. Specifically, related surveys such as Liang et al. (2022); Wing (2021); Emaminejad and Akhavian (2022) focus on specific aspects of trustworthy AI implementation and evaluation, namely data, formal methods, and robotics, respectively. In contrast, our article aims to provide a domain- and method-independent overview of trustworthy AI, which reflects the whole AI-lifecycle, including the evaluation phase. Secondly, concerning validation and evaluation schemes for trustworthy AI, existing technical conceptualizations of trustworthy AI such as Floridi (2021); Kaur et al. (2022); Li et al. (2023) have focused on technical and reliability-oriented requirements such as transparency, privacy, and robustness. In contrast, in our paper, we discuss methods and open challenges toward establishing and evaluating trustworthy AI also through the lens of human-centric and legal requirements such as fairness, accountability, and human agency. Therefore, to the best of our knowledge, our paper is the first to investigate all six requirements of trustworthy AI in a unified way by discussing implementation and evaluation aspects across the whole lifecycle of trustworthy AI and outlining open research challenges and issues for all six requirements.

Our article shows that while evaluation and validation methodologies for technical requirements like robustness may rely on established metrics and testing procedures (e.g., for model

accuracy), the assessment of human-centric considerations often requires more nuanced approaches that consider ethical, legal, and cultural factors. As such, our article emphasizes the need for further research to develop robust evaluation schemes that can be applied in research and practice across a variety of AI systems, particularly in high-risk domains where human values and rights are at stake (e.g., healthcare).

Next, in Section 2, we describe the relevant background for this article, including general definitions of AI and its lifecycle, and introduce the six requirements of trustworthy AI covered herein. After discussing each requirement separately in Section 3, the paper closes with a conclusion and an outlook into future research directions in Section 4.

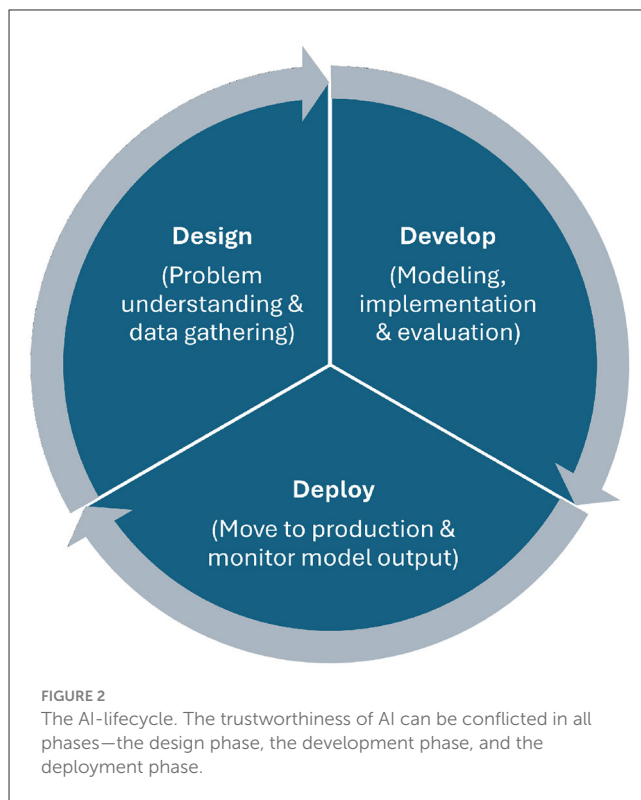
## 2 Background

In this section, we give a short overview of definitions and preliminaries relevant to our article, introduce the six requirements of trustworthy AI discussed, and describe the methodology of our investigation.

### 2.1 Definitions and preliminaries of trustworthy AI

For our understanding of AI in the context of this work, we adhere to the definition outlined in the EU AI Act (adopted text, Art 3(1),<sup>1</sup>), in which AI is defined as “a machine-based system designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.” This definition encompasses a broad spectrum of algorithmic implementations, from simple logistic regression models to complex machine-learning approaches. In the article at hand, we consider this spectrum of AI systems, recognizing the diverse challenges and requirements that are associated with ensuring its trustworthiness.

Additionally, we aim to consider the trustworthiness of AI from a holistic perspective that can be influenced in all phases of the AI-lifecycle, as thoroughly described in Haakman et al. (2021) (see Figure 2). These phases encompass the design, development, and deployment of AI-based systems and their designated tasks. While discussing the requirements of trustworthy AI, we refer to the phases where needed in this article. To comply with trustworthy AI, special attention should be paid to considering AI requirements in the *design phase* throughout requirements engineering, problem understanding, and data collection strategies, the *development phase*, comprising model implementation (e.g., optimizing feature weights), documentation, and evaluation, and finally, the *deployment phase* including the integration of the AI model into a production



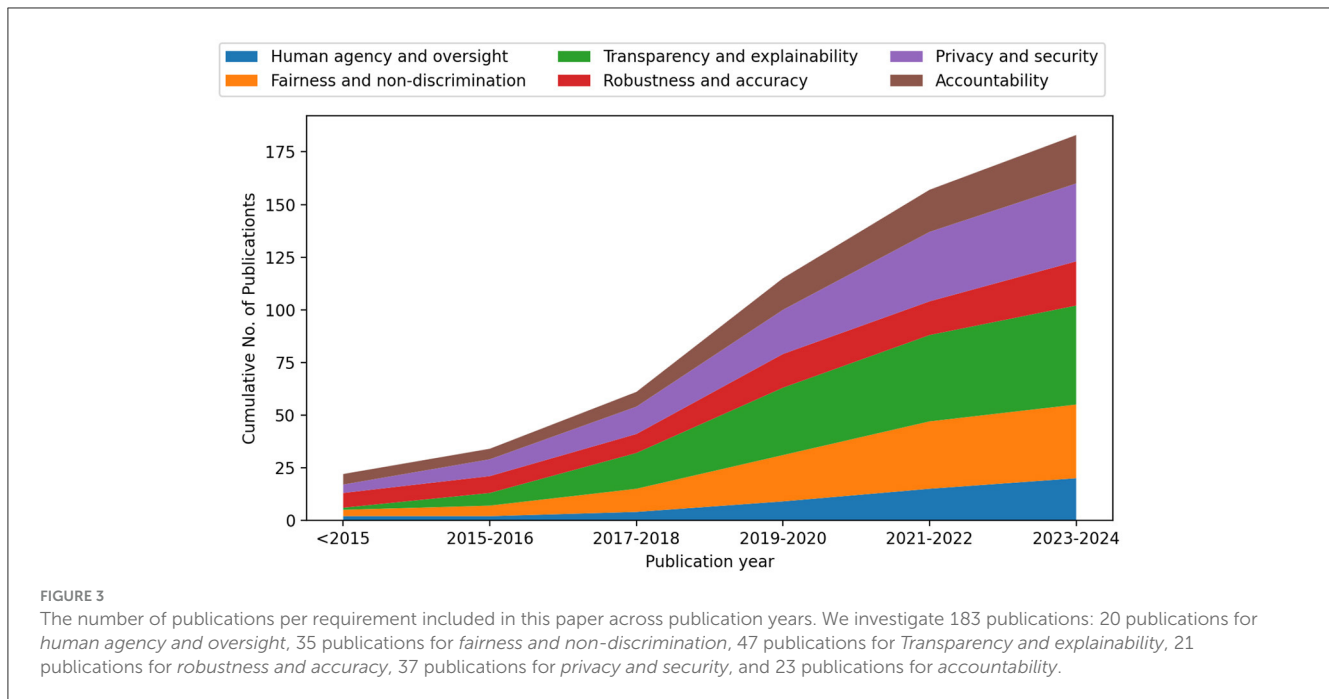
environment, and the continuous monitoring and updating of the model.

### 2.2 Requirements of trustworthy AI

Over the last years, various frameworks, guidelines, survey articles, and collections of requirements of trustworthy AI have been developed and published by researchers, governments, and private organizations (Smuha, 2019; Kaur et al., 2021; Floridi, 2021; Kaur et al., 2022; Li et al., 2023; Yeung, 2020). Although these investigations differ with respect to the exact wordings, they agree on four fundamental principles that need to be considered when developing and validating trustworthy AI: (1) respect for human autonomy, (2) fairness, (3) explicability, and (4) prevention of harm (Smuha, 2019; Kaur et al., 2021, 2022). In the following, we describe six requirements of trustworthy AI that are manifested within these principles.

Hereinafter, principle 1 (respect for human control) is mainly associated with **human agency and oversight** (requirement 1, see Section 3.1), which refers to sustaining the autonomy of humans affected by AI systems, given different levels of human-AI interaction. The second principle (fairness) aims for the equal treatment of all affected individuals and subpopulations (i.e., defined by age, gender, education, ...). **Fairness and non-discrimination** (requirement 2, see Section 3.2) describes the absence of bias in AI decisions that could result in unfair, unequal treatment that negatively affects certain people. Next, principle 3 (explicability) ensures the AI system is transparent and explainable.

<sup>1</sup> [https://www.europarl.europa.eu/doceo/document/TA-9-2024-0138\\_EN.pdf](https://www.europarl.europa.eu/doceo/document/TA-9-2024-0138_EN.pdf)



In particular, **transparency and explainability** (requirement 3, see Section 3.3) is defined as the understandability of an AI system and the provision of information to explain the AI model's decisions.

Finally, principle 4 (prevention of harm) should ensure that the AI system does not cause any harm to humans, society, or the environment. For example, it should be prevented that AI-based systems harm or, even worse, kill humans, which unfortunately has happened, e.g., in the aforementioned Uber car crash in 2018, in which a pedestrian was killed by a malfunctioned self-driving car (Kohli and Chadha, 2020). This principle includes a range of requirements, including technical and legal concerns. An essential technical requirement within this principle is **robustness and accuracy** (requirement 4, see Section 3.4), which is related to the performance of AI models and their ability to function in unexpected circumstances. Additionally, the principle "prevention of harm" is linked to **privacy and security** (requirement 5, see Section 3.5) that focuses on protecting the personal and sensitive information of users in AI systems and on preventing malicious attacks on AI models. Concerning legal aspects, **accountability** (requirement 6, see Section 3.6) entails the understanding of who is responsible for the decisions of AI systems and to ensure that mechanisms are in place to interfere with negative consequences. The literature (e.g., Kaur et al., 2021; Smuha, 2019) also discusses other important requirements of trustworthy AI within this principle, such as safety, reproducibility, sustainability, societal and environmental wellbeing, and data governance. While we highlight these requirements' significance, we believe they serve as overarching aspects that underpin the six other requirements of trustworthy AI discussed in this article. However, we strongly suggest a foundational commitment to safety, reproducibility, sustainability, societal and environmental wellbeing, and data

governance when developing and validating trustworthy AI. In the next section, we discuss in detail these six requirements outlined above.

## 2.3 Methodology

To collect relevant resources, we conduct an exploratory approach to define the research field and, for the main part, follow a semi-structured literature review (Snyder, 2019). This allows the consideration of interdisciplinary literature to (1) specify a comprehensive set of aspects of trustworthiness in AI, (2) synthesize available knowledge regarding these aspects that is relevant when aiming to design, implement, and evaluate trustworthy AI, and (3) identify open challenges and knowledge gaps in these regards. Thus, we conducted exploratory literature research on trustworthy AI in general, which resulted in the six herein-discussed requirements. Then, the following procedure was completed for each requirement: (1) a Scopus search for conference papers, articles, and reviews, (2) the screening of the 100 most relevant abstracts as ranked by Scopus, (3) the screening of the remaining papers, and the extraction of relevant content, (4) snowballing and additional search in Google Scholar to close information gaps.

The goal of this paper was not to cover all existing publications of the field but rather to generate a comprehensive understanding of relevant research directions and their existing challenges. Therefore, we excluded articles with over-specialization, such as solutions only applicable to specific use cases or domains, and articles with limited contributions. This resulted in a collection of 183 papers as illustrated in Figure 3.



**TABLE 1** Problem definitions of the six trustworthy AI requirements investigated in this paper.

Trustworthy AI requirement	Problem definition
Human agency and oversight	Sustaining the autonomy of humans affected by AI systems
Fairness and non-discrimination	Mitigating bias in AI decisions to prevent unfair treatment
Transparency and explainability	Improving the understandability of an AI system and its decisions
Robustness and accuracy	Sustaining the model's performance (in unexpected circumstances)
Privacy and security	Protecting personal information of users in AI systems
Accountability	Understanding who is responsible for the decisions of AI systems

### 3 Overview and discussion of trustworthy AI requirements

In the following, we discuss six requirements an AI-based system should meet to be considered trustworthy. Each requirement is first defined, then we describe methods to establish and evaluate it, and finally, we debate open issues and research challenges. Table 1 provides a complementary illustration of the problem definitions of these six requirements.

#### 3.1 Requirement 1: human agency and oversight

##### 3.1.1 Definition of human agency and oversight

The principle of human agency and oversight refers to the idea that AI systems should uphold individual autonomy and dignity, and need to operate in a way that allows for substantial human control over the AI system's impact on people and society. This principle further postulates that AI systems should contribute to a democratic, flourishing, and equitable society and allow for human supervision to foster fundamental rights and ethical norms (High-Level Expert Group on AI, 2019).

Although the terms *human agency* and *human oversight* are very alike and sometimes used as synonymous, they are not interchangeable (Bennett et al., 2023). In this paper, we understand the term and concept *human agency* as referring to the very broad idea that humans as intentional actors should be in control, particularly with respect to substantial and important parts of their lives (High-Level Expert Group on AI, 2019). AI systems shall not restrict this agency; rather, it would be desirable that through AI systems, human agency is increased. AI systems could, for instance, limit human agency by deceiving or manipulating users. However, users should be able to influence automated decisions and to fairly evaluate or question the AI system. Consequently, users who are impacted by AI systems or who oversee AI systems need to be able to acquire or to be equipped with related competencies and

skills (AI Literacy) to understand and engage with AI systems to a satisfying degree (Long and Magerko, 2020; Pammer-Schindler and Lindstaedt, 2022).

The term and concept of *human oversight* is more specifically related to how AI systems are used and suggests that AI systems do not operate entirely autonomously, but instead, humans should oversee the way AI systems "work" within a more extensive process. This concept, therefore, is concerned with forms of human-AI interaction or collaboration, postulating that humans should be in a supervisory and decision-making role. Human oversight activities include observing, interpreting, and interfering in AI operations. Human oversight can be understood as a specific approach to facilitating human agency.

Long and Magerko (2020) define AI literacy as "a set of competencies that enables individuals to evaluate AI technologies critically; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace." Being AI literate means having a basic understanding of AI that empowers users to better interact with AI systems as they are able to judge the outcomes provided and, at the same time, to retain autonomy and agency (Hermann, 2022). Facilitating AI literacy on a large scale is currently a subject both of research (Ng et al., 2021) and of public endeavor.

##### 3.1.2 Methods to establish human agency and oversight

Very broadly speaking, socio-technical and human-centric design methods (Baxter and Sommerville, 2010) are approaches for designing (AI-based) systems that can systematically consider human users and people impacted, from the very early stage of designing the systems (Dennerlein et al., 2020). The consideration of human factors in the design phase of the AI-lifecycle (see Figure 2) responds well to the complex dynamics of the issues. Humans can interact with AI in various ways, which may require different levels of human agency and oversight. Anders et al. (2022) propose to think of different patterns of human engagement with AI-based operations and decisions as follows, sorted along the decreasing involvement of humans:

- *Human-in-command*: Humans manage and oversee an AI system's overall operation, including its wider impact on society, economy, law, and ethics. Decisions at a high level are made on when and how to use an AI system (Anderson and Fort, 2022). For example, regulatory bodies set ethical guidelines for the deployment of AI systems in healthcare to ensure patient safety and data privacy.
- *Human-in-the-loop*: Humans can intervene in AI-based decisions as well as in different steps of the underlying (typically machine learning based) algorithms (Mosqueira-Rey et al., 2023; Munro, 2021). In medical diagnostics, for example, an AI system can suggest potential diagnoses, but a physician reviews and decides on the diagnosis and treatment plan. Note that often, this kind of human-computer collaboration is not possible or even desired (High-Level Expert Group on AI, 2019).

- *Human-on-the-loop*: Humans can intervene through (re-)designing the AI system or through involvement in system operations, where their role is primarily on monitoring (Anderson and Fort, 2022). For example, a human operator may remotely monitor the performance of an autonomous vehicle and intervene only in emergency situations.
- *Human-out-of-the-loop*: Humans do not intervene. This could mean allowing an AI-based system to work without human involvement for specific tasks or in a completely automated manner. For example, in fully automated sections of vehicle production assembly lines, car parts can be assembled without human interaction.

Other authors categorize the spectrum of involvement in the design and operations of AI-based systems differently. For instance, Fanni et al. (2020) suggests a distinction between active and passive agency. Passive agency occurs when there are limited or no communication features that provide explanations for the decisions made by the AI system, or when users are uninformed about the potential consequences of AI interventions. This passive agency relates to the concepts of human-out-of-the-loop and human-on-the-loop. In contrast, active agency refers to situations where humans play a critical role in the design and operations of the AI-based system. This relates to human-in-the-loop and human-in-command approaches. Wang B. Y. et al. (2023) proposes the level of involvement to be defined by decisions and actions undertaken by humans and by AI-based systems. The author provides an example pattern of human-AI interaction as “AI Suggests, Human Intelligence (HI) Decides”, which can be interpreted such that the AI is providing recommendations, but humans are taking the role of final decision makers. Overall, thinking about such patterns of human-AI interaction allows deciding—at the time of designing and using an AI-based system—what kind of interaction is desirable or possible. Finally, we introduce the notion of AI literacy as positively contributing to human agency and oversight. This applies to both the users of AI-based systems and the decision-makers responsible for regulating and deciding which systems are used and how or what AI-related competencies users need to deal with AI systems.

AI literacy can be obtained in two dominant ways. First, through education in AI, particularly about everyday activities and technology (Zimmerman, 2018). Such education can also be mediated by technology. For instance, researchers are engaging young learners in creative programming activities, including AI (Kahn and Winters, 2017; Zimmermann-Niefeld et al., 2019). Secondly, everyday AI systems could be designed to support users in being or becoming AI literate. Long and Magerko (2020) outlined 15 concrete design considerations to promote users’ understanding and learning when interacting with AI systems. For example, AI systems could provide visualizations and explanations of decision-making processes to enhance users’ comprehension. They could also offer users the opportunity to learn about the system’s reasoning processes by putting themselves “in the agent’s shoes”; to encourage users to investigate the used data in terms of source, data collection processes, and known limitations or encourage users “to be critical consumers of AI

technologies by questioning their intelligence and trustworthiness”. Such support of AI literacy by AI systems is precious, as complex knowledge is highly context- and activity-dependent, and transferring knowledge from one context or activity to another can be quite challenging (Eraut, 2004).

Overall, systems need to be designed to be understandable for humans (Long and Magerko, 2020; Ng et al., 2021). This relates to the long-standing concepts of usability and learnability of systems. Finally, some of the above-described concepts, such as explaining decision-making, are essential for supporting transparency and explainability (see Section 3.3).

### 3.1.3 Methods to evaluate human agency and oversight

Building upon the previous discussion of human agency, human oversight, and AI literacy and their interrelations, we propose considering evaluation as moving upward the hierarchy of dependencies:

- *AI literacy*: AI literacy of relevant stakeholders is considered a prerequisite for human agency and oversight. This level can be assessed, for instance, through knowledge or competency tests, by checking certifications and formal degrees, or by investigating educational opportunities that relevant stakeholders have accessed or utilized.
- *System understandability*: It is critical to evaluate whether the AI-based system or functionality communicates understandably how it operates and what effects user actions might have. This relates to established concepts of usability and learnability and newer, AI-specific concepts like transparency and explainability (see Section 3.3).
- *Human oversight*: It needs to be established whether and how the intended interaction pattern of human oversight is present in the evaluation setting. This evaluation level concerns observing the designed activity, with a focus on establishing whether humans are, reasonably and in an engaged manner, involved in the process, either in-the-loop, on-the-loop, or in-command.
- *Human agency*: This is probably the most challenging concept to verify. One could argue that the existence and evaluation of a human-centric and socio-technical (AI) design process implies a certain level of human agency. The discussions led during this process could provide insights into how human agency is conceptualized and implemented in the AI-based system. In the inverse, it may be possible to establish its absence, i.e., when it becomes noticeable that human agency is limited through an AI-based system or functionality.

### 3.1.4 Open issues and research challenges

Challenges concern the *conceptualization of human agency, oversight, and AI literacy* as interwoven concepts, and the operationalization in design patterns of (interfaces for) AI. Such developments will need to be made in relation to maturing technology such as increasing shop-floor automation at the

conjunction of Internet of Things and AI-enabled data analytics, or the usage of generative AI across many sectors of knowledge-based and creative work. Additionally, more and better synthesized *design-oriented knowledge that captures how to design for human agency, oversight, and AI literacy* is pending. To date, research is revisiting the value of these design principles, e.g., in the context of AI (Long and Magerko, 2020) or generative AI (Simkute et al., 2024; Weisz et al., 2024). However, these design principles need clearer examples of how to operationalize them concretely within applications. Finally, due to the unclear conceptualization and the plethora of different designs with little grouping into overarching design patterns, *the evaluation of human agency, oversight, and AI literacy* will remain challenging. Evaluations will also need to uncover how these concepts interact with design patterns, ethics, and trust in AI systems (High-Level Expert Group on AI, 2019), actual decisions made in the domain of interest, and the overall socio-technical system performance (i.e., how good are decisions in the broader context and for whom).

## 3.2 Requirement 2: fairness and non-discrimination

### 3.2.1 Definition of fairness and non-discrimination

As AI products are being increasingly used in various fields and domains, their influence and impact on society are discussed not only in the machine learning community (e.g., Righetti et al., 2019) but also among the general public. AI may negatively impact individuals and society by reproducing existing societal stereotypes that can adversely affect vulnerable groups (Dubal, 2023). The unjust treatment of specific populations or individuals is particularly concerning in sensitive fields such as criminal justice, employment, education, or health, as it can result in significant consequences such as being refused medical care (Seyyed-Kalantari et al., 2021) or educational opportunity (Chang et al., 2021). Previous instances of such misconduct have been documented, including Google Ads showing lower-paid jobs to women and minority groups,<sup>2</sup> Apple Card granting lower credit limits to women than equally qualified men,<sup>3</sup> and commercial facial recognition systems performing poorly for women with black skin (Buolamwini and Gebru, 2018).

People's perception of fairness strongly depends on the context, which can include various factors, such as socio-political views, personal preferences, or the particular context and use case (Saxena, 2019). Especially in AI systems, achieving fairness is a multifaceted problem. Algorithmic fairness describes the absence of bias in AI decisions that would favor or disadvantage a person or group in a way that is considered unfair in the context of the application (Ntoutsis et al., 2020; Srivastava et al., 2019). Bias, often also called “discriminatory” or “unfair” bias, refers to outcomes of disproportionate advantage or disadvantage for a specific group of individuals, i.e., “*systematic discrimination combined with an unfair outcome is considered bias*” (Bird et al., 2019). Consequently, we

refer to fairness as the absence of discriminatory or “unfair” bias toward individuals, items, or groups.

Although ethical concerns are often at the forefront of public discourse, “unfair” bias can significantly impact society and businesses, even in seemingly non-critical domains. Therefore, it is crucial to consider the various risks from a business perspective. According to a report by Fancher et al. (2021), biased AI bears risk for several negative consequences. These include missing out on potential business opportunities, damaging reputation, and facing regulatory and compliance issues. One example of missing out on opportunities is when a recommender system only benefits a particular user group (Kowald and Lacic, 2022). While the members of the advantaged group may find the system useful, other groups don't experience the same level of system performance and stop using the product. This results in a loss of potential customers for the platform provider. Another consequence of biased AI is reputational damage, especially when the technology fails to address sensitive societal issues. For instance, using face recognition software that only works well for parts of the ethnicities in the user population will likely lead to negative public perception and backlash against the company. Finally, in cases where anti-discrimination laws govern the use of AI, such as in the job market, unfair algorithms can lead to legal problems. For example, an HR system that discriminates based on gender, age, or race can result in fines and penalties for the company. More information on accountability can be found in Section 3.6.

Furthermore, the issue of bias and fairness is complex because bias is naturally inherent in human behavior (Houwer, 2019) and thus, “unfair” bias can be introduced in every stage of the AI-lifecycle, as illustrated in Section 2. This problem becomes even more challenging in evolving AI systems because they can potentially reinforce bias between the user population, data, and algorithm (Baeza-Yates, 2018). Thus, monitoring and addressing “unfair” bias throughout the entire AI-lifecycle is essential.

### 3.2.2 Methods to establish fairness and non-discrimination

A wide range of methods has been proposed to increase fairness in AI models (Bellamy et al., 2019; Barocas et al., 2021). Because fairness strongly depends on the context, making AI models fair means making them fair in a particular context, i.e., according to an appropriate definition of fairness (Srivastava et al., 2019). Depending on their application level in the AI-lifecycle, “bias mitigation approaches” or “fairness enhancing methods”, are commonly grouped into three categories (Bellamy et al., 2019; Pessach and Shmueli, 2023; Barocas et al., 2023):

1. *Pre-processing* concerns improving the training data's quality and balancing its composition in regard to protected groups. This can be applied independent of the AI algorithm. Pre-processing can regulate fairness in acceptance rates but does not cater to other fairness constraints. Examples of pre-processing algorithms include reweighting (Calders et al., 2009), optimized preprocessing (Calmon et al., 2017), learning fair representations (Zemel et al., 2013), and disparate impact removal (Feldman et al., 2015).

<sup>2</sup> <https://incidentdatabase.ai/cite/19/#r184>

<sup>3</sup> <https://incidentdatabase.ai/cite/92/#r2037>

2. *In-processing* describes the design and optimization of an AI algorithm toward an explicitly defined, fair solution. Thus, it incorporates fairness in the training algorithms themselves and can only be applied to specific algorithms/models that are well-understood (Srivastava et al., 2019).
3. *Post-processing* aims to adapt the AI model's results toward a balanced distribution for protected groups. Examples thereof include methods for calibration, constraint optimization, or setting thresholds for the maximum accuracy differences between groups (Pleiss et al., 2017).

While each approach has its particular pros and cons, all of them potentially negatively affect the models' accuracy (see Section 3.4).

### 3.2.3 Methods to evaluate fairness and non-discrimination

The auditing or evaluation of algorithmic fairness can, similar to the mitigation strategies, be approached according to the three main phases of the AI-lifecycle, i.e., design, development, and deployment (Koshiyama et al., 2021). Examples of what can be assessed are (1) population balance and fair representation in data (design), (2) the implementation of fairness constraints in modeling or the adherence to fairness metrics in evaluation (development), and (3) the adherence to fairness metrics in real-time monitoring (deployment) (Akula and Garibay, 2021).

Measures of algorithmic bias are a quantitative evaluation of the result set of the system at hand (Pessach and Shmueli, 2023). The highest level of separation between different definitions of fairness is between individual and group fairness, which are related to the legal concepts of disparate treatment and disparate impact, respectively (Barocas and Selbst, 2016).

- *Individual fairness* considers fairness on an individual level and requires treating similar individuals similarly.
- *Group fairness* calculates fairness on a group level, requiring different groups to be treated equally.

Furthermore, one can differentiate between three principal approaches: fairness in acceptance rates, fairness in error rates, and fairness in outcome frequency (Barocas et al., 2021). Verma and Julia (2018) provides an overview of the 20 most prominent definitions. One challenge, however, is selecting the "right" definition and metrics, as many different definitions of algorithmic fairness and related metrics exist. In many settings, these definitions contradict each other—thus, it is usually not possible for an AI model to be fair in all three aspects. The appropriate metrics must be selected for a given application and its particularities. Several software packages are available that implement important metrics. Popular open-source frameworks include AIF 360 (Bellamy et al., 2019), Fairlearn (Bird et al., 2020), and Aequitas (Saleiro et al., 2018).

### 3.2.4 Open issues and research challenges

Fairness is a concept highly context-dependent that, in practice, may require ethical consultation (John-Mathews, 2022).

No fairness definition applies to all contexts, and it seems an intrinsic—and unsolvable—challenge of the field to formulate legally compliant measurements mathematically (Wachter et al., 2021). What is perceived as fair or unfair also varies between different cultural and legal settings. It remains unclear how to tune the fairness of an AI application intended to be used in multiple cultural or legal contexts (Srivastava et al., 2019) and, more generally, how to apply and assess existing regulations, standards, and ethical constraints in practice (Costanza-Chock et al., 2022).

From a more technical perspective, ensuring fairness when combining different AI components poses a significant challenge. This can be particularly difficult when reusing AI tools or algorithms with limited access to code, or when exchanging data audited only for a specific use case or application context. In fact, it has been shown that measures of algorithmic fairness are sensitive to any alterations in the input data and to even simple changes in train-test splits (Friedler et al., 2019). In principle, monitoring fairness in AI systems that are in production is possible (e.g., Vasudevan and Kenthapadi, 2020). However, it is still much more demanding to define when fairness criteria are met and when not because the algorithm's performance may change over time (e.g., Lazer et al., 2014). The application of generative AI models presents additional challenges, especially in the context of language. Despite a significant body of research, it is still unclear how to effectively measure and evaluate their bias and how to transform these measurements to be suitable for application in various contexts or to consider different minority characteristics (Nemani et al., 2023).

Finally, there is no standard to determine the adequate trade-off between different fairness metrics nor between fairness and accuracy. While there have been attempts to show that the fairness-accuracy trade-off is rather an issue of historical bias in data (Dutta et al., 2020), it remains unclear how to generate an ideal, unbiased dataset as a standard in practice.

## 3.3 Requirement 3: transparency and explainability

### 3.3.1 Definition of transparency and explainability

Transparency and explainability are two related but distinct concepts. Explainability aims to enhance comprehension, build trust, and facilitate decision-making (Adadi and Berrada, 2018). In contrast, the goal of transparency is to ensure understandability and accountability (Lepri et al., 2018; Arias-Duart et al., 2022; McDermid et al., 2021). With transparent and explainable AI, users can better estimate the trustworthiness of AI systems since they can understand their inner workings and, consequently, their opportunities and limitations (Naiseh et al., 2023). However, there is no final consensus on the definition and scope of transparency to date. Therefore, in this paper, we will focus on explainability as a means to achieve transparency in AI models.

Furthermore, including explainability methods in an AI system offers additional benefits, such as allowing for debugging and expert verification of the AI system, which can foster task accuracy and efficiency (Weber et al., 2023) [e.g., in healthcare (Albahri et al., 2023; Hulsén, 2023)]. For example, Anders et al. (2022) improved



a model's prediction accuracy by utilizing explainability methods to identify dataset samples that led the model to learn spurious correlations. Young et al. (2019) used explainable AI to help experts verify that models for melanoma detection relied on the correct data aspects for their predictions. Other researchers have developed explainability methods to facilitate iterative learning and improvement of AI models by identifying patterns, biases, or errors in the model's decision-making processes (Ribeiro et al., 2016; Mehdiyev and Fettke, 2021).

### 3.3.2 Methods to establish transparency and explainability

Methods to make AI systems explainable are often summarized under the umbrella term “Explainable AI (XAI)” (Holzinger et al., 2022), and sometimes also interpretable AI (Molnar, 2020). XAI aims to increase transparency and explainability in AI systems to ensure trust, understanding, and accountability. With respect to the AI-lifecycle (see Section 3), XAI is relevant in all three phases: in the design phase, for understanding and incorporating stakeholders' explainability requirements; in the development phase, for understanding important data aspects, performing error analysis and model refinement; and finally, in the deployment phase for continuous model verification and enhancing user trust. XAI methods vary in their approach toward achieving these goals and can be categorized based on their *scope* and *model dependence*. The term *scope* refers to whether XAI methods can produce global or local explanations (Samek et al., 2021). Global explanations aim to explain the AI model as a whole and, thus, provide insight into the model's general decision-making process, e.g., SHAP (Lundberg and Lee, 2017). In contrast, local explanations focus on the model's decision-making process in regard to a single sample, making the explanation more specific [e.g., LIME (Ribeiro et al., 2016)]. Global explanations tend to be computationally more expensive, as they need to consider the entire training input space, whereas local explanations might work on one input sample.

Furthermore, XAI methods can be classified as “model-agnostic” or “model-specific” methods (Arrieta et al., 2020), according to their *model dependence*. Model-agnostic XAI methods can provide explainability and transparency for any AI system. In contrast, model-specific XAI methods are tailored to a single AI architecture, which may limit their compatibility with other AI systems. However, they tend to create more accurate and translucent (i.e., the extent to which the explanations rely on particularities of the inner workings of the AI system) explanations than model-agnostic methods (Carvalho et al., 2019). Additionally, model-specific XAI methods can not be implemented for every AI system. Meanwhile, model-agnostic XAI methods can be implemented on every AI system, as they often use the AI system as an oracle. They do so by probing the model many times to estimate the effects of the input on the model prediction, which can lead to expensive computations. Examples of model-agnostic XAI methods are SHAP (Lundberg and Lee, 2017), LIME (Ribeiro et al., 2016) and the broader category of counterfactual explanations (Guidotti, 2022); model-specific XAI methods are DeepLIFT (Shrikumar et al., 2017) and Integrated Gradients (Sundararajan et al., 2017).

It is a common belief that “white-box” models—models whose “inner workings” can be inspected—are immediately interpretable and transparent. However, they often have lower prediction accuracy than more complex models (Moreira et al., 2022). In addition, white-box models, e.g., Linear Regression and Decision Trees, often need extra steps to be used or treated as “full-fledged” XAI methods because in order to be most effective, XAI methods are required to be understandable and not overwhelming to their target users, meaning that they are specifically tailored to meet their requirements (Miller et al., 2017). Social science research regards high-quality explanations as a form of conversation and proposes explanation theories like temporal causality, social constructivism, and attribution theory (Mendoza et al., 2023).

Deep learning models have succeeded across various domains by utilizing computational units called neurons, ordered in sequential layers forming neural network (NN) models. NNs can autonomously learn meaningful internal features without manual feature engineering (LeCun et al., 2015). Consequently, to train the models, we often use raw data directly or include all features, regardless of complexity (Roy et al., 2015). However, while NNs map inputs directly to outcomes, they do not disclose how features are weighted in relation to the model's output (Zhao et al., 2015).

XAI methods can provide explanations in many different ways. *Feature attribution methods* generate values for each input feature, highlighting its importance to the AI model's predictions. However, these methods can be sensitive to input noise and correlated features, resulting in misleading conclusions (Adebayo et al., 2018). They are commonly presented textually (numerically), via bar charts (Ribeiro et al., 2016), or via heatmaps (Sundararajan et al., 2017). XAI methods can also provide explanations by visualizing the models' internals, e.g., activation maps in Convolutional Neural Networks, which can quickly become too complex when dealing with many neurons (Carter et al., 2019). *Counterfactual explanations* are a category of explanations that aims to answer the “why” question with “because if it was *something different*, it would be *this other thing* instead” (Guidotti, 2022). Counterfactuals are typically computationally intensive, and generating meaningful counterfactual examples depends on the task context (Artelt and Hammer, 2019).

Undoubtedly, explaining AI requires numerous considerations. Despite inherent limitations, each explanation technique enhances the explainability and transparency of AI systems, thereby advancing the overarching objective of fostering trustworthy and accountable AI applications. Many software libraries exist that make employing XAI methods straightforward. For Python, the libraries SHAP (Lundberg and Lee, 2017), LIME (Ribeiro et al., 2016), Captum (Kokhlikyan et al., 2020), and scikit-learn (Pedregosa, F. et al., 2011) are widespread and cover most XAI method categories. The DALEX (Biecek, 2018) library offers model-agnostic explanations for the programming language R.

### 3.3.3 Methods to evaluate transparency and explainability

Evaluating explanation methods is vital for assessing their correctness, efficacy, and practical utility. Various approaches for estimating the effectiveness and quality of explanations have been

introduced and can be divided into the following categories (Doshi-Velez and Kim, 2017):

- *Application-grounded evaluations* that involve human participants performing realistic tasks and offer insights into how XAI methods work in real-world scenarios.
- *Human-grounded evaluations* that use simplified tasks for human participants to assess the comprehensibility and usefulness of explanations provided by the AI systems.
- *Functionally-grounded evaluations* that rely on proxy tasks without human involvement, focusing on XAI algorithms' functionality and their scores against a pre-defined metric of interpretable quality.

While application and human-grounded approaches focus on the plausibility and usefulness of explanations to users, functionally-grounded evaluations estimate the *correctness* of XAI algorithms. Various properties of explanation methods can be examined to determine if they function correctly (Hedström et al., 2023). Some of these properties are:

- *Faithfulness* (Alvarez-Melis and Jaakkola, 2018b; Samek et al., 2016; Šimić et al., 2022) estimates how accurately explanation methods identify features in the input driving the model prediction.
- *Robustness* (Montavon et al., 2018; Alvarez-Melis and Jaakkola, 2018a) measures an explanation method's sensitivity to input perturbations.
- *Localization* (Selvaraju et al., 2017; Fong and Vedaldi, 2017) identifies if the explanation method focuses correctly on the desired regions of interest.
- *Complexity* (Bhatt et al., 2020; Nguyen and Martínez, 2020) measures the conciseness of explanations, where less complex explanations are deemed more interpretable than more complex ones.
- *Randomization* quantifies an explanation method's sensitivity to modifications of model parameters.
- *Axioms* (Adebayo et al., 2018; Kindermans et al., 2019) define criteria that an explanation method has to fulfill.

Hence, careful identification of evaluation aspects is necessary to address context-specific concerns, such as faithfulness, robustness, or comprehensibility. For a detailed overview of evaluation metrics for transparency and explainability, please also see Hulsén (2023), in which metrics such as simulatability, decomposability, coherence, or comprehensiveness are mentioned. Unfortunately, software libraries that offer metrics for validating explanation methods are scarce; among the few existing ones are Quantus (Hedström et al., 2023) and AI Explainability 360 (Arya et al., 2019).

### 3.3.4 Open issues and research challenges

The requirement for transparency and explainability of AI faces several open challenges. First and foremost, the research community needs to fully agree on a common, clear, and precise definition for transparency in AI systems, which currently leads to

ambiguity regarding what explanations should entail. For instance, properly calibrating AI explanations to instill the correct amount of trust in AI models is crucial but complex (Wang et al., 2019), as it requires a balance between providing understandable insights without oversimplifying or overwhelming users and, at the same time, without over or underselling the explained AI model's capabilities. Additionally, tailoring explanations for diverse user groups and individuals remains challenging, as different stakeholders require different explanations at varying levels of granularity and detail (Miller et al., 2017; Mendoza et al., 2023). Furthermore, the evaluation of transparency and explainability of AI models is challenging, and developing intuitive user interfaces for explanations poses a design challenge, requiring informative yet user-friendly interfaces that follow "XAI UI guidelines" (Liao et al., 2020; Wolf, 2019).

Finally, ensuring transparency and explainability in large language models and generative AI systems presents unique difficulties due to their complexity, and it is also unclear how their explanations should look like (Schneider, 2024).

## 3.4 Requirement 4: robustness and accuracy

### 3.4.1 Definition of robustness and accuracy

Robustness and accuracy are key properties of any AI system, and ensuring them is an essential part of the AI model development. Robustness and accuracy—in loose terms—refer to how "adequate" or "correct" the outputs of an AI model are. In contrast to other requirements—such as fairness or transparency—sufficient robustness and accuracy are required for *any* AI model, independent of its specific purpose (Huber, 2004).

AI training algorithms are typically designed for general problem settings (e.g., image-classification tasks). A specific AI model is then developed for a particular problem. In many problem settings, various models can be employed. However, in complex settings, identifying a suitable model becomes challenging, often requiring the application of a model despite uncertainties regarding its suitability. In such a case, it is important to use models that produce reasonable results even if they are used in settings they were originally not designed for. Under challenging conditions, some models may behave unpredictably and produce unstable outputs, whereas other models exhibit more constant behavior, and the quality of their outputs differs only slightly from the optimal setting. Clearly, it is preferable to use the latter class of models. More concretely, Huber (2004) outline three key properties an AI model should ideally possess. The model should: (1) achieve optimal or near-optimal results if it is applied in exactly the setting it was designed for; (2) degrade the quality of the results only slightly if it is subject to small deviations from the assumed setting; (3) not trigger nonsensical or dangerous outputs if it is applied in settings with large deviations from the assumed setting. An AI model is considered *robust* if it meets properties (2) and (3), while property (1) is essential to ensure a sufficiently *accurate* model.

It is optimal to achieve high values in both, accuracy and robustness, but this is rarely possible. Instead, there is typically a trade-off between robustness and accuracy. While some models

are more robust and can be applied across different settings, they come at the cost of lower accuracy. For example, a face detection algorithm could have very high accuracy in a highly specific setting, e.g., fixed camera type, fixed angle, and fixed lighting conditions, but might fail as soon as one of those parameters changes. A different model, on the other hand, might show slightly lower accuracy but perform similarly in various settings, e.g., different camera types and lighting conditions.

Considering robustness and accuracy are crucial in all phases of the AI-lifecycle (see Section 2). In the design phase, it is important to make choices that do not compromise the accuracy of the model, for example, in the selection of appropriate training and testing data. The development phase is particularly essential, where one of the core tasks of every AI system development is to ensure these qualities. Finally, in the deployment phase, robustness and accuracy metrics need constant monitoring—especially in the case of continuously changing systems—to ensure a well-working AI model (Hamon et al., 2020).

### 3.4.2 Methods to establish robustness and accuracy

Ensuring accuracy is the core of AI model development and part of all best practices. Additional core considerations are the appropriate choice of (1) target metrics the model is trained on, (2) data splitting techniques (e.g., train-test-splits), and (3) model selection methods. These three points are part of AI model developments that, in practice, are often done properly but not documented in a sufficient manner. To generate trust in AI models, it is essential to both deliver quality and document all relevant choices made in the process. These include, for example, the choice of suitable evaluation metrics, i.e., not only is it necessary to document the choice (e.g., “F1-score”), but also the reasoning for that choice (e.g., “classification problem with unbalanced data”) (Huber, 2004; Hamon et al., 2020).

Ensuring robustness can be done in two principal ways: (1) by restricting potential models to model types shown to be more robust (e.g., multilinear regression is generally more robust than deep learning), or (2) by explicitly evaluating model robustness and incorporating it in the model selection process. In the process of model selection, certain model types can be adapted to increase robustness. For example, fragile models can improve robustness by introducing mechanics that ignore certain data points or limit their effect, e.g., a drop-out layer in neural networks (Krizhevsky et al., 2012), or thresholds (Kim and Scott, 2012). This has the advantage that the model learns to rely less on specific data points and focus more on the general information depicted in the majority of the data. However, “reserved” data usage has a cost: The model has less data to work with, which puts it at a statistical disadvantage compared to fragile models that use all data. This cost is particularly high when the model is used in a setting where accuracy is more important than robustness (Fisher, 1922).

### 3.4.3 Methods to evaluate robustness and accuracy

In the AI literature, there are many different forms of robustness, e.g., robustness to domain shift (Blanchard et al., 2011;

Muandet et al., 2013; Gulrajani and Lopez-Paz, 2020), adversarial robustness (Nicolae et al., 2019; Xu et al., 2020), robustness to noise (Zhu and Wu, 2004; Garcia et al., 2015), robustness to non-adversarial perturbations (Hendrycks and Dietterich, 2019; Rusak et al., 2020; Scher and Trügler, 2023), and others. While some generic robustness scores have been proposed (Weng et al., 2018; Sharma et al., 2020), they do, in fact, only measure specific types of adversarial robustness. Therefore, there is no single unified notion of robustness. Accordingly, as with most other aspects of trustworthy AI, the type of robustness to be considered depends on the context.

Accuracy can be measured with a wide variety of metrics. In technical terms, accuracy is simply the fraction of correct outputs of an AI model (Naidu et al., 2023). The goal of measuring accuracy is to measure how “good” or “correct” the outputs of the AI model are. How “good” or “correct” is defined in practice and highly depends on the application at hand. Therefore, there is no single generally applicable accuracy metric. The choice of which metric or metrics are applicable depends on the type of problem the AI model attempts to solve (e.g., regression, classification, ranking, translation), and on the particular properties of the application. For instance, for classification tasks with balanced classes, accuracy is a useful metric. However, classification tasks with highly imbalanced data, this can be misleading, and metrics such as precision and recall are more appropriate. Regression tasks require very different evaluation metrics than classification tasks. Examples are root-mean-square error or mean absolute error. Which one is more appropriate again depends on the application at hand. The same holds for ranking tasks, which require yet other types of metrics (such as mean reciprocal rank or mean average precision). An overview of these commonly used metrics can be found in Poretschkin et al. (2023).

### 3.4.4 Open issues and research challenges

The ongoing surge in generative AI models has opened a new challenge for existing models with respect to accuracy and robustness. It has been shown that generative models are, to a certain extent, capable of producing adversarial examples that cause catastrophic outputs in existing fragile AI models (Han et al., 2023). Moreover, some of these examples are transferable from one system to the next and can be re-used to cause failure in a number of different AI systems (Wang Z. et al., 2023). Currently, there is no widely applicable easy strategy to address these issues. Existing fragile models need to be replaced with more robust models, and higher accuracy needs to be established for robust models, which are currently not able to compete with their fragile counterparts.

Providing accuracy is the core of machine learning and AI, and thus, methods ensuring AI applications are accurate need to be integrated with the development and improvement of the models. Accuracy evaluation metrics are well-established in statistics and machine learning, and their computation is generally straightforward. However, the choice of a proper metric and the definition of its thresholds are much more complex. While best practices exist, no formal guidelines are available. Also, despite a wide range of established accuracy metrics, there is a need for additional, new accuracy metrics that are specifically developed and

tailored to the particularities of distinct applications (Naidu et al., 2023)—and perhaps also tailored to permit robust solutions.

The field of robustness faces a variety of open issues and limitations. Robustness is a broad concept that current research strands do not necessarily cover in all aspects. In part, no specific methods—besides best-practice examples—are available for increasing robustness (e.g., robustness against noise). Especially the problem of adversarial attacks is constituted in a “cat and mouse game”: if a specific attack strategy is known, AI systems can be made robust against it by incorporating the attack in the training procedure—known as adversarial training. However, this does not guarantee robustness toward a new attack that has not yet been part of the training. Another open issue is caused by composite systems, where multiple AI components are combined or in situations where AI evaluation is part of a larger product/solution. Moreover, yet unanswered, is how to assess accuracy and robustness in evolving (learning) AI systems that are constantly updated (in some cases with every single user interaction) (Hamon et al., 2020). In general, data quality, as well as model training and selection, are very important for AI systems, as these aspects influence accuracy and robustness, among other qualities such as fairness. Nonetheless, currently, no unified quality concept is available, even though basic automated tests are feasible.

## 3.5 Requirement 5: privacy and security

### 3.5.1 Definition of privacy and security

Privacy and security are indispensable pillars supporting the trustworthiness and ethical use of AI systems. Privacy refers to the data used as input for the AI model and to protecting information that belongs to the data owner. This information must not be disclosed to any third parties and may only be disclosed to parties that the data owner defines. Security, on the other hand, pertains to the AI model itself and is linked to defending it against any malicious attacks that aim to impact or manipulate it in an undesired or harmful way. Both privacy and security risks can arise along the whole life cycle of AI systems (see Section 2). Existing countermeasures span a broad spectrum, encompassing methods from manipulating the input data of AI models for ensuring privacy and security and designing AI models that are by themselves private and secure to recent advances that allow the protection of AI models during the inference process, i.e., during deployment (Elliott and Soifer, 2022).

If AI is utilized in critical areas such as healthcare, autonomous vehicles, or national security, it may even endanger human safety. Incidents like the unintended memorization of sensitive information by large language models<sup>4</sup> highlight the tangible privacy and security risks associated with AI. These examples serve as a reminder of AI models' potential to compromise privacy and security inadvertently. In response, regulatory bodies, particularly in the European Union, have been proactive in updating legal frameworks to address these challenges: The General Data Protection Regulation (GDPR) and the AI Act are prime

examples of such regulatory efforts, aiming to establish clear guidelines for AI design, development, and deployment (Zaeem and Barber, 2020).

### 3.5.2 Methods to establish privacy and security

Understanding the weaknesses of AI systems and identifying the diverse kinds of attacks on privacy and security is critical for developing defense strategies and, consequently, evaluating their effectiveness. Attacks can be classified based on several aspects, including the attacker's capabilities and the attack goal. For example, attackers can deviate from the agreed protocol (active/malicious) or try to learn as much as possible without violating the protocol (passive/semi-honest/honest-but-curious). Moreover, an attacker may be assumed to have finite or infinite computational power. Based on the attacker's knowledge, one can differentiate between black-box attacks (which only access the model's output), white-box attacks (which access the full model), and gray-box attacks (which gain partial access). In the following, we classify attacks based on the attack goal, i.e., evasion attacks, poisoning and backdoor attacks, and privacy attacks (BSI, 2022; Bae et al., 2018):

1. *Evasion attacks* (including adversarial attacks) aim to mislead AI models through carefully crafted inputs, forcing incorrect predictions.
2. *Poisoning attacks* corrupt the training process, while *backdoor attacks* insert hidden triggers into models.
3. *Privacy attacks* seek to extract sensitive information from AI models. The most common privacy attacks include:
  - Membership inference attacks aim to determine whether a specific data sample was used in the training phase of the AI model.
  - Attribute inference attacks aim to infer sensitive attributes, e.g., the gender of individual records.
  - Model inversion attacks aim to infer features that characterize classes from the training data.
  - Model extraction and stealing attacks aim to reconstruct the model's behavior, architecture, and/or parameters.

Adilova et al. (2022) list best practices to defend against the aforementioned attacks. In the following, we briefly provide examples for each class of attacks. Countermeasures against *evasion attacks* include: (1) certification or verification of output bounds, i.e., utilizing certification methods to calculate guarantees on the output distribution to certify the AI model's robustness, (2) adversarial retraining, i.e., incorporating perturbed samples into the training process, (3) injection of randomness into training, i.e., using random transformations to protect against attacks, (4) use of more training data, i.e., enhancing adversarial robustness with larger and more diverse training datasets, (5) multi-objective optimization, i.e., not only optimizing for accuracy but balancing between adversarial robustness and task-specific accuracy, and (6) attack detection, i.e., implementing detection methods for malicious inputs. The risk of *backdoor* and *poisoning attacks* can effectively be mitigated by the following strategies: (1) use of trusted sources, i.e., ensuring reliability and trustworthiness of data models;

<sup>4</sup> <https://incidentdatabase.ai/cite/357>



(2) random data augmentation, i.e., employing data augmentation techniques to mitigate the effect of poisoned samples; (3) use of an auxiliary pristine dataset, i.e., supporting training with trusted data to dilute the impact of poisoned samples; (4) attack detection, i.e., applying techniques to identify poisoned samples or models, including analysis of data distributions and model inspections; (5) model cleaning, i.e., utilizing methods like pruning, retraining, or differential privacy to eliminate the influence of triggers or poisoned data; and (6) adversarial training, i.e., adapting adversarial training approaches to counter poisoning attacks, enhancing model resilience.

Overall, defending the security of AI models against a variety of attacks involves a multifaceted approach that combines diverse techniques and practices, highlighting the need for AI practitioners to continuously assess and update their defense strategies. Similarly, the development of privacy-enhancing technologies (PETs) has been instrumental in protecting AI models from **privacy attacks**. The following (incomprehensive) list of PETs details the most important technologies, which currently form the forefront of research in private and secure AI computations.

- *Homomorphic encryption (HE)* (Gentry, 2009; Rechberger and Walch, 2022; Smart, 2016; Phong et al., 2018) supports the performing of certain operations on encrypted data (i.e., without decrypting it). This allows privacy in cloud-based AI services to be maintained without exposing private data or model details. However, HE requires substantial computational resources and entails high computational costs.
- *Secure multi-party computation (MPC)* (Rechberger and Walch, 2022; Evans et al., 2018) allows collaborative computation without revealing individual inputs. In the context of AI, MPC is especially useful for collaborative learning and private classification, though it requires significant communication overhead for many participants.
- *Differential privacy (DP)* (Dwork, 2008; Dwork and Roth, 2014) bounds the maximum amount of information that an AI model's output discloses about an individual data point by incorporating curated noise into the computation. Specifically, noise can be added either to the input data, during the training process, or to the output (Friedman et al., 2016). While effective in various AI applications, including deep learning (Abadi et al., 2016) and recommender systems (Müllner et al., 2023), DP's main challenge is the trade-off between privacy protection and accuracy.
- *Federated learning (FL)* (Zhang et al., 2021; Li et al., 2020) is a machine learning approach that allows multiple clients, like mobile phones, to collaboratively learn a model by training locally and sharing updates with a central server. This method enhances privacy by keeping data local, although there is a risk of data reconstruction from model updates (Yin et al., 2021; Nasr et al., 2019; Ren et al., 2022).
- *Synthetic data* (Slokom, 2018; Liu et al., 2022) mimics real data's statistical features to enable the AI model to still learn the real data's features, but without using the real data. This offers a way to preserve privacy in data sharing, yet it is not immune to reconstruction risks (Stadler et al., 2022).

- *Transfer learning* (Zhuang et al., 2020), while not a PET *per se*, contributes to privacy by fine-tuning pre-trained models on new tasks with minimal data, reducing the need for large private datasets (Gao et al., 2019). Similar ideas are also employed by PETs based on meta-learning (Muellner et al., 2021).

The described defense methods can also be combined to increase privacy and security. For example, DP can mitigate the risk of reconstruction in FL (Wei et al., 2020) and synthetic data (Tai et al., 2022; Stadler et al., 2022).

### 3.5.3 Methods to evaluate privacy and security

The vulnerability of AI models to privacy and security attacks can be assessed using two complementary approaches: mathematical analysis and attack-based evaluation. *Mathematical analysis* offers formal proofs of privacy and security features within a system, much like cryptography, guaranteeing system security under certain assumptions (e.g., DP). This method is crucial, especially when introducing new privacy or security techniques, as it requires thorough checks for implementation errors and the appropriate selection of parameters. On the other hand, *attack-based evaluation* gives us practical insight into how an AI model reacts to various attack strategies. This method tests the model's vulnerability to different attacks and determines its resilience by using various metrics (Wagner and Eckhoff, 2018; Pendleton et al., 2016). These metrics might include the attacks' success rate, the effort required to breach the model (measured in iterations), the precision of the attack, and the smallest necessary data alterations to compromise the model successfully (BSI, 2022). The choice of metrics depends on the nature of the attack and on assumptions about the attacker's skills and knowledge. It is tailored to each specific scenario and model based on potential threats and existing literature. However, it is important to acknowledge the limitations of attack-based evaluations. While they can pinpoint specific weaknesses and vulnerabilities, they do not offer a comprehensive guarantee of privacy or security. Additionally, these evaluations only cover known attack scenarios, leaving the potential for undetected vulnerabilities against new or complex attack techniques.

### 3.5.4 Open issues and research challenges

Despite the existing countermeasures, the AI privacy and security field still faces numerous unresolved challenges. Many defense strategies cannot fully mitigate the models' vulnerability to attacks, especially not to adversarial and poisoning attacks. Additional challenges emerge with the increasing advancement of generative AI, particularly in models that rely heavily on unstructured data such as text. For example, establishing clear boundaries on what constitutes private information becomes increasingly difficult due to the inherent complexities of unstructured data (Brown et al., 2022). PETs often introduce trade-offs, such as increased computational demands (HE and MPC) (Moore et al., 2014), reduced prediction accuracy and increased unfairness (DP) (Abadi et al., 2016; Bagdasaryan et al.,

2019; Müllner et al., 2024), or a surge in communication overhead while having no privacy guarantees (FL) (Almanifi et al., 2023; Bagdasaryan et al., 2020). Therefore, integrating PETs smoothly into AI systems without compromising performance remains complex and requires further research. Just as for fairness and robustness, evaluating privacy and security when combining multiple AI components is challenging. Adding components that protect against one identified risk can even introduce new vulnerabilities (Debenedetti et al., 2023).

In general, fostering secure model sharing and privacy-preserving collaboration, developing standardized evaluation metrics, and preparing for advanced AI threats necessitate a collaborative approach among researchers, developers, and policymakers. Ongoing research and shared best practices will be crucial for building a secure, privacy-conscious AI ecosystem.

## 3.6 Requirement 6: accountability

### 3.6.1 Definition of accountability

Another key requirement for trustworthy AI is accountability. At its heart, accountability is the obligation to notify an authority of one's conduct and to justify it (Bovens, 2007; Brandsma and Schillemans, 2012; Novelli et al., 2023; Hauer et al., 2023; Wieringa, 2020), whereas responsibility includes explicit obligations defined in advance (Bivins, 2006) and can be seen as a subcategory of accountability (Gabriel et al., 2021). Liability is closely related to accountability and means legal responsibility, including sanctions for misbehavior. In this article, we, therefore, see liability as a sub-concept of accountability and solely use the term accountability.

From a conceptual perspective, accountability can also be defined as a virtue or as a mechanism (Bovens, 2010). Accountability “refers to the idea that one is responsible for their action—and as a corollary their consequences—and must be able to explain their aims, motivations, and reasons”<sup>5</sup>.

The definition of Bovens (2007) is widely used as the basis for addressing accountability and identifies the following key elements of accountability: actor, forum, relationship between these two, account, and consequences. The actors, as natural persons, groups or organizations (e.g., developers, deployers, manufacturers, or users of AI systems), shall be able to explain their actions (e.g., used models and data, intended use, planned outcomes, and potential malfunctions of AI systems) by certain criteria to the forum (e.g., a court, a supervisor, an auditor), that can “pose questions and pass judgments”. The relationship between actor and forum can vary and involves individual, hierarchical, collective and corporate accountability. Finally, there will be consequences (e.g., fines for non-compliance with rules).

### 3.6.2 Methods to establish accountability

When addressing accountability features such as context, the range of actions taken, the acting entity, the forum as the bearer of interests and imposed standards, processes, and implications must be considered (Bovens, 2007) to be able to achieve compliance,

report, oversight, and enforcement (Novelli et al., 2023). Thus, accountability is always relational (Bovens, 2007), contextual (Lewis et al., 2020) and involves single persons, other entities, as well as groups and societies. Depending on impacts, different levels of accountability are required (Cech, 2021).

Accountability systems range from hard law regulations over functional roles within organizations (Novelli et al., 2023) to social norms that, in turn, form the basis for decision-making and behavior (Gabriel et al., 2021). In the field of software development, responsibility involves maintaining quality in the design process (Eriksen, 2002), implementing tools for characterizing system failure (Nushi et al., 2018), as well as using transparency and inspection mechanisms (Hauer et al., 2023). So called “algorithmic accountability” is also described as the expectation that people along the AI-lifecycle (see Section 2) will comply with legislation and standards to ensure the proper and safe use of AI and involves not only the use, design, implementation, and consequences but the whole “socio-technical process” (Hauer et al., 2023; Novelli et al., 2023; Wieringa, 2020). Thus, accountability shall ensure compliance with requirements such as fairness, transparency, and robustness (Durante and Floridi, 2022; Novelli et al., 2023). Therefore, it also requires that mechanisms for auditability, minimization, and reporting of negative impacts, trade-offs, and redress are in place. Therefore, accountability must be ensured along the whole AI-lifecycle—in the design phase, the development phase, and the deployment phase.

### 3.6.3 Methods to evaluate accountability

Due to the versatility of accountability, its evaluation is challenging. Numerous approaches for evaluating accountability regarding AI systems are put forth.<sup>6</sup> For example, Tagiou et al. (2019) suggest a “a tool-supported framework for the assessment of algorithmic accountability” that focuses on both algorithmic and organizational aspects and Cech (2021) proposes the “Accountability Agency Framework (A3)” as an analytic lens as a qualitative, explorative, and complementary tool to assess algorithmic accountability, which is based on Bovens (2007)'s definition of accountability. Their framework encompasses four steps: requesting information, providing account, imposing consequences, and effective change. Additionally, it provides a series of guiding questions for assessing algorithmic accountability (Cech, 2021). Xia et al. (2024) proposed a granular AI Metrics catalog that includes process, resource, and product metrics and is specially designed for generative AI. Besides, numerous other, mainly contextualized frameworks, which range from accountability in organizations (Buhmann and Christian, 2019), public reason (Binns, 2018) and public service (Brown et al., 2019) to “AI robots accountability” (Toth et al., 2022) frameworks have been proposed. From a qualitative perspective, approaches that, for instance, take human rights into account are discussed (McGregor et al., 2019).

In general, Brandsma and Schillemans (2012) suggest a so-called “accountability cube” as a quantitative assessment tool for assessing accountability, considering three dimensions

<sup>5</sup> <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>

<sup>6</sup> In this article, we only present individual approaches as examples without any claim to completeness.

of accountability processes: information, discussion, and consequences/sanctions. Accordingly, accountability is “high” if there is much information, intensive discussions, and several opportunities to impose consequences. This approach can be applied in various contexts along the AI-lifecycle. Without any closer examination of the approach itself, we use the accountability cube to exemplify possible evaluation criteria of algorithmic accountability.

To start with, much information is given if people are aware of basic technical outlines, chances, and risks of AI in the respective context and know their obligations along the lifecycle of AI, including, for instance, information, documentation, and risk-assessment obligations. From our point of view, AI literacy is essential to this. Discussion is intensive if an informed exchange of views on AI systems and regulation—whether formal or informal—takes place between multiple stakeholders (e.g., policymakers, NGOs, technical experts, civil society, as well as companies and individuals). Besides, there shall be meaningful opportunities to explain actions (e.g., using certain design concepts/training data or using AI systems in certain situations). Finally, effective and proportionate consequences (e.g., penalties for non-compliance with rules and effective redress) shall be in scope. This, in turn, creates a need for clear and feasible rules.

Notably, the weight of these principles can vary. To exemplify, the “accountability rate” might still be high if there are clear non-binding standards with no legal consequences that are widely adhered whereas it is lower if there are binding rules that are not being followed due to societal rejection or inefficient enforcement. The weight of the principles might also vary in different contexts. For instance, in policymaking, intensive discussion might have a higher priority than in company internal processes.

### 3.6.4 Open issues and research challenges

In practice, evaluating algorithmic accountability poses severe problems. One of the biggest challenges is that algorithmic accountability is a “multifaceted and context-sensitive challenge” (Cech, 2021). At present, standards, standardized methods, and metrics covering different aspects of the AI-lifecycle, from design to deployment, are still incomplete and, therefore, do not provide sufficient legal security. Vague terms confront norm addressees with legal uncertainty when interpreting these norms. In turn, organizations are unable to implement sufficient accountability mechanisms within their organization.

On the one hand, accountability gaps arise if rules are inconsistent, unclear, or not feasible, and therefore, they lead to ineffective redress of victims. On the other hand, rules which are too strict generate accountability surpluses, which in turn decrease technological and economic growth (Bovens, 2007; Novelli et al., 2023). AI policymakers aim to close these accountability gaps that might arise due to the unpredictable, opaque nature of AI systems (Novelli et al., 2023; Busuioc, 2020). Several measures, like model certification, algorithmic impact assessments, real-world testing, and third-party audits, could foster accountability (Busuioc, 2020). Such measures are also included in the AI Act. For example, there are documentation, information, and transparency obligations for providers of high-risk AI systems, there are third-party checks, and testing in real-world laboratories is enabled. Rules are also

amendable according to technical changes, demonstrating effective change if needed.

Notably, developing sufficient rules, including ethical and technical standards, that cover the whole AI-lifecycle is challenging, as AI systems are complex and based on various programming methods, developing rapidly and can have wide-ranging effects on people (Cech, 2021). Generative AI systems seem to exacerbate this problem due to their large scale, complexity, and adaptability (Xia et al., 2024). Consequently, it is particularly difficult to find suitable metrics for evaluating the trustworthiness of generative AI. To ensure “actionable” accountability, both technical and non-technical aspects, among them legal and ethical aspects, must be considered (Stix, 2021). When creating rules and standards on AI, it is crucial to weigh up technical and economic aspects. An informed dialogue between policymakers, (technical) experts, and civil society is essential to reaching sufficient rules and avoiding unnecessary bureaucracy.

## 4 Conclusion and future research

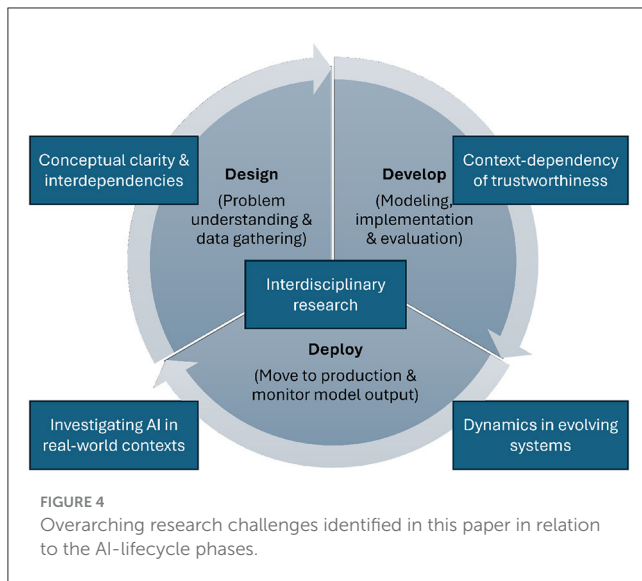
In this paper, we investigated the following six requirements of trustworthy AI: (1) human agency and oversight, (2) fairness and non-discrimination, (3) transparency and explainability, (4) robustness and accuracy, (5) privacy and security, and (6) accountability. With respect to our guiding research question introduced in Section 1 (i.e., *What is the current state of research regarding the establishment and evaluation of comprehensive—technical, human-centered, and legal—requirements of trustworthy AI?*), our findings confirm that ensuring AI systems meet these criteria is a complex endeavor requiring technical solutions, policy frameworks, and interdisciplinary collaboration. Additionally, our article demonstrates that while evaluation and validation methodologies for technical requirements, such as robustness, can often rely on well-established metrics and testing procedures (e.g., model accuracy), assessing human-centric considerations demands more nuanced approaches that take into account ethical, legal, and cultural factors. Therefore, we believe that our article complements existing surveys and assessment lists [e.g., ALTAI (Ala-Pietilä et al., 2020; Radclyffe et al., 2023)] of trustworthy AI.

This section further synthesizes our key observations across these very different aspects of AI systems in relation to their trustworthiness and discusses the implications of this overarching analysis. Additionally, Figure 4 visualizes these overarching research challenges in relation to the phases of the AI-lifecycle mentioned in Section 2.

- **Interdisciplinary research.** The interdisciplinary nature of trustworthy AI research becomes apparent when considering the different scientific foundations necessary to discuss the design, development, and deployment of trustworthy AI. This demand for interdisciplinarity is also recognized by initiatives like the human-centered AI (HCAI) workshops and sessions at AI conferences such as NeurIPS,<sup>7</sup> as well as the FAccT community’s work on fairness, accountability, and transparency in AI.<sup>8</sup> However, broader collaboration is

<sup>7</sup> <https://hcai-at-neurips.github.io/site/>

<sup>8</sup> <https://factconference.org/>



needed. In particular, insights from social scientists, ethicists, and policymakers can complement technical research, for example, in fairness definitions, explainability, or human oversight. At the same time, interdisciplinary research comes with its own challenges, such as distinct disciplinary-specific jargon. Thus, agreeing on standardized, cross-disciplinary terminology remains an ongoing challenge in various subjects.

- Conceptual clarity and interdependencies.** Across all trustworthy AI requirements, we see the need to sharpen definitions and to consider interdependencies and relationships between concepts. This involves understanding the potential trade-offs between requirements, such as fairness and accuracy or explainability and privacy. Such conceptual clarity and knowledge of interdependencies will help in designing trustworthy AI with regard to specific requirements while allowing for informed discussions of trade-offs. Therefore, it is essential to consider potential trade-offs and interdependencies already when designing trustworthy AI systems.
- Context-dependency of trustworthiness.** Our research indicates that AI requirements are very context-dependent. This means that any insights for developing trustworthy AI are challenging to transfer across different contexts due to cultural and application-specific aspects. Different interaction patterns between humans and AI will be appropriate in different contexts, and definitions of trustworthiness vary between societies and applications. This raises questions about the sufficiency of existing evaluation frameworks and suggests the need for new approaches that can better adapt to contextual differences. Additionally, if an audited algorithm is reused and fails to meet requirements in a different context, assigning responsibility becomes complex. Even more difficult is the handling of AI solutions that consist of multiple interacting components, i.e., composite systems. Understanding how different components interact and affect each other is crucial when algorithms are reused in conjunction with other components. Even if single

components are considered trustworthy, the results of their interplay potentially violate the requirements of trustworthy AI. Such uncertainties affect the licensing and use of software frameworks, which emphasizes the importance of developing licensing models that clearly outline accountability while promoting the responsible use of AI.

- Dynamics in evolving systems.** One of the emerging issues in trustworthy AI is the potential of learning unintended facets during deployment. In evolving systems (i.e., systems that learn during deployment), in particular, the capturing of biases may lead to trustworthiness issues, especially with respect to fairness and non-discrimination. Such biases are often cognitive biases of users, which are acquired through ongoing learning cycles and require more sophisticated research to form a deeper understanding of related patterns and furthermore, develop approaches for detection and mitigation. This concern also highlights the necessity of dynamic and adaptive evaluation and simulation frameworks. Since the majority of trustworthy AI evaluation schemes operate in a static manner, additional research is needed to investigate, monitor, and capture long-term dynamics of trustworthiness.
- Investigating trustworthy AI in real-world contexts.** Due to the complexity of AI systems and their contextual dependencies, it is crucial to study their functionality in real-life contexts to gain a deeper understanding of their impact. The involvement of human factors, such as how a system is used by different people and how this fits into a complex socio-technical context, makes real-world investigations very challenging from a methodological standpoint. However, for some requirements on AI systems, such as fairness or human agency, this may be particularly important to the extent that fully valid statements about these concepts may only be made after investigation in real-world contexts. Thus, monitoring the trustworthiness of AI is an ongoing investigation, especially after the system has been deployed in a real-world context.

Finally, and as outlined in Sections 3.1–3.6, the current developments around generative AI and LLMs introduce new challenges for establishing and evaluating trustworthy AI. Therefore, future research also needs to investigate how existing trustworthy AI methods and definitions (e.g., fairness metrics for binary classification problems) can be transformed into more general settings provided by generative AI and LLMs. We hope that our paper provides a reference point for both researchers and practitioners in the field of trustworthy AI and a starting point for future research directions addressing the open research challenges identified in this work and discussed in this section.

## Author contributions

DK: Conceptualization, Supervision, Writing – original draft, Writing – review & editing. SS: Investigation, Writing – original draft, Writing – review & editing. VP-S: Investigation, Writing – original draft, Writing – review & editing. PM: Investigation, Writing – original draft, Writing – review & editing. KW: Writing – original draft, Writing – review & editing. LD: Writing – original



draft, Writing – review & editing. AF: Writing – original draft, Writing – review & editing. MT: Writing – original draft, Writing – review & editing. IM: Writing – original draft, Writing – review & editing. IŠ: Writing – original draft, Writing – review & editing. VS: Writing – original draft, Writing – review & editing. AT: Writing – original draft, Writing – review & editing. EV: Writing – original draft, Writing – review & editing. RK: Writing – original draft, Writing – review & editing. TN: Investigation, Writing – original draft, Writing – review & editing. SK: Conceptualization, Supervision, Writing – original draft, Writing – review & editing.

## Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article. The work received funding from the TU Graz Open Access Publishing Fund.

## Acknowledgments

Parts of this manuscript are based on white papers created by SGS Digital Trust Services GmbH and Know Center Research GmbH, which are available online via Zenodo: <https://zenodo.org/>

## References

- Abadi, M., Chu, A., Goodfellow, I. J., McMahan, H. B., Mironov, I., Talwar, K., et al. (2016). “Deep learning with differential privacy,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, eds. E. R. Weippl, S. Katzenbeisser, C. Kruegel, A. C. Myers, and S. Halevi (Vienna: ACM), 308–318.
- Adadi, A., and Berrada, M. (2018). Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE Access* 6, 52138–52160. doi: 10.1109/ACCESS.2018.2870052
- Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., and Kim, B. (2018). “Sanity checks for saliency maps,” in *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18* (Red Hook, NY: Curran Associates Inc.), 9525–9536.
- Adilova, L., Böttinger, K., Danos, V., Jakob, S., Langer, F., Markert, T., et al. (2022). *Security of AI-Systems: Fundamentals - Adversarial Deep Learning*. Available at: [https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/KI/Security-of-AI-systems\\_fundamentals.html](https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/KI/Security-of-AI-systems_fundamentals.html) (accessed June, 2024).
- Akbar, M. A., Khan, A. A., Mahmood, S., Rafi, S., and Demi, S. (2024). Trustworthy artificial intelligence: a decision-making taxonomy of potential challenges. *Softw. Pract. Exp.* 54, 1621–1650. doi: 10.1002/spe.3216
- Akula, R., and Garibay, I. (2021). Audit and assurance of ai algorithms: a framework to ensure ethical algorithmic practices in artificial intelligence. *arXiv [preprint]*. doi: 10.48550/arXiv.2107.14046
- Ala-Pietilä, P., Bonnet, Y., Bergmann, U., Bielikova, M., Bonefeld-Dahl, C., Bauer, W., et al. (2020). *The Assessment List for Trustworthy Artificial Intelligence (ALTAI)*. Lausanne: European Commission.
- Albahri, A. S., Duhaim, A. M., Fadhel, M. A., Alnoor, A., Baqer, N. S., Alzubaidi, L., et al. (2023). A systematic review of trustworthy and explainable artificial intelligence in healthcare: Assessment of quality, bias risk, and data fusion. *Inf. Fus.* 96, 156–191. doi: 10.1016/j.inffus.2023.03.008
- Almanifi, O. R. A., Chow, C.-O., Tham, M.-L., Chuah, J. H., and Kanesan, J. (2023). Communication and computation efficiency in federated learning: a survey. *Int. Things* 22:100742. doi: 10.1016/j.iot.2023.100742
- Alvarez-Melis, D., and Jaakkola, T. S. (2018a). On the robustness of interpretability methods. *arXiv [preprint]*. doi: 10.48550/arXiv.1806.08049
- Alvarez-Melis, D., and Jaakkola, T. S. (2018b). “Towards robust interpretability with self-explaining neural networks,” in *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NeurIPS)*, 7786–7795.
- Anders, C. J., Weber, L., Neumann, D., Samek, W., Müller, K.-R., and Lapuschkin, S. (2022). Finding and removing clever hans: using explanation methods to debug and improve deep models. *Inf. Fus.* 77, 261–295. doi: 10.1016/j.inffus.2021.07.015
- Anderson, A., Maystre, L., Anderson, I., Mehrotra, R., and Lalmas, M. (2020). “Algorithmic effects on the diversity of consumption on spotify,” in *Proceedings of the Web Conference 2020* (New York, NY: ACM), 2155–2165.
- Anderson, M., and Fort, K. (2022). Human where? A new scale defining human involvement in technology communities from an ethical standpoint. *Int. Rev. Inf. Ethics.* 31. doi: 10.29173/irrie477
- Arias-Duart, A., Parés, F., Garcia-Gasulla, D., and Gimenez-Abalos, V. (2022). “Focus! rating xai methods and finding biases,” in *2022 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)* (New York, NY: IEEE), 1–8.
- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., et al. (2020). Explainable artificial intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible ai. *Inf. Fus.* 58, 82–115. doi: 10.1016/j.inffus.2019.12.012
- Artelt, A., and Hammer, B. (2019). On the computation of counterfactual explanations - a survey. *arXiv [preprint]*. doi: 10.48550/arXiv.1911.07749
- Arya, V., Bellamy, R. K., Chen, P.-Y., Dhurandhar, A., Hind, M., Hoffman, S. C., et al. (2019). One explanation does not fit all: a toolkit and taxonomy of ai explainability techniques. *arXiv [preprint]*. doi: 10.1145/3351095.3375667
- Bae, H., Jang, J., Jung, D., Jang, H., Ha, H., and Yoon, S. (2018). Security and privacy issues in deep learning. *arXiv [preprint]*. doi: 10.48550/arXiv.1807.11655
- Baeza-Yates, R. (2018). Bias on the web. *Commun. ACM* 61, 54–61. doi: 10.1145/3209581
- Bagdasaryan, E., Poursaeed, O., and Shmatikov, V. (2019). Differential privacy has disparate impact on model accuracy. *Adv. Neural Inf. Process. Syst.* 32.
- Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., and Shmatikov, V. (2020). “How to backdoor federated learning,” in *International Conference on Artificial Intelligence and Statistics (PMLR)*, 2938–2948.
- Barocas, S., Hardt, M., and Narayanan, A. (2021). *Fairness and Machine Learning: Limitations and Opportunities*. Available at: <https://fairmlbook.org/> (accessed June 19, 2019).
- Barocas, S., Hardt, M., and Narayanan, A. (2023). *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press.
- Barocas, S., and Selbst, A. D. (2016). Big data’s disparate impact. *Calif. L. Rev.* 104:671. doi: 10.2139/ssrn.2477899
- Baxter, G., and Sommerville, I. (2010). Socio-technical systems: from design methods to systems engineering. *Interact. Comput.* 23, 4–17. doi: 10.1016/j.intcom.2010.07.003
- records/11207961. Know Center Research GmbH is a COMET Center within the COMET—Competence Centers for Excellent Technologies Program and funded by BMK, BMAW, as well as the co-financing provinces Styria, Vienna and Tyrol. COMET is managed by FFG.

## Conflict of interest

DK, SS, VP-S, PM, KW, LD, AF, MT, IM, IŠ, VS, AT, EV, RK, and SK were employed by Know Center Research GmbH. TN was employed by SGS Digital Services GmbH.

## Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Bellamy, R. K., Dey, K., Hind, M., Hoffman, S. C., Houde, S., Kannan, K., et al. (2019). AI fairness 360: an extensible toolkit for detecting and mitigating algorithmic bias. *IBM J. Res. Dev.* 63: 2942287. doi: 10.1147/JRD.2019.2942287
- Bennett, D., Metatla, O., Roudaut, A., and Mekler, E. D. (2023). "How does HCI understand human agency and autonomy?," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (New York, NY: ACM), 1–18.
- Bhatt, U., Weller, A., and Moura, J. M. F. (2020). "Evaluating and aggregating feature-based model explanations," in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20, 3016-3022. International Joint Conferences on Artificial Intelligence Organization*, ed. C. Bessiere. (Washington, DC: AAAI Press).
- Biecek, P. (2018). Dalex: explainers for complex predictive models in R. *J. Mach. Learn. Res.* 19, 1–5.
- Binns, R. (2018). Algorithmic accountability and public reason. *Philos. Technol.* 31, 1–14. doi: 10.1007/s13347-017-0263-5
- Bird, S., Dudik, M., Edgar, R., Horn, B., Lutz, R., Milan, V., et al. (2020). *Fairlearn: A Toolkit for Assessing and Improving Fairness in AI*. Technical Report MSR-TR-2020-32, Microsoft. Washington, DC: ACM.
- Bird, S., Kenthapadi, K., Kiciman, E., and Mitchell, M. (2019). "Fairness-aware machine learning: practical challenges and lessons learned," in *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining* (New York, NY: ACM), 834–835.
- Bivins, T. H. (2006). "Responsibility and accountability," in *Ethics in Public Relations: Responsible Advocacy*, eds. K. Fitzpatrick and C. Bronstein (SAGE Publications, Inc.), 19–38.
- Blanchard, G., Lee, G., and Scott, C. (2011). Generalizing from several related classification tasks to a new unlabeled sample. *Adv. Neural Inf. Process. Syst.* 24.
- Bovens, M. (2007). Analysing and assessing accountability: a conceptual framework. *Eur. L. J.* 13, 447–468. doi: 10.1111/j.1468-0386.2007.00378.x
- Bovens, M. (2010). Two concepts of accountability: accountability as a virtue and as a mechanism. *West Eur. Polit.* 33, 946–967. doi: 10.1080/01402382.2010.486119
- Brandsma, G., and Schillemans, T. (2012). The accountability cube: measuring accountability. *J. Public Administ. Res. Theory* 23, 953–975. doi: 10.1093/jopart/mus034
- Brown, A., Chouldechova, A., Putnam-Hornstein, E., Tobin, A., and Vaithianathan, R. (2019). "Toward algorithmic accountability in public services: a qualitative study of affected community perspectives on algorithmic decision-making in child welfare services," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (New York, NY: ACM), 1–12.
- Brown, H., Lee, K., Miresheghallah, F., Shokri, R., and Tramèr, F. (2022). "What does it mean for a language model to preserve privacy?," in *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (ACM), 2280–2292.
- BSI (2022). *AI Security Concerns in a Nutshell*. Available at: [https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/KI/Practical\\_AI-Security\\_Guide\\_2023.html](https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/KI/Practical_AI-Security_Guide_2023.html) (accessed June, 2024).
- Buhmann, P., Paßmann, A., and Christian, F. (2019). Managing algorithmic accountability: Balancing reputational concerns, engagement strategies, and the potential of rational discourse. *J. Bus. Ethics* 163, 265–280. doi: 10.1007/s10551-019-04226-4
- Buolamwini, J., and Gebru, T. (2018). "Gender shades: Intersectional accuracy disparities in commercial gender classification," in *Proceedings of the 1st Conference on Fairness, Accountability and Transparency, volume 81 of Proceedings of Machine Learning Research*, eds. S. A. Friedler, and C. Wilson (PMLR), 77–91.
- Busuioc, M. (2020). Accountable artificial intelligence: Holding algorithms to account. *Public Adm. Rev.* 81:13293. doi: 10.1111/puar.13293
- Calders, T., Kamiran, F., and Pechenizkiy, M. (2009). "Building classifiers with independency constraints," in *2009 IEEE International Conference on Data Mining Workshops* (New York, NY: IEEE), 13–18.
- Calmon, F., Wei, D., Vinzamuri, B., Natesan Ramamurthy, K., and Varshney, K. R. (2017). "Optimized pre-processing for discrimination prevention," in *NIPS'17: Proceedings of the 31st International Conference on Neural Information Processing Systems* (MIT Press).
- Cao, L. (2022). AI in finance: challenges, techniques, and opportunities. *ACM Comp. Surv.* 55, 1–38. doi: 10.1145/3502289
- Carter, S., Armstrong, Z., Schubert, L., Johnson, I., and Olah, C. (2019). Activation atlas. *Distill* 4:e15. doi: 10.23915/distill.00015
- Carvalho, D. V., Pereira, E. M., and Cardoso, J. S. (2019). Machine learning interpretability: a survey on methods and metrics. *Electronics* 8:832. doi: 10.3390/electronics8080832
- Cech, F. (2021). The agency of the forum: mechanisms for algorithmic accountability through the lens of agency. *J. Respons. Technol.* 7–8:100015. doi: 10.1016/j.jrt.2021.100015
- Chang, H.-C. H., Bui, M., and McIlwain, C. (2021). Targeted ads and/as racial discrimination: Exploring trends in new york city ads for college scholarships. *arXiv [preprint]*. doi: 10.24251/HICSS.2022.348
- Chatila, R., Dignum, V., Fisher, M., Giannotti, F., Morik, K., Russell, S., et al. (2021). *Trustworthy AI. Reflections on Artificial Intelligence for Humanity* (New York, NY: Springer), 13–39.
- Chen, Z. (2023). Ethics and discrimination in artificial intelligence-enabled recruitment practices. *Human. Soc. Sci. Commun.* 10, 1–12. doi: 10.1057/s41599-023-02079-x
- Costanza-Chock, S., Raji, I. D., and Buolamwini, J. (2022). "Who audits the auditors? recommendations from a field scan of the algorithmic auditing ecosystem," in *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (New York, NY: ACM), 1571–1583.
- Debenedetti, E., Severi, G., Carlini, N., Choquette-Choo, C. A., Jagielski, M., Nasr, M., et al. (2023). Privacy side channels in machine learning systems. *arXiv [preprint]*. doi: 10.48550/arXiv.2309.05610
- Dennerlein, S., Wolf-Brenner, C., Gutounig, R., Schweiger, S., and Pammer-Schindler, V. (2020). "Guiding socio-technical reflection of ethical principles in tel software development: the srep framework," in *Addressing Global Challenges and Quality Education*, eds. C. Alario-Hoyos, M. J. Rodríguez-Triana, M. Scheffel, I., Arnedillo-Sánchez, and S. M. Dennerlein (Cham: Springer International Publishing), 386–391.
- Díaz-Rodríguez, N., Del Ser, J., Coeckelbergh, M., de Prado, M. L., Herrera-Viedma, E., and Herrera, F. (2023). Connecting the dots in trustworthy artificial intelligence: from AI principles, ethics, and key requirements to responsible ai systems and regulation. *Inf. Fus.* 99:101896. doi: 10.1016/j.inffus.2023.101896
- Doshi-Velez, F., and Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv [preprint]*. doi: 10.48550/arXiv.1702.08608
- Dubal, V. (2023). On algorithmic wage discrimination. *Columbia Law Rev.* 123, 1929–1992. doi: 10.2139/ssrn.4331080
- Durante, M., and Floridi, L. (2022). "A legal principles-based framework for ai liability regulation," in *The 2021 Yearbook of the Digital Ethics Lab*, eds. J. Mókander and M. Ziosi (New York, NY: Springer), 93–112.
- Dutta, S., Wei, D., Yueksel, H., Chen, P.-Y., Liu, S., and Varshney, K. (2020). "Is there a trade-off between fairness and accuracy? a perspective using mismatched hypothesis testing," in *International Conference on Machine Learning* (PMLR), 2803–2813.
- Dwork, C. (2008). "Differential privacy: a survey of results," in *Theory and Applications of Models of Computation, 5th International Conference, TAMC 2008, Xi'an, China, April 25-29, 2008. Proceedings, volume 4978 of Lecture Notes in Computer Science*, eds. M. Agrawal, D. Du, Z. Duan, and A. Li (Springer), 1–19.
- Dwork, C., and Roth, A. (2014). The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* 9, 211–407. doi: 10.1561/04000000042
- Elliott, D., and Soifer, E. (2022). AI technologies, privacy, and security. *Front. Artif. Intell.* 5:826737. doi: 10.3389/frai.2022.826737
- Emaminejad, N., and Akhavan, R. (2022). Trustworthy AI and robotics: implications for the aec industry. *Automat. Construct.* 139:104298. doi: 10.1016/j.autcon.2022.104298
- Eraut, M. (2004). Informal learning in the workplace. *Stud. Contin. Educ.* 26, 247–273. doi: 10.1080/15803704200025245
- Eriksén, S. (2002). "Designing for accountability," in *Proceedings of the Second Nordic Conference on Human-Computer Interaction* (New York, NY: ACM), 177–186.
- Evans, D., Kolesnikov, V., and Rosulek, M. (2018). A pragmatic introduction to secure multi-party computation. *Found. Trends Priv. Secur.* 2, 70–246. doi: 10.1561/9781680835090
- Fancher, D., Ammanath, B., Holdowsky, J., and Natasha, B. (2021). *Deloitte. Insights AI Model Bias Can Damage Trust More Than You May Know. But it Doesn't Have To* (London: Deloitte Development LLC).
- Fanni, R., Steinkogler, V. E., Zampedri, G., and Pierson, J. (2020). "Active human agency in artificial intelligence mediation," in *Proceedings of the 6th EAI International Conference on Smart Objects and Technologies for Social Good* (Gent: EAI), 84–89.
- Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., and Venkatasubramanian, S. (2015). "Certifying and removing disparate impact," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (New York, NY: ACM), 259–268.
- Fisher, R. A. (1922). On the mathematical foundations of theoretical statistics. *Philos. Transact. R. Soc. London Ser. A* 222, 309–368. doi: 10.1098/rsta.1922.0009
- Floridi, L. (2021). Establishing the rules for building trustworthy AI. *Ethics Govern. Policies Artif. Intell.* 41–45. doi: 10.1007/978-3-030-81907-1\_4
- Fong, R. C., and Vedaldi, A. (2017). "Interpretable explanations of black boxes by meaningful perturbation," in *Proceedings of the IEEE International Conference on Computer Vision* (New York, NY: IEEE), 3429–3437.
- Friedler, S. A., Scheidegger, C., Venkatasubramanian, S., Choudhary, S., Hamilton, E. P., and Roth, D. (2019). "A comparative study of fairness-enhancing interventions in machine learning," in *Proceedings of the Conference on Fairness, Accountability, and Transparency* (New York, NY: ACM), 329–338.
- Friedman, A., Berkovsky, S., and Kaafar, M. A. (2016). A differential privacy framework for matrix factorization recommender systems. *User Model. User Adapt. Interact.* 26, 425–458. doi: 10.1007/s11257-016-9177-7

- Gabriel, A., León, E. K. C., and Wilkins, A. (2021). Accountability increases resource sharing: effects of accountability on human and AI system performance. *Int. J. Human-Comp. Interact.* 37, 434–444. doi: 10.1080/10447318.2020.1824695
- Gao, D., Liu, Y., Huang, A., Ju, C., Yu, H., and Yang, Q. (2019). “Privacy-preserving heterogeneous federated transfer learning,” in *2019 IEEE International Conference on Big Data (Big Data)* (New York, NY: IEEE), 2552–2559.
- García, L. P. F., de Carvalho, A. C. P. L. F., and Lorena, A. C. (2015). Effect of label noise in the complexity of classification problems. *Neurocomputing* 160, 108–119. doi: 10.1016/j.neucom.2014.10.085
- Gentry, C. (2009). *A Fully Homomorphic Encryption Scheme*. Stanford, CA: Stanford University.
- Guidotti, R. (2022). Counterfactual explanations and how to find them: literature review and benchmarking. *Data Mining Knowl. Disc.* 38, 1–55. doi: 10.1007/s10618-022-00831-6
- Gulrajani, I., and Lopez-Paz, D. (2020). In search of lost domain generalization. *arXiv [preprint]*. doi: 10.48550/arXiv.2007.01434
- Haakman, M., Cruz, L., Huijgens, H., and van Deursen, A. (2021). AI lifecycle models need to be revised: an exploratory study in fintech. *Empir. Softw. Eng.* 26, 1–29. doi: 10.1007/s10664-021-09993-1
- Hamon, R., Junklewitz, H., and Sanchez, I. (2020). *Robustness and Explainability of Artificial Intelligence*. Luxembourg: Publications Office of the European Union 207:2020.
- Han, S., Lin, C., Shen, C., Wang, Q., and Guan, X. (2023). Interpreting adversarial examples in deep learning: a review. *ACM Comp. Surv.* 55, 1–38. doi: 10.1145/3594869
- Hauer, M., Krafft, T., and Zweig, K. (2023). Overview of transparency and inspectability mechanisms to achieve accountability of artificial intelligence systems. *Data Policy* 5:30. doi: 10.1017/dap.2023.30
- Hedström, A., Weber, L., Krakowczyk, D., Bareeva, D., Motzkus, F., Samek, W., et al. (2023). Quantus: an explainable ai toolkit for responsible evaluation of neural network explanations and beyond. *J. Mach. Learn. Res.* 24, 1–11.
- Hendrycks, D., and Dietterich, T. (2019). Benchmarking neural network robustness to common corruptions and perturbations. *arXiv [preprint]*. doi: 10.48550/arXiv.1903.12261
- Hermann, E. (2022). Artificial intelligence and mass personalization of communication content—an ethical and literacy perspective. *New Media Soc.* 24, 1258–1277. doi: 10.1177/14614448211022702
- High-Level Expert Group on AI (2019). *Ethics Guidelines for Trustworthy AI*. Brussels: Report, European Commission.
- Holzinger, A., Saranti, A., Molnar, C., Biecek, P., and Samek, W. (2022). “Explainable ai methods—a brief overview,” in *International Workshop on Extending Explainable AI Beyond Deep Models and Classifiers* (New York, NY: Springer), 13–38.
- Houwer, J. D. (2019). Implicit bias is behavior: a functional-cognitive perspective on implicit bias. *Perspect. Psychol. Sci.* 14, 835–840. doi: 10.1177/1745691619855638
- Huber, P. J. (2004). *Robust Statistics*. Vol. 523. John Wiley & Sons.
- Hulsen, T. (2023). Explainable artificial intelligence (XAI): concepts and challenges in healthcare. *AI* 4, 652–666. doi: 10.3390/ai4030034
- John-Mathews, J.-M. (2022). Some critical and ethical perspectives on the empirical turn of AI interpretability. *Technol. Forecast. Soc. Change* 174:121209. doi: 10.1016/j.techfore.2021.121209
- Kahn, K., and Winters, N. (2017). “Child-friendly programming interfaces to AI cloud services,” in *12th European Conference on Technology Enhanced Learning* (New York, NY: Springer), 566–570.
- Kaur, D., Uslu, S., and Durresi, A. (2021). “Requirements for trustworthy artificial intelligence—a review,” in *Advances in Networked-Based Information Systems: The 23rd International Conference on Network-Based Information Systems (NBIS-2020) 23* (New York, NY: Springer), 105–115.
- Kaur, D., Uslu, S., Rittichier, K. J., and Durresi, A. (2022). Trustworthy artificial intelligence: a review. *ACM Comp. Surv.* 55, 1–38. doi: 10.1145/3491209
- Kim, J., and Scott, C. D. (2012). Robust kernel density estimation. *J. Mach. Learn. Res.* 13, 2529–2565.
- Kindermans, P.-J., Hooker, S., Adebayo, J., Alber, M., Schütt, K. T., Dähne, S., et al. (2019). “The (Un)reliability of saliency methods,” in *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, eds. W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, and K. R. Müller (Cham: Springer International Publishing), 267–280.
- Kohli, P., and Chadha, A. (2020). “Enabling pedestrian safety using computer vision techniques: a case study of the 2018 Uber Inc. self-driving car crash,” in *Advances in Information and Communication: Proceedings of the 2019 Future of Information and Communication Conference (FICC), Volume 1* (New York, NY: Springer), 261–279.
- Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Alsallakh, B., Reynolds, J., et al. (2020). Captum: a unified and generic model interpretability library for pytorch. *arXiv [preprint]*. doi: 10.48550/arXiv.2009.07896
- Koshiyama, A., Kazim, E., Treleaven, P., Rai, P., Szpruch, L., Pavey, G., et al. (2021). Towards algorithm auditing: a survey on managing legal, ethical and technological risks of ai, ml and associated algorithms. *SSRN Electron. J.* doi: 10.2139/ssrn.3778998
- Kowald, D., and Lacic, E. (2022). “Popularity bias in collaborative filtering-based multimedia recommender systems,” in *International Workshop on Algorithmic Bias in Search and Recommendation* (New York, NY: Springer), 1–11.
- Kowald, D., Schedl, M., and Lex, E. (2020). “The unfairness of popularity bias in music recommendation: a reproducibility study,” in *Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part II 42* (New York, NY: Springer), 35–42.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* 25. doi: 10.1145/3065386
- Lazer, D., Kennedy, R., King, G., and Vespignani, A. (2014). The parable of google flu: traps in big data analysis. *Science* 343, 1203–1205. doi: 10.1126/science.1248506
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521, 436–444. doi: 10.1038/nature14539
- Lepri, B., Oliver, N., Letouzé, E., Pentland, A., and Vinck, P. (2018). Fair, transparent, and accountable algorithmic decision-making processes: The premise, the proposed solutions, and the open challenges. *Philos. Technol.* 31, 611–627. doi: 10.1007/s13347-017-0279-x
- Lewis, D., Hogan, L., Filip, D., and Wall, P. J. (2020). Global challenges in the standardization of ethics for trustworthy AI. *J. ICT Standard.* 8, 123–150. doi: 10.13052/jicsts2245-800X.823
- Li, B., Qi, P., Liu, B., Di, S., Liu, J., Pei, J., et al. (2023). Trustworthy AI: from principles to practices. *ACM Comp. Surv.* 55, 1–46. doi: 10.1145/3555803
- Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. (2020). Federated learning: challenges, methods, and future directions. *IEEE Signal Process. Mag.* 37, 50–60. doi: 10.1109/MSP.2020.2975749
- Liang, W., Tadesse, G. A., Ho, D., Fei-Fei, L., Zaharia, M., Zhang, C., et al. (2022). Advances, challenges and opportunities in creating data for trustworthy AI. *Nat. Mach. Intell.* 4, 669–677. doi: 10.1038/s42256-022-00516-1
- Liao, Q. V., Gruen, D., and Miller, S. (2020). “Questioning the ai: informing design practices for explainable ai user experiences,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (New York, NY: ACM), 1–15.
- Liu, F., Cheng, Z., Chen, H., Wei, Y., Nie, L., and Kankanhalli, M. (2022). “Privacy-preserving synthetic data generation for recommendation systems,” in *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (New York, NY: ACM), 1379–1389.
- Long, D., and Magerko, B. (2020). “What is ai literacy? competencies and design considerations,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (New York, NY: ACM), 1–16.
- Lundberg, S. M., and Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* 30, 4765–4774.
- Madiega, T. (2021). *Artificial Intelligence Act*. Brussels: European Parliament: European Parliamentary Research Service.
- McDermid, J. A., Jia, Y., Porter, Z., and Habli, I. (2021). Artificial intelligence explainability: the technical and ethical dimensions. *Philos. Transact. R. Soc. A* 379:20200363. doi: 10.1098/rsta.2020.0363
- McGregor, L., Murray, D., and Ng, V. (2019). International human rights law as a framework for algorithmic accountability. *Int. Comp. Law Q.* 68, 309–343. doi: 10.1017/S0020589319000046
- Mehdiyev, N., and Fettek, P. (2021). Explainable artificial intelligence for process mining: a general overview and application of a novel local explanation approach for predictive process monitoring. *Interpret. Artif. Intell.* 937, 1–28. doi: 10.1007/978-3-030-64949-4\_1
- Mendoza, I. G., Sabol, V., and Hoffer, J. G. (2023). “On the importance of user role-tailored explanations in industry 5.0,” in *VISIGRAPP (2: HUCAPP)* (Setúbal: SciTePress), 243–250.
- Miller, T., Howe, P., and Sonenberg, L. (2017). Explainable AI: beware of inmates running the asylum or: How i learnt to stop worrying and love the social and behavioural sciences. *arXiv [preprint]*. doi: 10.48550/arXiv.1712.00547
- Molnar, C. (2020). *Interpretable Machine Learning*. Lulu Press.
- Montavon, G., Samek, W., and Müller, K.-R. (2018). Methods for interpreting and understanding deep neural networks. *Digit. Signal Process.* 73, 1–15. doi: 10.1016/j.dsp.2017.10.011
- Moore, C., O’Neill, M., O’Sullivan, E., Doröz, Y., and Sunar, B. (2014). “Practical homomorphic encryption: a survey,” in *2014 IEEE International Symposium on Circuits and Systems (ISCAS)* (New York, NY: IEEE), 2792–2795.
- Moreira, C., Chou, Y.-L., Hsieh, C., Ouyang, C., Jorge, J., and Pereira, J. M. (2022). Benchmarking counterfactual algorithms for XAI: from white box to black box. *arXiv [preprint]*. doi: 10.48550/arXiv.2203.02399



- Mosqueira-Rey, E., Hernández-Pereira, E., Alonso-Ríos, D., Bobes-Bascarán, J., and Fernández-Leal, Á. (2023). Human-in-the-loop machine learning: a state of the art. *Artif. Intell. Rev.* 56, 3005–3054. doi: 10.1007/s10462-022-10246-w
- Muandet, K., Balduzzi, D., and Schölkopf, B. (2013). “Domain generalization via invariant feature representation,” in *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, ICML'13 (JMLR.org)*, 10–18.
- Mueller, P., Kowald, D., and Lex, E. (2021). “Robustness of meta matrix factorization against strict privacy constraints,” in *Advances in Information Retrieval: 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28–April 1, 2021, Proceedings, Part II 43* (New York, NY: Springer), 107–119.
- Müllner, P., Lex, E., Schedl, M., and Kowald, D. (2023). Differential privacy in collaborative filtering recommender systems: a review. *Fron. Big Data* 6:1249997. doi: 10.3389/fdata.2023.1249997
- Müllner, P., Lex, E., Schedl, M., and Kowald, D. (2024). “The impact of differential privacy on recommendation accuracy and popularity bias,” in *European Conference on Information Retrieval* (New York, NY: Springer), 466–482.
- Munro, R. (2021). *Human-in-the-Loop Machine Learning: Active Learning and Annotation for Human-Centered AI*. Manning.
- Naidu, G., Zuva, T., and Sibanda, E. M. (2023). “A review of evaluation metrics in machine learning algorithms,” in *Computer Science On-line Conference* (New York, NY: Springer), 15–25.
- Naiseh, M., Al-Thani, D., Jiang, N., and Ali, R. (2023). How the different explanation classes impact trust calibration: the case of clinical decision support systems. *Int. J. Hum. Comput. Stud.* 169:102941. doi: 10.1016/j.ijhcs.2022.102941
- Nasr, M., Shokri, R., and Houmansadr, A. (2019). “Comprehensive privacy analysis of deep learning: passive and active white-box inference attacks against centralized and federated learning,” in *2019 IEEE Symposium on Security and Privacy (SP)* (New York, NY: IEEE), 739–753.
- Nemani, P., Joel, Y. D., Vijay, P., and Liza, F. F. (2023). Gender bias in transformers: a comprehensive review of detection and mitigation strategies. *Nat. Lang. Process. J.* 6:100047. doi: 10.1016/j.nlp.2023.100047
- Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., and Qiao, M. S. (2021). Conceptualizing ai literacy: an exploratory review. *Comp. Educ. Artif. Intell.* 2:100041. doi: 10.1016/j.caeai.2021.100041
- Nguyen, A.-P., and Martínez, M. R. (2020). On quantitative aspects of model interpretability. *arXiv [preprint]*. doi: 10.48550/arXiv.2007.07584
- Nicolae, M.-I., Sinn, M., Tran, M. N., Buesser, B., Rawat, A., Wistuba, M., et al. (2019). Adversarial robustness Toolbox v1.0.0. *arXiv [preprint]*. doi: 10.48550/arXiv.1807.01069
- Novelli, C., Taddeo, M., and Floridi, L. (2023). Accountability in artificial intelligence: what it is and how it works. *AI Soc.* 39, 1–12. doi: 10.2139/ssrn.4180366
- Ntoutsis, E., Fafalios, P., Gadiraju, U., Iosifidis, V., Nejd, W., Vidal, M.-E., et al. (2020). Bias in data-driven artificial intelligence systems—an introductory survey. *Wiley Interdiscip. Rev.* 10:e1356. doi: 10.1002/widm.1356
- Nushi, B., Kamar, E., and Horvitz, E. (2018). “Towards accountable AI: hybrid human-machine analyses for characterizing system failure,” in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, Vol. 6* (Washington, DC: AAAI Press), 126–135.
- Pammer-Schindler, V., and Lindstaedt, S. (2022). AI literacy f<sup>ur</sup> entscheidungsträgerinnen im strategischen management. *Wirtschaftsinformatik Manag.* 14, 140–143. doi: 10.1365/s35764-022-00399-2
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830.
- Pendleton, M., Garcia-Lebron, R., Cho, J.-H., and Xu, S. (2016). A survey on systems security metrics. *ACM Comp. Surv.* 49, 1–35. doi: 10.1145/3005714
- Pessach, D., and Shmueli, E. (2023). “Algorithmic fairness,” in *Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook* (New York, NY: Springer), 867–886.
- Phong, L. T., Aono, Y., Hayashi, T., Wang, L., and Moriai, S. (2018). Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Trans. Inf. Forensics Secur.* 13, 1333–1345. doi: 10.1109/TIFS.2017.2787987
- Pleiss, G., Raghavan, M., Wu, F., Kleinberg, J., and Weinberger, K. Q. (2017). On fairness and calibration. *Adv. Neural Inf. Process. Syst.* 30.
- Poretschkin, M., Schmitz, A., Akila, M., Adilova, L., Becker, D., Cremers, A. B., et al. (2023). Guideline for trustworthy artificial intelligence-ai assessment catalog. *arXiv [preprint]*. doi: 10.48550/arXiv.2307.03681
- Radclyffe, C., Ribeiro, M., and Wortham, R. H. (2023). The assessment list for trustworthy artificial intelligence: a review and recommendations. *Front. Artif. Intell.* 6:1020592. doi: 10.3389/frai.2023.1020592
- Rajpurkar, P., Chen, E., Banerjee, O., and Topol, E. J. (2022). AI in health and medicine. *Nat. Med.* 28, 31–38. doi: 10.1038/s41591-021-01614-0
- Rechberger, C., and Walch, R. (2022). “Privacy-preserving machine learning using cryptography,” in *Security and Artificial Intelligence* (New York, NY: Springer), 109–129.
- Ren, H., Deng, J., and Xie, X. (2022). Grmn: generative regression neural network—a data leakage attack for federated learning. *ACM Transact. Intell. Systems. Technol.* 13, 1–24. doi: 10.1145/3510032
- Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). “Why should i trust you?” Explaining the predictions of any classifier,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (New York, NY: ACM), 1135–1144.
- Righetti, L., Madhavan, R., and Chatila, R. (2019). Unintended consequences of biased robotic and artificial intelligence systems [ethical, legal, and societal issues]. *IEEE Robot. Automat. Mag.* 26, 11–13. doi: 10.1109/MRA.2019.2926996
- Roy, D., Murty, K. R., and Mohan, C. K. (2015). “Feature selection using deep neural networks,” in *2015 International Joint Conference on Neural Networks (IJCNN)* (New York, NY: IEEE), 1–6.
- Rusak, E., Schott, L., Zimmermann, R. S., Bitterwolf, J., Bringmann, O., Bethge, M., et al. (2020). “A simple way to make neural networks robust against diverse image corruptions,” in *European Conference on Computer Vision* (New York, NY: Springer), 53–69.
- Saleiro, P., Kuester, B., Hinkson, L., London, J., Stevens, A., Anisfeld, A., et al. (2018). Aequitas: a bias and fairness audit toolkit. *arXiv [preprint]*. doi: 10.48550/arXiv.1811.05577
- Samek, W., Binder, A., Montavon, G., Lapuschkin, S., and Müller, K.-R. (2016). Evaluating the visualization of what a deep neural network has learned. *IEEE Transact. Neural Netw. Learn. Syst.* 28, 2660–2673. doi: 10.1109/TNNLS.2016.2599820
- Samek, W., Montavon, G., Lapuschkin, S., Anders, C. J., and Müller, K.-R. (2021). Explaining deep neural networks and beyond: a review of methods and applications. *Proc. IEEE* 109, 247–278. doi: 10.1109/JPROC.2021.3060483
- Saxena, N. A. (2019). “Perceptions of fairness,” in *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (New York, NY: AAAI; ACM), 537–538.
- Scher, S., and Trügler, A. (2023). Testing robustness of predictions of trained classifiers against naturally occurring perturbations. *arXiv [preprint]*. doi: 10.48550/arXiv.2204.10046
- Schneider, J. (2024). Explainable generative AI (genXAI): a survey, conceptualization, and research agenda. *Artif. Intell. Rev.* 57:289. doi: 10.1007/s10462-024-10916-x
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D. (2017). “Grad-CAM: visual explanations from deep networks via gradient-based localization,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (New York, NY: IEEE), 618–626.
- Seyyed-Kalantari, L., Zhang, H., McDermott, M. B., Chen, I. Y., and Ghassemi, M. (2021). Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nat. Med.* 27, 2176–2182. doi: 10.1038/s41591-021-01595-0
- Sharma, S., Henderson, J., and Ghosh, J. (2020). “CERTIFAI: a common framework to provide explanations and analyse the fairness and robustness of black-box models,” in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* (New York, NY: ACM), 166–172.
- Shrikumar, A., Greenside, P., and Kundaje, A. (2017). “Learning important features through propagating activation differences,” in *International Conference on Machine Learning* (New York, NY: PMLR), 3145–3153.
- Šimić, I., Sabol, V., and Veas, E. (2022). “Perturbation effect: a metric to counter misleading validation of feature attribution,” in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (New York, NY: ACM), 1798–1807.
- Simkute, A., Tankelevitch, L., Kewenig, V., Scott, A. E., Sellen, A., and Rintel, S. (2024). Ironies of generative AI: understanding and mitigating productivity loss in human-ai interactions. *arXiv [preprint]*. doi: 10.1080/10447318.2024.2405782
- Slokum, M. (2018). “Comparing recommender systems using synthetic data,” in *Proceedings of the 12th ACM Conference on Recommender Systems* (New York, NY: ACM), 548–552.
- Smart, N. P. (2016). *Cryptography Made Simple. Information Security and Cryptography*. New York, NY: Springer.
- Smuha, N. A. (2019). The eu approach to ethics guidelines for trustworthy artificial intelligence. *Comp. Law Rev. Int.* 20, 97–106. doi: 10.9785/crl-2019-200402
- Snyder, H. (2019). Literature review as a research methodology: an overview and guidelines. *J. Bus. Res.* 104, 333–339. doi: 10.1016/j.jbusres.2019.07.039
- Srivastava, M., Heidari, H., and Krause, A. (2019). “Mathematical notions vs. human perception of fairness: a descriptive approach to fairness for machine learning,” in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (New York, NY: ACM), 2459–2468.



- Stadler, T., Oprisanu, B., and Troncoso, C. (2022). "Synthetic data - anonymisation groundhog day," in *31st USENIX Security Symposium, USENIX Security 2022*, eds. K. T. B. Butler, and K. Thomas (Boston, MA: USENIX Association), 1451–1468.
- Stix, C. (2021). Actionable principles for artificial intelligence policy: three pathways. *Sci. Eng. Ethics* 27:15. doi: 10.1007/s11948-020-00277-3
- Sundararajan, M., Taly, A., and Yan, Q. (2017). "Axiomatic attribution for deep networks," in *International Conference on Machine Learning (PMLR)*, 3319–3328.
- Tagiou, E., Kanellopoulos, Y., Aridas, C., and Makris, C. (2019). "A tool supported framework for the assessment of algorithmic accountability," in *2019 10th International Conference on Information, Intelligence, Systems and Applications (IISA)* (New York, NY: IEEE), 1–9.
- Tai, B., Li, S., Huang, Y., and Wang, P. (2022). "Examining the utility of differentially private synthetic data generated using variational autoencoder with tensorflow privacy," in *27th IEEE Pacific Rim International Symposium on Dependable Computing, PRDC 2022* (Beijing: IEEE), 236–241.
- Thiebess, S., Lins, S., and Sunyaev, A. (2021). Trustworthy artificial intelligence. *Electron. Mark.* 31, 447–464. doi: 10.1007/s12525-020-00441-4
- Toth, Z., Caruana, R., Gruber, T., and Loebbecke, C. (2022). The dawn of the ai robots: towards a new framework of ai robot accountability. *J. Bus. Ethics* 178, 895–916. doi: 10.1007/s10551-022-05050-z
- Van den Broek, E., Sergeeva, A., and Huysman, M. (2021). When the machine meets the expert: an ethnography of developing ai for hiring. *MIS Q.* 45:16559. doi: 10.25300/MISQ/2021/16559
- Vasudevan, S., and 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, NY: ACM), 2773–2780.
- Verma, S., and Julia, R. (2018). "Fairness definitions explained," in *Proceedings of the International Workshop on Software Fairness (FairWare'18)* (ACM), 1–7.
- Wachter, S., Mittelstadt, B., and Russell, C. (2021). Why fairness cannot be automated: bridging the gap between eu non-discrimination law and AI. *Comp. Law Sec. Rev.* 41:105567. doi: 10.1016/j.clsr.2021.105567
- Wagner, I., and Eckhoff, D. (2018). Technical privacy metrics: a systematic survey. *ACM Comp. Surv.* 51, 1–38. doi: 10.1145/3168389
- Wang, B. Y., Boell, S. K., Riemer, K., and Peter, S. (2023). "Human agency in ai configurations supporting organizational decision-making," in *ACIS 2023 Proceedings* (Atlanta, GA: AIS), 1–22.
- Wang, D., Yang, Q., Abdul, A., and Lim, B. Y. (2019). "Designing theory-driven user-centric explainable AI," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (New York, NY: ACM), 1–15.
- Wang, Z., Yang, H., Feng, Y., Sun, P., Guo, H., Zhang, Z., et al. (2023). "Towards transferable targeted adversarial examples," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (New York, NY: IEEE), 20534–20543.
- Weber, L., Lopuschkin, S., Binder, A., and Samek, W. (2023). Beyond explaining: opportunities and challenges of XAI-based model improvement. *Inf. Fus.* 92, 154–176. doi: 10.1016/j.inffus.2022.11.013
- Wei, K., Li, J., Ding, M., Ma, C., Yang, H. H., Farokhi, F., et al. (2020). Federated learning with differential privacy: algorithms and performance analysis. *IEEE Trans. Inf. Forensics Secur.* 15, 3454–3469. doi: 10.1109/TIFS.2020.2988575
- Weisz, J. D., He, J., Muller, M., Hofer, G., Miles, R., and Geyer, W. (2024). "Design principles for generative ai applications," in *Proceedings of the CHI Conference on Human Factors in Computing Systems* (New York, NY: ACM), 1–22.
- Weng, T. W., Zhang, H., Chen, P. Y., Yi, J., Su, D., Gao, Y., et al. (2018). "Evaluating the robustness of neural networks: an extreme value theory approach," in *6th International Conference on Learning Representations, ICLR 2018*, 1–18.
- Wieringa, M. (2020). "What to account for when accounting for algorithms: a systematic literature review on algorithmic accountability," in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT\* '20* (New York, NY: Association for Computing Machinery), 1–18.
- Wing, J. M. (2021). Trustworthy AI. *Commun. ACM* 64, 64–71. doi: 10.1145/3448248
- Winter, P. M., Eder, S., Weissenböck, J., Schwald, C., Doms, T., Vogt, T., et al. (2021). Trusted artificial intelligence: towards certification of machine learning applications. *arXiv [preprint]*. doi: 10.48550/arXiv.2103.16910
- Wolf, C. T. (2019). "Explainability scenarios: towards scenario-based XAI design," in *Proceedings of the 24th International Conference on Intelligent User Interfaces* (New York, NY: ACM), 252–257.
- Xia, B., Lu, Q., Zhu, L., Lee, S. U., Liu, Y., and Xing, Z. (2024). "Towards a responsible AI metrics catalogue: a collection of metrics for ai accountability," in *Proceedings of the IEEE/ACM 3rd International Conference on AI Engineering - Software Engineering for AI, CAIN '24* (New York, NY: Association for Computing Machinery), 100–111.
- Xu, H., Ma, Y., Liu, H.-C., Deb, D., Liu, H., Tang, J.-L., et al. (2020). Adversarial attacks and defenses in images, graphs and text: A Review. *Int. J. Automat. Comp.* 17, 151–178. doi: 10.1007/s11633-019-1211-x
- Yeung, K. (2020). Recommendation of the council on artificial intelligence (OECD). *Int. Legal Mater.* 59, 27–34. doi: 10.1017/ilm.2020.5
- Yin, H., Mallya, A., Vahdat, A., Álvarez, J. M., Kautz, J., and Molchanov, P. (2021). "See through gradients: image batch recovery via gradinversion," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual* (New York, NY: IEEE), 16337–16346.
- Young, K., Booth, G., Simpson, B., Dutton, R., and Shrapnel, S. (2019). "Deep neural network or dermatologist?," in *Interpretability of Machine Intelligence in Medical Image Computing and Multimodal Learning for Clinical Decision Support: Second International Workshop, iMIMIC 2019, and 9th International Workshop, ML-CDS 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 17, 2019, Proceedings 9* (New York, NY: Springer), 48–55.
- Zaeem, R. N., and Barber, K. S. (2020). The effect of the gdpr on privacy policies: recent progress and future promise. *ACM Transact. Manag. Inf. Syst.* 12, 1–20. doi: 10.1145/3389685
- Zemel, R., Wu, Y., Swersky, K., Pitassi, T., and Dwork, C. (2013). "Learning fair representations," in *International Conference on Machine Learning (PMLR)*, 325–333.
- Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., and Gao, Y. (2021). A survey on federated learning. *Knowl. Based Syst.* 216:106775. doi: 10.1016/j.knsys.2021.106775
- Zhao, L., Hu, Q., and Wang, W. (2015). Heterogeneous feature selection with multi-modal deep neural networks and sparse group lasso. *IEEE Transact. Multimed.* 17, 1936–1948. doi: 10.1109/TMM.2015.2477058
- Zhu, X., and Wu, X. (2004). Class noise vs. attribute noise: a quantitative study. *Artif. Intell. Rev.* 22, 177–210. doi: 10.1007/s10462-004-0751-8
- Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., et al. (2020). A comprehensive survey on transfer learning. *Proc. IEEE* 109, 43–76. doi: 10.1109/JPROC.2020.3004555
- Zimmerman, M. (2018). *Teaching AI: Exploring new frontiers for learning*. Arlington, VA: International Society for Technology in Education.
- Zimmermann-Niefeld, A., Turner, M., Murphy, B., Kane, S. K., and Shapiro, R. B. (2019). "Youth learning machine learning through building models of athletic moves," in *Proceedings of the 18th ACM International Conference on Interaction Design and Children* (New York, NY: ACM), 121–132.
- Zou, Y., and Schaub, F. (2018). "Concern but no action: consumers' reactions to the equifax data breach," in *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (New York, NY: ACM), 1–6.