

# On the Need to Rethink Trust in AI Assistants for Software Development: A Critical Review

Sebastian Baltes<sup>†\*</sup>, Timo Speith<sup>†‡\*</sup>, Brenda Chiteri<sup>†</sup>, Seyedmoein Mohsenimofidi<sup>†</sup>,  
Shalini Chakraborty<sup>‡</sup>, Daniel Buschek<sup>‡</sup>

<sup>†</sup>Heidelberg University, Germany

<sup>‡</sup>University of Bayreuth, Germany

<sup>†</sup>{sebastian.baltes, brenda.chiteri, s.mohsenimofidi}@uni-heidelberg.de

<sup>‡</sup>{timo.speith, s.chakraborty, daniel.buschek}@uni-bayreuth.de

**Abstract**—Trust is a fundamental concept in human decision-making and collaboration that has long been studied in philosophy and psychology. However, software engineering (SE) articles often use the term *trust* informally; providing an explicit definition or embedding results in established trust models is rare. In SE research on AI assistants, this practice culminates in equating trust with the likelihood of accepting generated content, which, in isolation, does not capture the full conceptual complexity of trust. Without a common definition, true secondary research on trust is impossible. The objectives of our research were: (1) to present the psychological and philosophical foundations of human trust, (2) to systematically study how trust is conceptualized in SE and the related disciplines human-computer interaction and information systems, and (3) to discuss limitations of equating trust with content acceptance, outlining how SE research can adopt existing trust models to overcome the widespread informal use of the term trust. We conducted a literature review across disciplines and a critical review of recent SE articles with a focus on trust conceptualizations. We found that trust is rarely defined or conceptualized in SE articles. Related disciplines commonly embed their methodology and results in established trust models, clearly distinguishing, for example, between *initial trust* and *trust formation* and between *appropriate* and *inappropriate trust*. On a meta-scientific level, other disciplines even discuss whether and when trust can be applied to AI assistants at all. Our study reveals a significant maturity gap of trust research in SE compared to other disciplines. We provide concrete recommendations on how SE researchers can adopt established trust models and instruments to study trust in AI assistants beyond the acceptance of generated software artifacts.

**Index Terms**—Software Engineering, AI Assistants, Large Language Models, Human-Computer Interaction, Information Systems, Trust, Literature Review, Critical Review

## I. INTRODUCTION

Since the launch of ChatGPT in November 2022 and the subsequent hype around generative artificial intelligence (GenAI), AI assistants that help humans generate images, text, or source code have become an integral part of the daily work of many knowledge workers around the world [100]. GenAI tools, for example ChatGPT and GitHub Copilot, support software engineers in tasks such as understanding existing code, fixing bugs, or implementing new features [74], [147]. GitHub claims that, as of February 2023, Copilot is “*behind an average of 46 percent of developers’ code across*

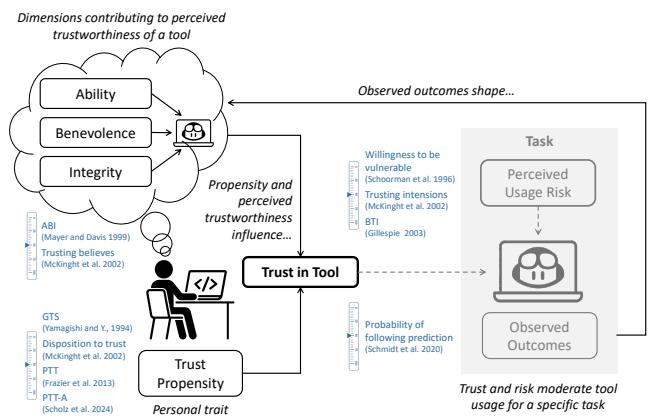


Fig. 1. Trust model by Mayer et al. applied to AI assistants with references to existing measurement instruments [43], [48], [97], [102], [132], [134], [135].

all programming languages.”<sup>1</sup> The IDE integration of GitHub Copilot and similar AI development assistants allows tool vendors to collect fine-grained usage data to understand which suggestions developers accept, which they reject, and which they modify before committing a change [176].

Previous work coauthored by employees of companies with commercial interests in increasing tool adoption, including Microsoft, Google, and Amazon, investigated what they refer to as **trust** in AI assistants [24], [31], [66], [107], [125], [161]. In some of those articles, the concept of *trust* is operationalized as tool adoption in general or, more specifically, as a high **acceptance rate** for generated content. In short, they equate trust with a high probability of developers accepting generated software artifacts, while ignoring aspects like trust formation and the antecedents of trust (see Figure 1). Compared to existing conceptualizations of trust in foundational disciplines such as psychology and philosophy, but also more applied disciplines such as human-computer interaction (HCI) and information systems (IS), we consider this view to be **too simplistic**. It **disregards decades of research on trust**.

Although embedding software engineering (SE) research in existing models and theories of trust is not widespread (see Section III), our community does have a history of discussing

\*The first two authors contributed equally to this work.

<sup>1</sup>GitHub Blog: GitHub Copilot now has a better AI model

**trust in source code.** For example, in his Turing Award lecture in 1984, Ken Thompson addressed trust in foreign code, stating: “*You can’t trust code that you did not totally create yourself... No amount of source-level verification or scrutiny will protect you from using untrusted code*” [154]. Given the widespread adoption of AI-based code generation and the persistent issue of hallucinations [14], it becomes imperative to adopt a more **holistic perspective on trust**, one that extends beyond the notion of accepting AI-generated content. Combining Ken Thompson’s statement with the increasing amount of generated code, this leads to a critical question: *Can developers genuinely trust generated software artifacts, or is there a conceptual difference between trust and mere acceptance of generated content?*

Our analysis of conceptualizations of trust in other disciplines reveals that viewing trust as mere acceptance of generated content is problematic for two key reasons. First, it **lacks a solid theoretical foundation** because it is unclear whether what is being measured is actually aligned with common conceptualizations of trust. People’s behavior, and hence their likelihood of accepting generated content, can be influenced not only by the perceived trustworthiness of the AI they are using, but also by other factors (see Figure 1 for an example). Factors that affect trust are not necessarily directly related to trust [82], [98]. Such factors include people’s specific situation (e.g., time pressure), the type of AI system with which they are interacting, the specific task they are performing, and the potential consequences of decisions.

Second, viewing trust as mere acceptance of generated content is **potentially harmful**. Even if trust could be measured through acceptance alone, this would not imply that the measured trust is appropriate trust. Given the potential maintainability challenges posed by generated code [89], security risks associated with misplaced trust in AI-generated code [75], [80], [93], [113], [115], [116], [129], [137], and the potential negative impact on developers’ learning process [138], [169], researcher need to distinguish **appropriate** from **inappropriate trust**. The fact that many developers distrust the accuracy of AI assistants<sup>2</sup> but nevertheless adopt these tools indicates that there is more to adoption than just trusting a tool.

We advocate for a broader interdisciplinary approach to understanding trust in SE. Our goal is to **motivate the SE research community to rethink how we conceptualize and measure trust**, especially in the context of AI assistants for software development. We would like to establish that software cannot be developed with the goal of “increasing trust” without knowing what trust actually entails. Without a scientifically validated trust model, trust becomes an empty marketing concept or, even worse, a concept used for *ethics washing* (see Section II-C). Without an established trust model, there is no reference point to compare claims of “increased trust” to and check whether they are actually true. More specifically, without a clear understanding of the concept of trust, it is difficult to consider trust in design goals, which in turn makes it difficult to design interactive AI assistants for

software development in a way that addresses concrete trust-related challenges.

Note that in this paper, we focus on **human or human-like trust**, not trust in a more technical sense (e.g., the implementation of chains of trust in a security context). Our target audience are SE researchers planning to design studies on trust as well as editors and reviewers of SE research who are responsible for ensuring our communities’ epistemological and methodological standards when it comes to researching interdisciplinary topics such as trust. To start a community discourse on rethinking trust and to support researchers who want to adopt and adapt established trust models, we address the following research questions (RQs):

- RQ1** How is trust conceptualized in foundational disciplines and applied disciplines related to SE?
- RQ2** How is trust conceptualized in SE?
- RQ3** How can we adopt existing trust models for SE research to study trust in AI assistants?

After introducing the foundations of human trust based on research in psychology and philosophy (Section II), we present the results of a literature review (Section III) on conceptualizations of trust in HCI and IS (**RQ1**) as well as SE (**RQ2**). We complement that literature review with a critical review of trust conceptualizations in recent SE articles (**RQ2**) and then synthesize the results of the two literature reviews into a discussion on how to rethink trust in AI assistants for software development (**RQ3**), including actionable recommendations for SE researchers (Section IV). We then discuss threats to validity (Section V) and conclude the paper (Section VI).

## II. BACKGROUND: FOUNDATIONS OF TRUST

*Trust* and the related concept of *trustworthiness* have been discussed and conceptualized long before disciplines such as SE, HCI, or IS recognized these concepts. Therefore, before we dive into our cross-disciplinary literature review of trust, we outline how philosophy and psychology, that is, disciplines that have been studying the foundations of human behavior for a long time, conceptualize *human trust*.

First, it is important to distinguish *trust* from *trustworthiness*. Both philosophy and psychology agree that *trust* is an attitude that one entity (the *trustor*) has toward another entity (the *trustee*) that it hopes will be *trustworthy*. Hence, leaving *perceived trustworthiness* aside, *trustworthiness* is considered a property of the *trustee* and not an attitude [82], [104], [105]. Furthermore, both disciplines agree that *trust* and *trustworthiness* are contextual matters: a *trustor* may trust a *trustee* in one situation or for one specific task while, at the same time, not trusting them in another situation or for another task. Typically, the *trustor* is human, while it is debatable whether the *trustee* must also be human (see Section II-B).

### A. Psychology

In psychology, two models have been very influential in the discourse on trust. Those models have also been adopted in other fields. In 1995, Mayer et al., surveying the literature on trust, created the *ability, benevolence, and integrity model of trust* [98] (see Figure 2). They defined trust as “*the*

<sup>2</sup>2025 Stack Overflow Developer Survey

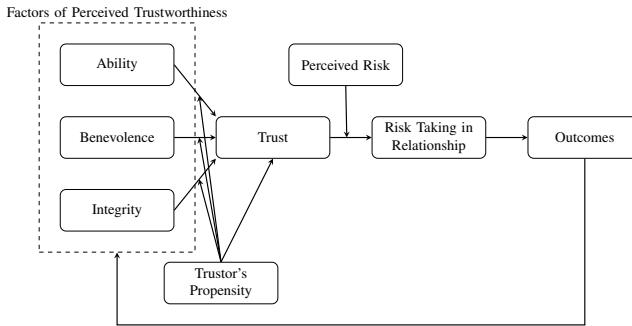


Fig. 2. Trust model by Mayer et al. (own illustration based on [98]).

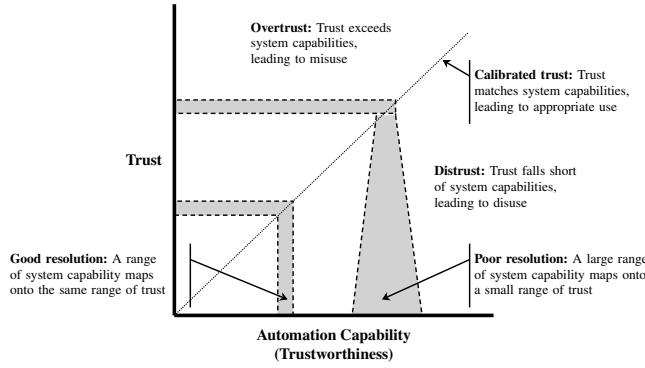


Fig. 3. Trust model by Lee and See (own illustration based on [82]).

willingness of a party to be vulnerable to the actions of another party based on the expectation that the other will perform a particular action important to the trustor, irrespective of the ability to monitor or control that other party" [98, p. 712].

In this model, *ability* refers to the trustee's competence in a given domain, *benevolence* captures the trustee's (assumed) goodwill and positive intentions toward the trustor, and *integrity* refers to the trustee's adherence to ethical principles the trustor agrees with. The model suggests that trust is high when all these three components of trustworthiness are perceived by the trustor as high, while deficiencies in any factor can weaken trust. How much (changes in) the factors affect trust depends on the trustor's *propensity to trust*. This propensity also influences the trustor's *initial level of trust*. Trust can manifest itself in specific situations, depending on the level of trust and the *risk* perceived by a trustor when they trust the trustee to perform a certain task. Depending on the outcome of this risk-taking, the perception of the three factors can change.

Lee and See take this model, expand it and apply it to automation contexts [82]. They introduce the concepts of *calibration*, *resolution*, and *specificity* (see Figure 3). *Calibration* concerns the crucial question of whether the trust a trustor has corresponds to the real capabilities, that is, the actual trustworthiness of a system. Ill-calibrated trust may lead to disuse (in the case of *dis-/undertrust*) or misuse (in the case of *overtrust*) of a system. *Resolution* is about how changes in system capabilities change levels of trust. For example, with low resolution, high changes in capability would lead to small changes in trust. Finally, *specificity* reflects how granular the trust is. Specificity emphasizes the context-dependence of trust

by stating how much trust depends on the capabilities of a particular system or the situation. According to Lee and See, good calibration, high resolution, and high specificity of trust can mitigate misuse and disuse of automation systems and enhance human-automation partnerships. Although they also introduce a model on the interaction of context, agent characteristics, and cognitive properties with the appropriateness of trust, their work is mostly cited for making the above distinctions.

### B. Philosophy

McLeod has analyzed philosophical debates on trust and distills three requirements for trust [104]. In summary, trust requires that one can:

- 1) be *vulnerable* to others, in particular to betrayal,
- 2) rely on others to be *competent* to do what one wishes to trust them to do, and
- 3) rely on them to be *willing* to do it.

For most philosophers, trust is a kind of *reliance*, although it involves "*some extra factor*" [104]. This extra factor is often determined by how philosophers interpret the three requirements listed above.

For example, regarding the factor of *vulnerability to betrayal*, trust is often connected to *anti-monitoring* (see, e.g., [10], [73]). The idea here is that if a person monitors an entity, this person cannot be betrayed by that entity. At most, the person could be disappointed if the entity does not perform as intended. Here, there is an interesting difference to Mayer et al.'s theory of trust, which explicitly decouples monitoring from trust. However, their theory is flexible enough to capture the same intuition concerning monitoring situations. For example, an employee may devalue the *benevolence* of (and thus the trust in) their employer due to constant monitoring.

The factors of *competence* and *willingness* (i.e., motivation) have received the most attention in the debate, with many competing theories that exist [104]. These theories range from the trustee acting out of self-interest [53] or goodwill [10], [68] to theories of networks of trust [44]. In summary, while there is a minimum consensus of the three requirements mentioned above, philosophers differ so much in the exact formulations of these requirements that some even assume that there are different types of trust. This would mean that trust is not one form of reliance but several different ones [30], [130], [143].

Philosophers place more emphasis on the *preconditions of trust* and the conditions of its *formation* than psychologists. Therefore, in addition to the concept of *calibrated trust* (which they call *well-grounded trust*), they also introduce the concepts of *justification* and *plausibility* trust [104]. *Justification* is about the indicators that have been consulted for trust formation, that is, whether it is based on good evidence. If a person happens to trust the system in the right way, but does so by chance without any indications, then philosophers would not speak of justified trust. Finally, *plausibility* is about whether the preconditions for trust are satisfied. For example, if two parties are antagonistic, there can be no trust between them.

*Justified* trust and *well-grounded* trust are independent of each other. A person might be justified in their trust, for instance, by receiving reports on the trustee's performance.

However, these reports might be forged, significantly overestimating actual performance, so that the ensuing trust would be ungrounded. Similarly, the above example of the trustor trusting the trustee to the right level by chance is a case of well-grounded but unjustified trust.

### C. Philosophy and Psychology on Trust and AI

Philosophical and psychological discussions about trust continue to this day, often in specific areas of application. AI is one of them, especially in recent times. There are three discussions that are worth highlighting here. In particular, with regard to the *plausibility of trust*, there are discussions as to whether it makes sense to apply the term trust to AI at all. From a philosophical perspective, current AI is not conscious, so it does not make sense to speak of it ‘betraying’ us.

This discussion even made it into the media when the *Guidelines for Trustworthy AI* of the high-level AI expert group set up by the European Commission were published. One of the members of this group, the German philosopher Thomas Metzinger, gave an interview shortly before the guidelines were published, in which he complained about the strong involvement of companies in the design of the guidelines. He claimed that these would use the principle of *trustworthy AI* for *ethics washing*, because in his opinion, the term trustworthy AI is conceptual nonsense.<sup>3</sup> While some authors agree with Metzinger in that the concept of trust cannot be applied to AI systems, for example, since they do not possess emotive states or cannot be held responsible [37], [124], others defend the possibility to trust AI, for instance, by showing how philosophical theories of trust can be adapted to AI [30], [62], [110], [126], [158] or by differentiating trust in AI from trust in humans [173] (see [38] for an overview of the debate).

This leads to a second, more practice-oriented discussion: How is trustworthiness operationalized? Irrespective of the discussion as to whether the concept of trust can be meaningfully applied to AI, the *Guidelines for Trustworthy AI* had strong effects on the subsequent operationalization of AI trustworthiness (and with that trust in AI). By specifying requirements that are essential for trustworthiness, e.g. technical robustness, safety, transparency, and accountability, they provide fertile ground for psychological and philosophical research to operationalize these requirements more precisely (see, e.g., [86], [96]). Encompassing some of these requirements, especially technical robustness, *computational reliability* came to be an important factor for trustworthiness in philosophical discussions (see, e.g., [41]).

The last noteworthy discussion revolves around the relationship between explainable AI (XAI) and trust in AI. Increasing trust is one of—if not the—primary goal of XAI [70], [81]. Kästner et al. refer to this as the *explainability-trust hypothesis* [70]. However, studies show that providing explanations does not consistently increase trust, and in many cases, even decreases it [70], [81]. Here, trust calibration plays an important role, as explanations that reveal system errors should lead to lower trust. Unfortunately, psychological biases

influence this process, as people sometimes trust a system simply because of the illusion of an explanation [40]. On the same topic, Ferrario and Loi argue, drawing on the concept of anti-monitoring, that XAI cannot directly contribute to trust, as it constitutes a form of monitoring [41]. However, they argue that XAI can indirectly support trust because human-AI interaction that is supported by XAI is generally more trustworthy.

Regarding biases and the possibility of generating problematic content, there is now more and more research that argues that, for a system to be adequately trusted, it is not just the system that needs to function correctly. Rather, for a system to be perceived as benevolent, there must be regulation, auditing processes, and established standards that govern the entire engineering process [94], [117], [164]. Winfield and Jirotka, for instance, argue that ethical governance, consisting of ethical codes of conduct, ethics training, responsible innovation, and transparency, is essential to build trust in AI systems [164]. Similarly, Manzini et al. argue that user trust requires evidence on the functionality of an AI assistant at the level of its design, the practices of the developing organization, and oversight by external bodies [94].

#### RQ1 (TRUST CONCEPTUALIZATIONS):

In **psychology**, the *ability, benevolence, and integrity* model of trust is common. How much (changes in) these factors affect trust depends on the trustor’s *propensity to trust*, which again affects the trustor’s *initial level of trust*. A central aspect is the *risk* perceived by a trustor when they trust the trustee to perform a certain task. The model was later extended to cover *calibration*, *resolution*, and *specificity* of trust. E.g., ill-calibrated trust may lead to *disuse* (in the case of *undertrust*) or *misuse* (in the case of *overtrust*) of a system.

In **philosophy**, *anti-monitoring* (not being able to monitor the trustee) is essential. Philosophers place more emphasis on the *preconditions of trust* and the conditions of *trust formation* than psychologists. They refer to *calibrated trust* as *well-grounded trust* and add the concepts *justification* and *plausibility* of trust. There are ongoing discussions whether trust can be meaningfully applied to AI.

## III. LITERATURE REVIEW (RQ1 AND RQ2)

After outlining the foundations of trust, we now turn to the methodology and results of our review of trust conceptualizations across the disciplines IS, HCI, and SE.

### A. Methodology

The goal of our research was to understand how trust is conceptualized in SE (RQ2) compared to other disciplines (RQ1) and, based on these conceptualizations, to explore how existing trust models can be adapted to study AI assistants in software development (RQ3). To achieve this, we conducted a literature review across disciplines.

<sup>3</sup>EU guidelines: Ethics washing made in Europe (Tagesspiegel)

1) **Literature review across disciplines:** We decided to use *Google Scholar* for the literature review, because portals such as the *ACM Digital Library* and the *IEEE Xplore Digital Library*, which are commonly used for literature reviews in SE, do not cover common publication venues in IS. We used the Python library *scholarly* to automate the retrieval of articles using a common search query across disciplines:

FIELD (trust OR trustworthiness OR trustworthy)

The values for FIELD were ‘software engineering’, ‘human-computer interaction’, and ‘(management OR information systems)’. We selected HCI and IS because they are closely related to SE, but we expected them to have a more mature view on trust and its conceptualizations, in particular because both HCI and IS have a strong connection to psychology. For the IS search queries, we added ‘management’ because information systems and management have overlapping publication venues. In fact, the most reputable IS venue is called *Management Information Systems Quarterly* (see the rankings linked from the inclusion criteria listed below). Since our goal was to apply existing trust models to AI assistants, we further retrieved AI-specific articles by adding the following to the before-mentioned queries:

(ai OR genai OR artificial intelligence)

For each discipline, we created two lists based on these queries (e.g., HCI and HCI-AI, see Table II) and retrieved 50 articles from each list, selecting those with the highest search rank as of March 2025, resulting in a total of 300 articles to review. We then identified and marked duplicates across the lists. We assigned one author to each discipline who then reviewed the articles according to the inclusion and exclusion criteria we defined. We organized weekly meetings to discuss unclear cases and applied the following criteria for the literature review across disciplines that we conducted to address **RQ1** and **RQ2**:

- **Include** papers published in the target discipline (use common rankings such as *ICORE*<sup>4</sup> for SE or the *ABDC JQL*<sup>5</sup> for IS; do not filter by rating, only exclude venues that are not listed in discipline-specific rankings).
- **Include** papers published in the main or research tracks of conferences and regular journal papers. Include literature reviews and special issues on human trust (analyzed separately, see Section III-B4).
- **Exclude** other editorials, short papers (< 5 pages), vision papers, textbooks, and dissertations.
- **Exclude** papers that do not discuss human/human-like trust but trust in a technical sense (e.g., in a security context).

For SE, we included 13 unique primary research papers (one paper was present in both lists), seven literature reviews (of which two were also part of the IS list; only three were SE-specific), and one special issue that overlaps with IS. For HCI, we included 22 unique primary research papers (five papers were present in both lists) and five literature reviews (two present in both lists; two overlapping with IS). For IS,

we included 24 unique primary research papers (no overlap), four literature reviews (two also present in the SE lists, three without a strong IS focus), and three special issues on trust (one overlapping with SE). Among the excluded SE papers were two HCI papers, which we added to the list of HCI papers (hence we have 24 HCI research papers in total). Similarly, the HCI list contained one IS paper, which we added to the list of IS papers (hence we have 25 IS research papers in total). Across the SE, HCI, and IS lists, we identified 15 articles that were published in philosophy venues (including ethics and fairness) or psychology venues. We read those papers and either integrated them into the background (Section II) or the discussion (Section IV) section.

Table II lists all the included papers. Row PH+PS lists all philosophy and psychology papers that we integrated as mentioned above. For the other papers, we downloaded the corresponding PDF files, searched for all occurrences of the substring “trust,” and extracted all paragraphs that conceptualized trust either directly or indirectly. We also kept track of which trust models the papers referred to (if any) and whether they customized (adapted or extended) existing trust models. We report the results of this analysis in Section III-B.

Many of the excluded SE papers were either not full papers or published in non-SE venues. This, in conjunction with the fact that only few SE papers contained explicit conceptualizations of trust, contributed to the decision to follow up with a critical review of recently published articles on trust in leading SE venues (see below).

2) **Critical review of recent SE literature:** A critical review focuses on evaluating key issues and analyzing a sample of studies with shared characteristics [12]. The issue we intended to evaluate is the lack of trust conceptualization in SE research (related to **RQ2**). To this end, we used *dblp SPARQL* to retrieve all articles that met the inclusion/exclusion criteria:

- **Include** papers in leading SE conferences (ICSE, ASE, FSE) and journals (TSE, TOSEM) published recently, i.e., just before or after the GenAI hype (2022–2024).
- **Include** papers published in the main or research tracks of conferences and regular full journal papers.
- **Exclude** editorials, special issues, short papers, vision papers, papers from co-located events, etc.

For the 2670 papers that met our inclusion criteria, we retrieved all abstracts from the ACM and IEEE digital libraries and searched for the substring “trust” in the title and abstract. Of the 2670 papers, 58 mentioned trust. We downloaded the corresponding PDF files and manually searched for and extracted implicit and explicit conceptualizations of trust. We report the results of this analysis in Section III-C.

3) **Data availability:** We provide all retrieval and analysis scripts and all metadata we retrieved from *Google Scholar* and *dblp* as part of our supplementary material [13]. Moreover, we share the annotated lists of papers, including exclusion reasons, and all extracted excerpts. All included research papers are referenced in Table II.

### B. Trust Conceptualizations across Disciplines

In this section, we report the results of our literature review across disciplines. Note that with this review, our goal was

<sup>4</sup><https://portal.core.edu.au/conf-ranks/>

<sup>5</sup><https://abdc.edu.au/abdc-journal-quality-list/>

TABLE I  
COMMON TRUST MODELS ACROSS DISCIPLINES.

Trust Model	Summary	HCI	HCI-AI	IS	IS-AI
Mayer et al. (1995) [98]	Multi-disciplinary model of organizational trust.	[28], [33], [52], [112], [119], [122], [146]	[15], [52], [119], [146]	[85], [145], [151], [174]	[7], [36], [132], [153]
Lee&See (2004) [82]	Trust in automation systems and technology.	[28], [112], [119], [146], [163]	[15], [119], [146], [149], [172]	[145]	[7], [153]
McKnight et al. (2011) [101]	Model of trust in technology grounded in system-level traits.	[52], [146]	[52], [146]	None	[153]
Gulati et al. (2019) [52]	Empirically validated model for measuring intelligent systems trust.	[119], [146]	[119], [146]	None	None
McAllister (1995) [99]	Dual-factor model of organizational interpersonal trust.	[112], [122]	[15]	None	None
Rousseau et al. (1998) [123]	Integration of organizational trust across disciplines.	[28], [122]	None	[32]	None
Madsen&Gregor (2000) [92]	Human-computer trust model, focus on affective & cognitive trust.	[79], [112]	[15]	None	None
McKnight et al. (2002) [103]	Model of initial trust in e-commerce settings.	None	None	[32], [57], [84], [85], [152]	[159]
Gefen et al. (2003) [46]	Model of trust formation in e-commerce settings.	None	None	[57], [145], [150]	[157]

not to achieve a complete coverage of all articles that describe conceptualizations of trust in the target disciplines. Our goal was to identify whether articles in the selected disciplines commonly refer to established trust models, with the goal of comparing SE articles to related disciplines. Although we acknowledge this as a limitation of our research, the difference between the SE articles and the HCI/IS articles was so clear that our observations would most likely not change considerably when scaling up the analysis. Table I lists the trust models that at least two articles referred to. The models proposed by Mayer et al. [98] and Lee and See [82] are used across HCI and IS, while McKnight et al. [103] and Gefen et al. [46] are IS-specific. The other models listed in the table are primarily referenced in HCI, but less frequently than the Mayer et al. [98] and Lee and See [82] models.

1) **Information Systems and Management (IS):** In IS research, trust has been studied since the 1990s [144]. The articles we included in our review of IS literature were published in the time span from 2000 to 2024. Of the 24 research papers we analyzed, only four did not embed their studies in existing trust models. It was common for IS papers to customize existing trust models (13 out of 25 articles).

Several articles discussed that “*trust is a dynamic concept that develops over time*” [85]. The factors and processes that lead to *trust formation* change gradually [98], and can include aspects such as people’s perceptions about a company or brand [145]. This means that trust in a company can transfer to trust in a particular tool, a concept called *trust transference* [148], [160]. Saffarizadeh et al. applied this concept to AI agents and noted that “*trust in an AI creator can be transferred to its AI agent when there is a perceived meaningful association between them*” [127].

IS articles clearly distinguish *initial trust*, which is important to “*create a trustworthy first impression of a new system and encourage adoption*” [85], from the before-mentioned process of trust formation and later stages such as *knowledge-based trust*, “*where the individual knows the other party well enough to predict the party’s behavior in a situation*” [153]. From a tool perspective, the later phases of trust are some-

TABLE II  
INCLUDED ARTICLES OF OUR LITERATURE REVIEW; CR: CRITICAL REVIEW, ADDED: FROM SEARCH RESULTS OF OTHER DISCIPLINE.

Area	Research Articles	Literature Reviews	Special Issues	Added
HCI	[1], [27], [28], [33], [52], [61], [67], [79], [112], [119], [122], [140], [146]	[8], [51], [65], [108]	None	[163]
HCI-AI	[6], [15], [28], [45], [52], [60], [61], [119], [140]–[142], [146], [149], [167], [172]	[8], [28], [105]	[69]	[139]
IS	[5], [32], [50], [57], [84], [85], [95], [114], [145], [150]–[152], [174]	None	[17], [18]	None
IS-AI	[7], [35], [36], [54], [64], [77], [127], [132], [153], [155], [157]	[49], [71], [90], [170]	[91]	[159]
SE	[16], [20], [21], [23], [29], [55], [63], [66], [72], [87], [156], [162]	[26], [59]	None	None
SE-AI	[3], [66]	[71], [83], [88], [90], [175]	[91]	None
SE-CR	[4], [22], [78], [111], [165]	None	None	None
PH+PS	[30], [38], [56], [62], [86], [96], [110], [117], [118], [124], [133], [161], [164], [169]	None	None	None

times referred to as “*postadoption*”, where aspects such as *predictability* and *helpfulness* of a tool [121] play an important role [152]. McKnight et al.’s model is frequently referred to as a common IS model for initial trust, whereas Gefen et al.’s model focuses on trust formation.

In McKnight et al.’s model, based on Mayer et al., trust in vendors of e-commerce websites is a multidimensional construct centered around *trusting beliefs*, that is, “*perceptions*

of the competence, benevolence, and integrity of the vendor” and *trusting intentions*, that is, the “willingness to depend” or, in other words, the “decision to make oneself vulnerable to the vendor” [103]. In addition to a trust model, McKnight et al. also introduced instruments to measure the concepts in their model (see Figure 1 and Section IV).

Gefen et al. models consumers’ online purchase intentions as a combination of perceived usefulness, ease-of-use, and trust in the vendor [46]. They show that online trust is built through “(1) a belief that the vendor has nothing to gain by cheating, (2) a belief that there are safety mechanisms built into the web site, and (3) by having a typical interface, (4) one that is, moreover, easy to use.”

An important observation by Mayer et al. is repeated in IS articles: humans differ in their general tendency to be willing to trust others [85], that is, their *propensity* [32], [98] or *disposition* to trust [84]. These individual differences are a potential confounding factor in studies with developers investigating trust in AI assistants.

Two other central constructs are *trusting intention*, which is “*the trustor’s willingness to depend on the trustee*” [85] and a person’s *trusting beliefs* [153], that is “*favorable object-specific beliefs*” [152] that capture “*the trustor’s perception of whether the trustee has the desired attributes to be trusted*” [85] (see also McKnight et al.’s model). Moreover, IS literature discusses how the trusting beliefs proposed by Mayer et al. (benevolence, ability, and integrity) can foster technology acceptance [19]. Troshani et al. categorize those trusting beliefs into *cognitive trust* and *affective trust* [157] (see also Komiak and Benbasat [76]). Cognitive trust captures “*a customer’s confidence or willingness to rely on a service provider’s competence and reliability*” while *affective trust* is the “*confidence that one places in another party on the basis of feelings generated by the level of care and concern as demonstrated by the other party*.” In summary, competence (i.e., ability in Mayer et al.’s model) and predictability influence cognitive trust while benevolence and integrity influence affective trust [157].

The IS literature discusses the validity of applying models of human trust to technology [32], [85], [127], which is referred to as *plausibility of trust* in philosophy (see Section II-C). An aspect to consider is that humans have long been known to anthropomorphize technology [109]. To study trust relationships between humans and human-like tools, researchers have adapted the concept of *interpersonal trust* to technology, for example, to study trust in automation or autonomous systems (see the description of Lee and See’s trust model in Section III-B2). However, researchers have also argued that such models do not necessarily apply to GenAI tools due to their nondeterminism [7].

IS articles further discuss that increasing automation and complexity leads to systems being perceived as opaque and sophisticated, hence trust becomes more and more important [82]. This situation leads to uncertainty and a *perceived risk* of using an automation tool, which could hinder adoption [25], [152]. However, Chopra and Wallace also write that “*trust can only arise when there exists a state of dependence between the trustor and trustee, and when acting on this*

*dependence entails risk, i.e., the trustor possesses uncertainty about the outcomes and vulnerability to a potential loss if the outcomes are undesirable*” [32]. This is certainly true for software engineers who use GenAI tools to develop software. The role of risk and vulnerability appears also in the psychology literature, in particular Mayer et al. (see Figure 1), as well as in the philosophical debate summarized in Section II-B.

In the context of AI for decision-making, both *mistakenly denied trust* and *unfounded trust* can occur [132]. Schmidt et al. summarize these situations as follows: “*If humans increasingly leverage AI to inform, derive, and justify decisions, it also becomes important to quantify when, how, why, and under which conditions they tend to overly trust or mistrust those systems*” [132]. This view corresponds to the concepts of *undertrust* and *overtrust* as described by Lee and See [82].

Regarding *trust quantification*, Schmidt et al. note that “*trust is measured differently in different research fields*” [132]. However, they also mention that a common approach used across disciplines is to ask study subjects about their *trusting beliefs* or *trusting intentions* toward a certain entity, for which standardized survey instruments exist (cf. [47]). An exemplary operationalization of trust in decision-making based on observed behavior is calculating the probability that users follow a model’s predictions by utilizing usage data [120], [131]. This resembles our motivating example of software developers accepting code suggestions. We discuss trust operationalization and measurement in more detail in Section IV.

2) **Human-Computer Interaction (HCI):** Trust has been studied in HCI since the late 1980s [106]. Our review of HCI literature includes articles published between 1996 and 2024. Among the 24 unique HCI research papers we included, eight articles embedded their trust definitions in Mayer et al.’s model. Three articles explicitly cite the definition, one paraphrases it, and four reference components such as *ability*, *integrity*, or *benevolence* without adopting the complete model. Furthermore, nine papers refer to Lee and See’s trust model, and another five papers incorporate elements from both models. Among the twelve papers building upon existing trust models, six extended the models either by developing new models or by adopting the definitions and further exploring specific aspects of trust. The remaining six papers refer to the models without modification. In contrast, eleven papers do not explicitly refer to existing trust models.

Several HCI articles conceptualize trust as a *context-dependent phenomenon* shaped by system behavior, user interaction, and situational factors. This perspective is especially prevalent in domains such as AI agents, automated decision support systems, and human-robot interaction. For example, Sousa et al. [146] propose the *human-centered trustworthy framework*, which adapts Gulati et al.’s human–computer trust model [52] and integrates trustworthiness components from Mayer et al. [98], along with constructs from McKnight et al. [101] and Lee and See [82]. In this framework, trust is approached as a multilevel concept that includes *user predispositions*, *interface features* such as transparency and control, and *institutional factors* such as data governance and accountability. This layered view underscores the relational and evolving nature of trust in human–AI interaction.

Pinto et al. [119] developed and validated a human–robot interaction trust scale by adapting Gulati et al.’s model [52]. Trust is conceptualized following Mayer et al. [98]’s definition as the *willingness to be vulnerable to another party*, and is further informed by Lee and See’s view of *trust development over time* [82]. Their adapted model incorporates dimensions such as *reciprocity*, *competence*, *benevolence*, *predictability*, *honesty*, *trust predisposition*, and *willingness to be vulnerable*, tailored to human–robot interaction.

Nothdurft et al. [112] conceptualize trust as a multidimensional and dynamic construct that evolves through interaction. They draw on Mayer et al. [98] to describe trustworthiness in terms of *ability*, *integrity*, and *benevolence*, and incorporate Lee and See’s view of trust in automation [82] as a *calibrated attitude* shaped by system performance and user experience. Their conceptualization also reflects Madsen and Gregor’s adaptation [92] of McAllister’s definition [99], which emphasizes user confidence and willingness to act on the recommendations of intelligent systems. Building on this foundation, Nothdurft et al. [112] propose an adaptive explanation architecture that adjusts system transparency and feedback to support appropriate levels of trust in automated systems.

Riegelsberger et al. [122] conceptualize trust as a *psychological state* characterized by a *positive expectation* that one’s *vulnerability* will not be exploited. Like Nothdurft et al. [112], they build on Mayer et al. [98] and McAllister [99], but extend the discussion by emphasizing the *social grounding* of trust. Their framework is embedded in a trustor–trustee model adapted from the *trust game* [9], where trust is shaped by *contextual cues*, *reputation*, and *internalized norms*. They argue that trust in mediated interactions is not only influenced by *individual traits* but also by *perceived trustworthiness* of systems shaped through interface design and social context.

In contrast, several papers reference foundational trust models without adapting or extending them. Banovic et al. [15] cite Mayer et al. [98], McAllister [99] and Madsen and Gregor [92] to examine how untrustworthy AI systems can still elicit user trust, but do not propose a new framework or modify the referenced models. Similarly, Yokoi et al. [172] apply Lee and See [82] definition to investigate trust in AI-based medical decision-making, without extending it conceptually. Suen and Hung [149] adopt Lee and See’s trust model to frame user trust in AI-based video interviews. Although the model is applied to a specific domain, its theoretical structure remains unchanged. Kraus et al. [79] primarily build on Madsen and Gregor [92] trust model, while only briefly mentioning Mayer et al. and Lee and See. These models are not embedded in their conceptual framework and rather serve as peripheral citations. Similarly, Cai et al. [28] adopt Mayer et al.’s and Lee and See’s definitions of trust to conceptualize user trust in human–computer interaction, but use the models primarily to validate the Chinese version of the human–computer trust scale in a new cultural context.

3) **Software Engineering (SE):** Trust is a critical element in SE that influences the adoption of software tools, processes, and collaborative practices. However, our review of 13 articles published between 2002 and 2024 reveals that SE research typically does not engage with established foundational trust

models. In fact, only one paper [63] cited Mayer et al.’s trust definition and based their research on a related trust model. Furthermore, one paper [162] implicitly referred to established trust models by mentioning *risk*, that is, the aspect of *vulnerability* (see Section II-A). The other papers either did not explicitly define trust or used custom definitions, diverging from disciplines such as IS and HCI, where established trust models are more commonly applied.

An example of SE’s customized approach to trust is the PICSE framework proposed by Johnson et al. [66]. It describes a structured approach to understanding trust in software tools, identifying five key factors: *personal*, *interaction*, *control*, *system*, and *expectation*. The authors refer to Fogg and Tseng’s credibility framework [42] from HCI and Rousseau et al.’s trust literature review [123] from IS. However, these trust models did not contribute to PICSE but are discussed after introducing the framework.

Jalali et al. adapted Schultz’s situational trust model [136], which extends Mayer et al.’s model, to analyze trust dynamics in global software engineering [63]. In line with HCI and IS research, their model underlines the dynamic nature of trust, distinguishing *initial trust building*, where trust forms based on past interactions, from *trust evolution*, where trust fluctuates during the project based on expectations and behaviors.

Wang and Redmiles applied evolutionary game theory to examine trust and cooperation in distributed teams. Their study does not explicitly model trust development or draw on established trust theories [162]. Similarly, Calinescu et al. propose dynamic assurance cases for engineering trustworthy self-adaptive software, emphasizing system reliability and correctness rather than interpersonal trust [29]. Akbar et al. take a different approach, presenting a taxonomy of decision-making challenges in trustworthy AI [3]. Although these studies address trust-related concerns, they do so in domain-specific ways, emphasizing technical, ethical, or structural factors rather than explicitly modeling trust development based on established theories.

Bratthall and Jørgensen address trust in single data source software engineering case studies [23]. Bertram et al. investigate *trust perceptions* in code review through an eye-tracking study, demonstrating how the perceived origin of code (human versus machine-generated) influences review behaviors and trust judgments [20]. However, both papers do not explicitly reference or build upon established trust models.

4) **Literature Reviews and Special Issues:** Among the included papers were 11 literature reviews. In addition, we found three editorials for special issues on trust. Describing these articles in detail is beyond the scope of this paper. However, we briefly summarize them because they reflect the scientific discourse on trust in the corresponding communities.

As part of our review, we found five literature reviews of trust that were **not limited to just one of our target disciplines** SE, HCI, or IS. Kaur et al.’s review centers around requirements for trustworthy AI [71]. They refer to trust definitions across disciplines, including psychology, sociology, and economy, mentioning *integrity* and *reliability* as an “*agreement*” across these disciplines. They define trustworthy AI as a “*framework to ensure that a system is worthy of*

being trusted based on the evidence concerning its stated requirements” [71]. Yang and Wibowo examined 142 studies published between 2015 and 2022, identifying components, influencing factors, and outcomes of users’ trust in AI [170]. Li et al.’s review outlines aspects of AI trustworthiness such as *robustness, explainability, transparency, and fairness*, organized along the lifecycle of AI systems [83]. However, they do not discuss established trust definitions or models. Gulati et al. reviewed 47 studies on trust in technology published in IS and HCI venues and found 17 different theories that were integrated from disciplines such as psychology and economics [51]. They conclude that “*the intricacies of how trust is formed and maintained in online environments still necessitates further investigation*” and that the development of standardized and empirically validated instruments for trust measurement is crucial. Finally, Lockey et al.’s review focuses on *antecedents of trust* in AI [90]. They identify five AI-specific trust challenges: *transparency and explainability, accuracy and reliability, automation versus augmentation, anthropomorphism and embodiment, and mass data extraction* [90]. An interesting observation regarding AI assistants for software development is that “*over-anthropomorphism may lead to overestimation of the AI’s capabilities*” [90], increasing risk [34], decreasing trust [11], and potentially leading to manipulation [128]. This is certainly an underexplored perspective on AI assistants for software development, especially when focusing on human factors of software security.

Three literature reviews focused on trust research in **HCI**. Jeon’s review discussed “*the effects of emotions on trust in the context of technology use*” [65]. Their review indicates that positive emotions lead to higher trust. Bach et al. reviewed 23 empirical studies that focus on user trust in AI-enabled systems [8]. They describe trust definitions, factors that influence user trust, and measurement methods, reporting that seven studies explicitly defined trust, six of which referring to Mayer et al. [98] or Lee and See [82]. Eight studies conceptualized trust, while nine studies neither defined nor conceptualized trust. They found surveys, interviews, and focus groups to be the most commonly used methods to assess user trust. Bach et al. conclude that it is important to select a definition of user trust that aligns with the specific study context. Mehrotra et al. reviewed 65 articles to assess the various definitions of appropriate trust in human-AI interaction, providing an overview of concepts and definitions and identifying similarities and differences between them [105]. They conclude that four common methods for building appropriate trust are improving *system transparency*, considering user *cognition and perception*, using *guidelines or models* to achieve calibrated trust in AI, and considering the entire *continuum of trust*, including *over-, under-, mis-, and dis-trust*.

One literature review had a clear **IS** focus. Glikson and Woolley synthesized two decades of empirical research on the determinants of human trust in AI across disciplines [49]. The authors motivate how trust in AI differs from other technologies and identify an AI’s capabilities as an important antecedent of trust development. In addition, they propose a framework that identifies factors that influence users’ *cognitive* and *emotional* trust. Regarding the development of *cognitive*

*trust*, they highlight the importance of an AI’s *tangibility, transparency, reliability, and immediacy*. For *emotional trust*, on the other hand, and AI’s *anthropomorphism* is a major factor. Glikson and Woolley also critically reflect on past studies by pointing to the diversity of trust measure used and the lack of longer-term studies in “*higher stakes environments*” [49].

For **IS**, we further identified three special issues on trust. One of them was published in 2010 in their flagship journal *MIS Quarterly*, focusing on “*novel perspectives*” on trust in IS. The fact that this special issue was published 15 years ago is another indication of the maturity of trust research in that discipline. The other IS-specific special issues focused on trust in online environments [18] and trust in AI [91].

For **SE**, we included two literature reviews. Liu et al. conducted a tertiary study reviewing 141 secondary studies on trustworthy AI and software, describing trustworthiness as a “*highly abstract concept comprising related quality attributes*” [88]. The review focuses on these quality attributes and compares them between trustworthy AI and other trustworthy software. Liu et al. do not focus on the conceptualization of trust and trustworthiness. Hou and Jansen reviewed 112 articles on trust in software ecosystems, identifying different definitions of concepts such as *software trust, system trust, or software service trust*. However, their focus was on compiling trust definitions, not unifying or aligning them [59].

#### RQ2 (TRUST CONCEPTUALIZATIONS):

In **information systems** (IS), trust has been studied since the 1990s; in **human-computer interaction** (HCI) since the late 1980s. Both IS and HCI research is usually embedded in existing trust models (see Table I). The Mayer et al. (1995) model is common across IS and HCI, Lee and See (2004) is more common in HCI. Customizing or extending existing trust models is common in both disciplines. Important concepts mentioned in IS research are *initial trust, trust formation, and trust transference* (e.g., trust in a company transferring to a tool). Trust is sometimes studied in the context of *tool adoption* and *technology acceptance*, but as one of several factors. IS research commonly adopts or develops *validated instruments* to measure trust-related concepts. HCI research frames trust as a *context-dependent* phenomenon. **Software engineering** (SE) research typically does not engage with established trust models; trust as a concept is either not or only informally defined.

#### C. Critical Review of Trust Conceptualizations in Software Engineering Research

Of the 58 articles that we analyzed as part of our critical review, only one article referred to an established trust model, and only three articles provided trust definitions. This underlines the problematic situation suggested by our literature across disciplines (see Section III-B3). There is a significant gap in the maturity of trust research in SE compared to IS and HCI. Trust as a concept is often discussed in a context-specific manner without embedding concepts and definitions

in established trust theories or models. This lack of embedding in established trust theories from other disciplines hinders interdisciplinary comparisons and research toward a deeper understanding of what constitutes trust in SE-specific settings.

Noller et al. are the only authors who explicitly reference an established trust model, incorporating Lee and See's model of human-automation trust [111]. Alami et al. define trust as "*the unyielding belief that the person is truthful and reliable*" and in the context of their research as "*the willingness of the community to rely on the contributor, a prerequisite for considering her code change*" [4]. This conceptualization frames trust as a social and reputation-based construct, essential in collaborative environments such as open-source software development, where trust is built through ongoing contributions. Although this framing is reasonable, an embedding in existing trust models would allow researchers to identify the particularities of trust in open-source software compared to other settings.

Winter et al. [165] implicitly define trust by referencing dispositional trust as proposed by Hoff and Bashir [58], which describes an individual's inherent tendency to trust automation, independent of context or system performance. This positions trust as a psychological trait that varies between individuals. However, it is important to note that trust was only a peripheral topic, not a central focus of their study. Kou et al. discuss trust in the context of large language models (LLMs) for code generation, suggesting that users assess trustworthiness based on the model's explainability and how well its predictions align with their expectations [78]. However, they neither explicitly define trust nor consider the broader explainability discussion (see Section II). Instead, they cite previous work by Boggust et al., who explored how users determine whether to trust a machine learning model based on its ability to convey behavior in a way that aligns with their expectations [22]. Neither study provides a formal definition of trust; rather, both frame it in terms of user perceptions and interpretability. Given the amount of related work on trustworthiness and explainability of AI-based systems in philosophy (see Section II-C) and IS (see Section III-B1), it is surprising that this article is again not broadly embedded in established trust conceptualizations.

#### RQ2 (TRUST IN SE RESEARCH):

Our critical review confirmed that there is a **significant maturity gap** in trust-related research in SE compared to IS and HCI. Of the 58 analyzed articles, only one referred to an established trust model, and only three provided trust definitions. This lack of embedding in established trust models hinders interdisciplinary comparisons and research toward a deeper understanding of what constitutes trust in SE-specific settings.

#### IV. ADOPTING TRUST MODELS IN SE (RQ3)

To answer **RQ3**, we first discuss problems with the way SE research currently handles trust research and then describe how the common trust model of Mayer et al. [98] could be adopted to study trust in AI assistants for software development.

TABLE III  
SUMMARY OF THE TENETS OF TRUST AS DISCUSSED IN PHILOSOPHY AND PSYCHOLOGY COMPARED TO SOFTWARE ENGINEERING, WHERE TRUST HAS BEEN OPERATIONALIZED AS ARTIFACT ACCEPTANCE.

Tenets of trust in philosophy (PH) and psychology (PS)	Operationalizing trust as artifact acceptance...
Trust is an attitude, not a behavior (PH, PS).	...treats trust as a behavior.
Trust is context-dependent (PH, PS).	...just focuses on one context.
Trust should be well-calibrated or well-grounded (PH, PS).	...does not measure calibration or grounding.
Trust should have a high resolution (PS).	...does not measure resolution.
Trust should have a high specificity (PS).	...could at most measure low specificity.
Trust should be justified (PH).	... ignores justification.
Trust should be plausible (PH).	... ignores plausibility.

#### A. Current State of Trust Research in SE

Our research was motivated by the fact that several SE research articles suggested approaches to increase 'trust' with the aim of increasing tool adoption (see Section I). Although previous research in other fields such as IS suggests that trust is indeed related to tool adoption and usage (see Section III-B1), the corresponding research articles thoroughly define and conceptualize trust and related concepts such as *trust transference*, *trusting beliefs*, or a user's *propensity to trust*. These conceptualizations originate in mature trust models, in particular the one proposed by Mayer et al. in 1995 [98]. However, in terms of trust and AI, Thiebes et al. note that "*trust in general is a complex phenomenon that has sparked many scholarly debates in recent decades*", therefore "*it is not surprising that the conceptualization of trust in AI and what makes AI trustworthy—as of today—remains inconclusive and highly discussed in research and practice*" [153]. It is unfortunate that the SE research community is not a major contributor to that interdisciplinary discourse because, as our review of SE literature has shown, SE research (and practice, see [118]) on trust mainly uses ad hoc trust definitions without embedding studies in established trust models.

The detachment of SE research from established trust models is problematic, as evidenced by the equation of trust with the acceptance of generated content, which is shortsighted for several reasons. Drawing on the foundations of trust from psychology and philosophy that we introduced in Section II, Table III shows how such an operationalization disregards the tenets of trust discussed in these disciplines. Most importantly, both disciplines see trust as an *attitude* and not a *behavior*. Therefore, measuring acceptance disregards factors that influence a person's behavior, such as their situation (e.g., time pressure), the type of AI system they are interacting with, their specific task, the potential consequences of decisions, and many more. Still, even if acceptance measured trust, it would not measure *calibrated* or *warranted trust*. Hence, the measured trust could be completely detached from the actual capabilities of an AI assistant. However, given the fact that

GenAI-based software development assistants will continue to hallucinate [14], the potentially resulting maintainability challenges [89], security risks [75], [80], [93], [113], [115], [116], [129], [137], and negative impact on learning [138], [169], it would not only be morally but also pragmatically better to find out whether developers' trust in such tools is actually calibrated or warranted. After all, maintainability issues or security vulnerabilities introduced by relying on generated code could cause major financial and reputational damage to individuals and companies.

It was interesting for us to see that IS and HCI primarily focus on the psychological side of trust. We posit that discussion in SE can also benefit from incorporating philosophical insights on trust. For example, the discussion whether trust in AI is even conceptually possible is relevant in the sense that it critically questions whether engineering for trust is not just a marketing phenomenon or, in the worst case, even ethicswashing (see Section II-C). A similar phenomenon can be observed in the debates on XAI, where the expectation is that XAI invariably leads to greater trust, something that could not be proven in studies [70]. In this sense, the concept of *justification* is also important for SE, as it can help ensure that trust in AI assistants for software development is formed appropriately, rather than, for example, being the result of misleading marketing rhetoric. It is surprising that the trust conceptualizations in SE are still so underdeveloped, as tools from SE can be extremely important for creating trustworthy systems. For example, Ahuja et al. argue that SE processes and tools, in particular a detailed requirements analysis, monitoring instruments, and automated testing, are essential to create systems that can really be trusted [2].

### B. General Role of Trust in SE

In software development, negative outcomes, such as bugs or security vulnerabilities introduced by using AI assistants, materialize with a certain delay, and are often borne by someone else. Human reviewers must invest more time reviewing large amounts of generated code, security teams must catch vulnerabilities in production, bugs slow users down, and the organization may face a reputation loss due to data leakage or outages. This creates an incentive landscape in which an individual developer can rationally feel a low personal risk even when the system-level risk is unchanged or even increased. In trust terms, this can lead to an unjustifiably low perceived risk and potentially even to an unjustifiably high trust propensity.

Since developers are rewarded for throughput (e.g., resolved tickets), and costs of defects accumulate as technical debt, teams might drift into thinking (and trusting) that "it usually works."—while long-term maintenance costs are accumulating. This makes trust in software contexts unusually sensitive to incentives, governance, and feedback frequency. Such factors are different in settings where the consequences are immediate and personal. In SE, however, a developer might have already left the company before the problems, which their trust in AI assistants has caused, materialize.

Moreover, the object that is trusted can shift in an SE context. Software development is embedded in institutional-

ized safety nets that involve code reviews, software tests, and compliance checks, often automated or supported by continuous integration pipelines. These processes can be risk-reducing, so developers may come to trust the pipeline rather than the AI assistant ("*our tests will catch potential issues*"). This mindset can lead to "*trust laundering*", where weak trust in an AI output is converted into or substituted by greater trust in the infrastructure. The effect of this is that trust becomes procedural and distributed. It is not "*I trust the AI assistant*" but "*I trust our socio-technical infrastructure to detect and absorb errors*."

Another philosophical and legal issue is that AI assistants lead to ambiguous authorship and accountability. When AI suggestions are integrated, it can become even more unclear who is responsible for an issue: Is it the developer who used the AI assistant, the reviewer who approved the change, or the team whose tests did not catch the bug? If accountability becomes more diffuse, developers may feel less personally exposed, which again raises trust even though organizational risk remains the same. In contrast, in teams with strong code ownership, norms, and guidelines, the same ambiguity can reduce trust because accepting AI output threatens clear boundaries of responsibility. One could argue that trust really depends on the institutional practices in place.

### C. An Example of Adopting Trust Models in SE

The dimensions of perceived trustworthiness proposed by Mayer et al. [98], that is, *ability* (A), *benevolence* (B), and *integrity* (I) can be easily transferred to AI assistants for software development. As mentioned before, *ability* requires AI assistants to generate correct and reliable code that is appropriate for a given context. *Benevolence* reflects interaction patterns that AI vendors implement to ensure that AI assistants are perceived as supportive, and *integrity* models the data curation and guardrails that AI vendors and regulatory bodies implement to avoid biases or problematic code that, for example, contains security vulnerabilities. In the following, we discuss existing instruments that SE researchers can use, with a particular focus on our own adaption of Mayer et al.'s trust model as presented in Figure 1. Please note that this is one example for adapting an established trust model to an SE context. If concepts such as *calibrated trust* are important for a particular SE research project, the model by Lee and See might be more suitable [82]. It is important that SE researchers know enough about trust conceptualizations to understand the implications of adapting particular trust models themselves. We hope that our article serves as a starting point for this.

As already mentioned, operationalizing trust is difficult; one needs to be specific about the actual concepts that one intends to capture (e.g., the *perceived trustworthiness* of a tool or a person's general *propensity to trust*). Fortunately, validated instruments exist to assess Mayer et al.'s ABI dimensions of perceived trustworthiness. SE researchers can adapt these instruments to study trust in AI assistants [97], [102]. There further exist various instruments to assess a person's *propensity to trust* [43], [102], [134], [168]. Since perceived *risk* and *vulnerability* are central to Mayer et al.'s model, we also want

to refer to instruments designed to measure related concepts such as a person's *willingness to be vulnerable* or their *trusting intentions* [48], [102], [135].

As mentioned above, IS literature has discussed operationalizing trust based on observed behavior by calculating the probability that users follow a model's prediction based on usage data [120], [131]. Since this is closely related to our motivating example of software developers accepting code suggestions, we want to discuss how such metrics can be included in our representation of Mayer et al.'s trust model (Figure 1). Schmidt et al. argue that behavioral metrics might work well as proxies of trust when comparing different treatments (e.g., different ways of presenting AI suggestions) or comparing different groups based on the same treatment [132]. However, factors such as participant knowledge and confidence, and (perceived) task difficulty are known to affect advice acceptance [171]. Over time, the usage of tools moderates users' bias toward self-reliance in decision-making and could further influence the observed behavior [39].

In the context of Mayer et al.'s model, the acceptance rate for AI suggestions can be interpreted as an indicator of the tool's "observed outcome." Studying such data in detail is crucial, as acceptance of AI suggestions can be a factor by which people form their trust; that is, in Mayer et al.'s model, acceptance rates can contribute to the *perceived ability* of the system. Support for this view comes, for example, from Wang et al.'s study. They have developed and tested design concepts, including a dashboard showing a user's personal GitHub Copilot acceptance rate, finding that it helped developers align their expectations with an AI assistant's ability [161].

A general limitation of the questionnaire-based instruments referenced above is that answering survey questions can never fully capture trust as conceptualized by Mayer et al. They cannot capture critical situations in which a person actually made itself "*vulnerable to the actions of another party*" [98] (here: an AI assistant). As mentioned above, proxies of *behavioral trust* can complement questionnaire-based instruments. However, they also cannot capture actual behavioral trust. A general example of a situation in which behavioral trust occurs would be that parents leave their children with a certain babysitter as an indicator of trusting the babysitter [132]. Another limitation is that human behavior and decision-making are complex and can be affected by factors such as monetary constraints or a lack of alternatives. Therefore, researchers must be careful when using behavior as a proxy for trust, considering that there might be different drivers for the observed behavior.

#### D. Recommendations

Our general recommendation is that SE researchers should keep in mind that focusing on individual factors such as *propensity to trust* or *behavioral metrics* is generally not sufficient to capture a holistic notion of trust. An approach to mitigating threats to validity is to combine different factors and instruments, ideally based on an established trust model. We want to conclude this section with a list of actionable recommendations for researchers who want to conduct studies on trust in an SE context. These recommendations are based on

our thorough description of the foundations of trust (Section II) and our literature review across disciplines (Section III-B):

- 1) Before designing concrete studies on trust in AI assistants, carefully read the literature on **trust beyond SE**. We hope that our article serves as a useful starting point for this.
- 2) When reporting study design and results in a paper, embed them in **established trust models** (see Table I).
- 3) Use **established terminology** when discussing trust, for example, distinguishing *overtrust/undertrust* from *calibrated trust* or distinguishing *initial trust* from *trust formation* (see Section II and Section III-B1).
- 4) Do not equate artifact acceptance with trust. When designing empirical studies to assess developers' trust in AI assistants, clearly outline which **aspects of trust** are being measured using which instruments (e.g., *cognitive trust* versus *affective trust*). Use **validated instruments** whenever possible (see trust operationalizations discussed in this article) and openly discuss their limitations.
- 5) Consider participants' *trust disposition* and *trusting beliefs* as **confounding factors**; consider the *transferable* nature of trust. For example, trust in a particular AI vendor can propagate to a specific AI assistant that this vendor offers.
- 6) Clearly describe the **study context** and reflect on its generalizability. Given how central *risk* and being *vulnerable* are for trusting behavior, a study in a low risk scenario might not generalize to a high risk situation, for example, in the context of safety-critical software.

#### RQ3 (RECOMMENDATIONS FOR SE):

SE researchers should treat trust as a *multidimensional* and *context-dependent* phenomenon. Focusing on individual factors is generally not sufficient to holistically capture trust. Therefore, SE researchers should build on established *trust models* and *measurement instruments*. Section IV-D provides further recommendations.

## V. THREATS TO VALIDITY

As with any empirical study, our methodology is subject to certain limitations that may affect the completeness, accuracy, and generalizability of our findings. In the following, we identify and discuss potential threats to the validity of our methodology, including biases in paper selection, interpretation of trust conceptualizations, and the applicability of findings across disciplines. We classify the validity threats in three categories as proposed by Wohlin et al. [166].

*Internal Validity* concerns whether our study design and execution introduced biases that may affect the correctness of our findings. A potential threat to internal validity is that we may have missed indirect or implicit definitions of trust as well as references to trust models during our literature review. Additionally, by focusing on the highest-ranked articles from our queries, we may have unintentionally excluded relevant papers ranked lower. We acknowledge the selective inclusion based on an opaque ranking algorithm provided by Google Scholar as a limitation. However, Google Scholar is used by academics across disciplines and allowed us to use the

same sampling approach for all disciplines. We complemented the review of SE literature based on Google Scholar search results by conducting a critical review of recently published SE papers, confirming our earlier results.

Future work should extend the literature review by using discipline-specific retrieval approaches. One related discipline that studies AI assistants for software development [115], [116], [129], sometimes referring to human trust [75], [113], is the security research community. We did not explicitly exclude security-related research in our review, only mentions of trust in a more technical context, as clarified in Section I. However, at the same time, we did not explicitly include security research venues. Extending the literature review to consider security research as well is a promising direction for future work.

*External Validity* concerns the generalizability of our findings. Our literature review focused on SE, HCI, and IS; our description of the foundations of trust is rooted in philosophy and psychology. We did not cover trust conceptualizations in other disciplines, such as sociology. However, given that our literature review primarily revealed trust models rooted in psychology, these foundations seem to be the most relevant ones for the considered disciplines. As we only considered English-language publications, our findings may not reflect cultural variations in trust.

*Construct Validity* concerns whether we accurately captured and analyzed the concept of trust as intended. Given the complexity of trust and its multidimensional nature, one potential threat is our own limited understanding of the concept. To mitigate this, we recruited coauthors from philosophy and HCI and collected feedback from another colleague with a background in psychology and IS.

## VI. CONCLUSION

Our review of SE, IS, and HCI literature has shown that SE articles, contrary to IS and HCI, often use the term ‘trust’ informally and that providing an explicit definition of trust or embedding results in established trust models is rare. However, without a common definition, true secondary research on trust is impossible. In this paper, we have discussed the consequences of the lack of common trust models in SE research, and how the SE research community can learn from other disciplines. In Section IV, we present an example of an adaption of an established trust model in an SE context and provide actionable recommendations for SE researchers planning to study trust in AI assistants. Investigating the relationship between trust-related concepts for which established instruments exist and behavioral metrics such as artifact acceptance in the context of AI assistants for software development is a promising direction for future work. We hope that our paper contributes to more mature and interdisciplinarily embedded SE trust research in the future.

## REFERENCES

- [1] E. Adar, D. S. Tan, and J. Teevan, “Benevolent deception in human computer interaction,” in *ACM CHI 2013*, 2013.
- [2] M. K. Ahuja, M. Belaid, P. Bernabé, et al., “Opening the software engineering toolbox for the assessment of trustworthy AI,” in *NeHuAI 2020*, ser. CEUR Workshop Proceedings, vol. 2659, 2020.
- [3] M. A. Akbar, A. A. Khan, S. Mahmood, S. Rafi, and S. Demi, “Trustworthy artificial intelligence: A decision-making taxonomy of potential challenges,” *Softw. Pract. Exp.*, vol. 54, no. 9, 2024.
- [4] A. Alami, R. Pardo, M. L. Cohn, and A. Wasowski, “Pull request governance in open source communities,” *IEEE Trans. Software Eng.*, vol. 48, no. 12, 2022.
- [5] D. K. Allen, D. Colligan, A. Finnie, and T. Kern, “Trust, power and interorganizational information systems: The case of the electronic trading community translease,” *Inf. Syst. J.*, vol. 10, no. 1, 2000.
- [6] S. Amershi, D. S. Weld, M. Vorvoreanu, et al., “Guidelines for human-ai interaction,” in *ACM CHI 2019*, S. A. Brewster, G. Fitzpatrick, A. L. Cox, and V. Kostakos, Eds., 2019, p. 3.
- [7] O. Asan, A. E. Bayrak, and A. Choudhury, “Artificial intelligence and human trust in healthcare: Focus on clinicians,” *J. Med. Internet. Res.*, vol. 22, no. 6, 2020.
- [8] T. A. Bach, A. Khan, H. Hallock, G. Beltrão, and S. C. Sousa, “A systematic literature review of user trust in ai-enabled systems: An HCI perspective,” *Int. J. Hum. Comput. Interact.*, vol. 40, no. 5, 2024.
- [9] M. Bacharach and D. Gambetta, “Trust in signs.,” in (Russell Sage foundation series on trust, vol. 2.), Russell Sage foundation series on trust, vol. 2. New York, NY, US: Russell Sage Foundation, 2001.
- [10] A. Baier, “Trust and antitrust,” *Ethics*, vol. 96, no. 2, 1986.
- [11] A. L. Baker, E. Phillips, D. Ullman, and J. R. Keebler, “Toward an understanding of trust repair in human-robot interaction: Current research and future directions,” *ACM Trans. Interact. Intell. Syst.*, vol. 8, no. 4, 2018.
- [12] S. Baltes and P. Ralph, “Sampling in software engineering research: A critical review and guidelines,” *Empir. Softw. Eng.*, vol. 27, no. 4, 2022.
- [13] S. Baltes, T. Speith, B. Chiteri, S. Mohsenimofidi, S. Chakraborty, and D. Buschek, *Rethinking trust in ai assistants for software development: A critical review (supplementary material)*, Zenodo, Apr. 2025. DOI: 10.5281/zenodo.15031354. [Online]. Available: <https://doi.org/10.5281/zenodo.15031354>.
- [14] S. Banerjee, A. Agarwal, and S. Singla, *Llms will always hallucinate, and we need to live with this*, 2024. arXiv: 2409.05746.
- [15] N. Banovic, Z. Yang, A. Ramesh, and A. Liu, “Being trustworthy is not enough: How untrustworthy artificial intelligence (AI) can deceive the end-users and gain their trust,” *Proc. ACM Hum. Comput. Interact.*, vol. 7, no. CSCW1, 2023.
- [16] S. Becker, W. Hasselbring, A. Paul, et al., “Trustworthy software systems: A discussion of basic concepts and terminology,” *ACM SIGSOFT SEN*, vol. 31, no. 6, 2006.

- [17] I. Benbasat, D. Gefen, and P. A. Pavlou, "Introduction to the special issue on novel perspectives on trust in information systems," *MIS Quarterly*, vol. 34, no. 2, 2010.
- [18] I. Benbasat, D. Gefen, and P. A. Pavlou, "Special issue: Trust in online environments," *J. Manag. Inf. Syst.*, vol. 24, no. 4, 2008.
- [19] I. Benbasat and W. Wang, "Trust in and adoption of online recommendation agents," *J. Assoc. Inf. Syst.*, vol. 6, no. 3, 2005.
- [20] I. Bertram, J. Hong, Y. Huang, W. Weimer, and Z. Sharafi, "Trustworthiness perceptions in code review: An eye-tracking study," in *ACM/IEEE ESEM 2020*, 2020.
- [21] V. D. Bianco, L. Lavazza, S. Morasca, and D. Taibi, "A survey on open source software trustworthiness," *IEEE Softw.*, vol. 28, no. 5, 2011.
- [22] A. W. Boggust, B. Hoover, A. Satyanarayan, and H. Strobel, "Shared interest: Measuring human-ai alignment to identify recurring patterns in model behavior," in *ACM CHI 2022*, 2022.
- [23] L. Bratthall and M. Jørgensen, "Can you trust a single data source exploratory software engineering case study?" *Empir. Softw. Eng.*, vol. 7, no. 1, 2002.
- [24] A. Brown, S. D'Angelo, A. Murillo, C. Jaspan, and C. Green, "Identifying the factors that influence trust in ai code completion," in *AIware 2024*, 2024.
- [25] H. G. Brown, M. S. Poole, and T. L. Rodgers, "Interpersonal traits, complementarity, and trust in virtual collaboration," *J. Manag. Inf. Syst.*, vol. 20, no. 4, 2004.
- [26] B. Buhnova, D. Halasz, D. Iqbal, and H. Bangui, "Survey on trust in software engineering for autonomous dynamic ecosystems," in *ACM SAC 2023*, 2023.
- [27] J. Burgoon, J. Bonito, B. Bengtsson, C. Cederberg, M. Lundeberg, and L. Allspach, "Interactivity in human-computer interaction: A study of credibility, understanding, and influence," *Computers in Human Behavior*, vol. 16, no. 6, 2000.
- [28] J. Cai, R. Wu, H. Hua, G. Sun, Y. Zhang, and X. F. and, "Psychometric properties of the chinese version of human-computer trust (hct) scale in the chinese sample of adults," *Int. J. Hum.-Comput. Int.*, vol. 41, no. 1, 2025.
- [29] R. Calinescu, D. Weyns, S. Gerasimou, M. U. Iftikhar, I. Habli, and T. Kelly, "Engineering trustworthy self-adaptive software with dynamic assurance cases," *IEEE Trans. Software Eng.*, vol. 44, no. 11, 2018.
- [30] M. Chen, "Trust and trust-engineering in artificial intelligence research: Theory and praxis," *Philosophy & Technology*, vol. 34, no. 4, 2021.
- [31] R. Cheng, R. Wang, T. Zimmermann, and D. Ford, "“it would work for me too”: How online communities shape software developers' trust in ai-powered code generation tools," *ACM Trans. Interact. Intell. Syst.*, vol. 14, no. 2, 2024.
- [32] K. Chopra and W. A. Wallace, "Trust in electronic environments," in *HICSS 2003*, 2003.
- [33] C. L. Corritore, B. Kracher, and S. Wiedenbeck, "Online trust: Concepts, evolving themes, a model," *Int. J. Hum. Comput. Stud.*, vol. 58, no. 6, 2003.
- [34] K. E. Culley and P. Madhavan, "A note of caution regarding anthropomorphism in HCI agents," *Comput. Hum. Behav.*, vol. 29, no. 3, 2013.
- [35] S. D. Das, P. K. Bala, and A. N. Mishra, "Towards defining a trustworthy artificial intelligence system development maturity model," *J. Comput. Inf. Syst.*, vol. 64, no. 6, 2024.
- [36] A. R. Dennis, A. Lakhial, and A. Sachdeva, "Ai agents as team members: Effects on satisfaction, conflict, trustworthiness, and willingness to work with," *J. Manag. Inf. Syst.*, vol. 40, no. 2, 2023.
- [37] J. Dorsch and O. Deroy, "Quasi-metacognitive machines: Why we don't need morally trustworthy ai and communicating reliability is enough," *Philosophy & Technology*, vol. 37, no. 2, 2024.
- [38] J. M. Durán and G. Pozzi, "Trust and trustworthiness in ai," *Philosophy & Technology*, vol. 38, no. 1, 2025.
- [39] M. T. Dzindolet, S. A. Peterson, R. A. Pomranky, L. G. Pierce, and H. P. Beck, "The role of trust in automation reliance," *Int. J. Hum. Comput. Stud.*, vol. 58, no. 6, 2003.
- [40] M. Eiband, D. Buschek, A. Kremer, and H. Hussmann, "The impact of placebo explanations on trust in intelligent systems," in *Extended Abstracts of ACM CHI*, 2019.
- [41] A. Ferrario and M. Loi, "How explainability contributes to trust in ai," in *ACM FAccT 2022*, 2022.
- [42] B. J. Fogg and H. Tseng, "The elements of computer credibility," in *CHI 1999*, 1999.
- [43] M. L. Frazier, P. D. Johnson, and S. Fainshmidt, "Development and validation of a propensity to trust scale," *J. Trust Res.*, vol. 3, no. 2, 2013.
- [44] K. Frost-Arnold, "Imposters, tricksters, and trustworthiness as an epistemic virtue," *Hypatia*, vol. 29, no. 4, 2014.
- [45] Ö. Ö. Garibay, B. Winslow, S. Andolina, et al., "Six human-centered ai grand challenges," *Int. J. Hum. Comput. Interact.*, vol. 39, no. 3, 2023.
- [46] D. Gefen, E. Karahanna, and D. W. Straub, "Trust and TAM in online shopping: An integrated model," *MIS Quarterly*, vol. 27, no. 1, 2003.
- [47] D. Gefen and D. Straub, "Managing user trust in b2c e-services," *e-Service*, vol. 2, no. 2, 2003.
- [48] N. Gillespie, "Survey measures of trust in organizational contexts: An overview," *Handbook of Research Methods on Trust*, 2015.
- [49] E. Glikson and A. W. Woolley, "Human trust in artificial intelligence: Review of empirical research," *Acad. Manag. Ann.*, vol. 14, no. 2, 2020.
- [50] M. Grimsley and A. Meehan, "E-government information systems: Evaluation-led design for public value and client trust," *Eur. J. Inf. Syst.*, vol. 16, no. 2, 2007.

- [51] S. Gulati, J. McDonagh, S. Sousa, and D. Lamas, “Trust models and theories in human–computer interaction: A systematic literature review,” *Computers in Human Behavior Reports*, vol. 16, p. 100495, 2024.
- [52] S. Gulati, S. C. Sousa, and D. Lamas, “Design, development and evaluation of a human-computer trust scale,” *Behav. Inf. Technol.*, vol. 38, no. 10, 2019.
- [53] R. Hardin, “Trust and trustworthiness,” *Russell Sage Foundation*, 2002.
- [54] A. Hasija and T. L. Esper, “In artificial intelligence (ai) we trust: A qualitative investigation of ai technology acceptance,” *J. Bus. Logist.*, vol. 43, no. 3, 2022.
- [55] A. E. Hassan, D. Lin, G. K. Rajbahadur, et al., “Rethinking software engineering in the era of foundation models: A curated catalogue of challenges in the development of trustworthy fmware,” in *ACM FSE 2024 Companion Proceedings*, 2024.
- [56] H. He, J. Gray, A. Cangelosi, Q. Meng, T. M. McGinnity, and J. Mehnen, “The challenges and opportunities of human-centered AI for trustworthy robots and autonomous systems,” *IEEE Trans. Cogn. Dev. Syst.*, vol. 14, no. 4, 2022.
- [57] H. Hoehle, S. Huff, and S. Goode, “The role of continuous trust in information systems continuance,” *J. Comput. Inf. Syst.*, vol. 52, no. 4, 2012.
- [58] K. A. Hoff and M. N. Bashir, “Trust in automation: Integrating empirical evidence on factors that influence trust,” *Hum. Factors*, vol. 57, no. 3, 2015.
- [59] F. Hou and S. Jansen, “A systematic literature review on trust in the software ecosystem,” *Empir. Softw. Eng.*, vol. 28, no. 1, p. 8, 2023.
- [60] W. Huo, Z. Zhang, J. Qu, et al., “Speciesism and preference of human-artificial intelligence interaction: A study on medical artificial intelligence,” *Int. J. Hum. Comput. Interact.*, vol. 40, no. 11, 2024.
- [61] W. Huo, G. Zheng, J. Yan, L. Sun, and L. Han, “Interacting with medical artificial intelligence: Integrating self-responsibility attribution, human-computer trust, and personality,” *Comput. Hum. Behav.*, vol. 132, 2022.
- [62] A. Jacovi, A. Marasovic, T. Miller, and Y. Goldberg, “Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in AI,” in *ACM FAccT 2021*, 2021.
- [63] S. Jalali, Ç. Gencel, and D. Smite, “Trust dynamics in global software engineering,” in *ACM/IEEE ESEM 2010*, 2010.
- [64] M. Janssen, P. Brous, E. Estevez, L. S. Barbosa, and T. Janowski, “Data governance: Organizing data for trustworthy artificial intelligence,” *Gov. Inf. Q.*, vol. 37, no. 3, 2020.
- [65] M. Jeon, “The effects of emotions on trust in human-computer interaction: A survey and prospect,” *Int. J. Hum. Comput. Interact.*, vol. 40, no. 22, 2024.
- [66] B. Johnson, C. Bird, D. Ford, N. Forsgren, and T. Zimmermann, “Make your tools sparkle with trust: The PICSE framework for trust in software tools,” in *SEIP@ICSE 2023*, 2023.
- [67] A. N. Joinson, U. Reips, T. Buchanan, and C. B. P. Schofield, “Privacy, trust, and self-disclosure online,” *Hum. Comput. Interact.*, vol. 25, no. 1, 2010.
- [68] K. Jones, “Second-hand moral knowledge,” *The Journal of Philosophy*, vol. 96, no. 2, 1999.
- [69] L. P. R. Jr., G. Bansal, N. P. Melville, and T. Stafford, “Introduction to the special issue on AI fairness, trust, and ethics,” *AIS Trans. Hum. Comput. Interact.*, vol. 12, no. 4, 2020.
- [70] L. Kästner, M. Langer, V. Lazar, A. Schomäcker, T. Speith, and S. Sterz, “On the relation of trust and explainability: Why to engineer for trustworthiness,” in *IEEE REW 2021*, 2021.
- [71] D. Kaur, S. Uslu, K. J. Rittichier, and A. Durresi, “Trustworthy artificial intelligence: A review,” *ACM Comput. Surv.*, vol. 55, no. 2, 2023.
- [72] I. Keivanloo and J. Rilling, “Software trustworthiness 2.0 - A semantic web enabled global source code analysis approach,” *J. Syst. Softw.*, vol. 89, 2014.
- [73] A. Keren, “Trust and belief: A preemptive reasons account,” *Synthese*, vol. 191, no. 12, 2014.
- [74] R. Khojah, M. Mohamad, P. Leitner, and F. G. de Oliveira Neto, “Beyond code generation: An observational study of chatgpt usage in software engineering practice,” *Proc. ACM Softw. Eng.*, vol. 1, no. FSE, 2024.
- [75] J. H. Klemmer, S. A. Horstmann, N. Patnaik, et al., “Using AI assistants in software development: A qualitative study on security practices and concerns,” in *ACM SIGSAC CCS 2024*, B. Luo, X. Liao, J. Xu, E. Kirda, and D. Lie, Eds., 2024, pp. 2726–2740.
- [76] S. Y. X. Komiak and I. Benbasat, “The effects of personalization and familiarity on trust and adoption of recommendation agents,” *MIS Quarterly*, vol. 30, no. 4, 2006.
- [77] P. Korzyński, S. C. e Silva, A. M. Górska, and G. Mazurek, “Trust in ai and top management support in generative-ai adoption,” *J. Comput. Inf. Syst.*, 2024.
- [78] B. Kou, S. Chen, Z. Wang, L. Ma, and T. Zhang, “Do large language models pay similar attention like human programmers when generating code?” *Proc. ACM Softw. Eng.*, vol. 1, no. FSE, 2024.
- [79] M. Kraus, N. Wagner, and W. Minker, “Effects of proactive dialogue strategies on human-computer trust,” in *ACM UMAP 2020*, 2020.
- [80] A. Kudriavtseva, N. A. Hotak, and O. Gadyatskaya, “My code is less secure with gen AI: surveying developers’ perceptions of the impact of code generation tools on security,” in *ACM/SIGAPP SAC 2025*, J. Hong, S. Battiato, C. Esposito, J. W. Park, and A. Przybylek, Eds., 2025, pp. 1637–1646.
- [81] M. Langer, D. Oster, T. Speith, et al., “What do we want from explainable artificial intelligence (xai)?—a stakeholder perspective on xai and a conceptual model guiding interdisciplinary xai research,” *Artificial Intelligence*, vol. 296, 2021.

- [82] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *Human Factors*, vol. 46, no. 1, 2004.
- [83] B. Li, P. Qi, B. Liu, et al., "Trustworthy AI: from principles to practices," *ACM Comput. Surv.*, vol. 55, no. 9, 2023.
- [84] X. Li, T. J. Hess, and J. S. Valacich, "Using attitude and social influence to develop an extended trust model for information systems," *Data Base*, vol. 37, no. 2-3, 2006.
- [85] X. Li, T. J. Hess, and J. S. Valacich, "Why do we trust new technology? A study of initial trust formation with organizational information systems," *J. Strateg. Inf. Syst.*, vol. 17, no. 1, 2008.
- [86] Y. Li, B. Wu, Y. Huang, and S. Luan, "Developing trustworthy artificial intelligence: Insights from research on interpersonal, human-automation, and human-ai trust," *Front. Psychol.*, vol. 15, 2024.
- [87] N. Limam and R. Boutaba, "Assessing software service quality and trustworthiness at selection time," *IEEE Trans. Software Eng.*, vol. 36, no. 4, 2010.
- [88] B. Liu, G. Li, H. Zhang, Y. Jin, Z. Wang, and D. Shao, "The gap between trustworthy AI research and trustworthy software research: A tertiary study," *ACM Comput. Surv.*, vol. 57, no. 3, 2025.
- [89] Y. Liu, T. Le-Cong, R. Widyasari, et al., "Refining chatgpt-generated code: Characterizing and mitigating code quality issues," *ACM Trans. Softw. Eng. Methodol.*, vol. 33, no. 5, 2024.
- [90] S. Lockey, N. M. Gillespie, D. Holm, and I. A. Someh, "A review of trust in artificial intelligence: Challenges, vulnerabilities and future directions," in *HICSS 2021*, 2021.
- [91] R. Lukyanenko, W. Maass, and V. C. Storey, "Trust in artificial intelligence: From a foundational trust framework to emerging research opportunities," *Electron. Mark.*, vol. 32, no. 4, 2022.
- [92] M. Madsen and S. Gregor, "Measuring human-computer trust," in *11th Australasian Conference on Information Systems*, vol. 53, 2000.
- [93] V. Majdinasab, M. J. Bishop, S. Rasheed, A. M. Dakhel, A. Tahir, and F. Khomh, "Assessing the security of github copilot's generated code - A targeted replication study," in *IEEE SANER 2024*, 2024, pp. 435–444.
- [94] A. Manzini, G. Keeling, N. Marchal, K. R. McKee, V. Rieser, and I. Gabriel, "Should users trust advanced ai assistants? justified trust as a function of competence and alignment," in *ACM FAccT 2024*, 2024.
- [95] J. Mao, J. Lee, and C. Deng, "Vendors' perspectives on trust and control in offshore information systems outsourcing," *Inf. Manag.*, vol. 45, no. 7, 2008.
- [96] J. Mattioli, H. Sohier, A. Delaborde, et al., "An overview of key trustworthiness attributes and kpis for trusted ml-based systems engineering," *AI and Ethics*, vol. 4, no. 1, 2024.
- [97] R. C. Mayer and J. H. Davis, "The effect of the performance appraisal system on trust for management: A field quasi-experiment," *J. Appl. Psychol.*, vol. 84, no. 1, 1999.
- [98] R. C. Mayer, J. H. Davis, and F. D. Schoorman, "An integrative model of organizational trust," *Acad. Manag. Rev.*, vol. 20, no. 3, 1995.
- [99] D. J. McAllister, "Affect-and cognition-based trust as foundations for interpersonal cooperation in organizations," *Academy of management journal*, vol. 38, no. 1, pp. 24–59, 1995.
- [100] D. McGeorge, *The ChatGPT Revolution: How to Simplify Your Work and Life Admin with AI*. Wiley, 2023.
- [101] D. H. McKnight, M. Carter, J. B. Thatcher, and P. F. Clay, "Trust in a specific technology: An investigation of its components and measures," *ACM Trans. Manag. Inf. Syst.*, vol. 2, no. 2, 12:1–12:25, 2011. DOI: 10.1145/1985347.1985353. [Online]. Available: <https://doi.org/10.1145/1985347.1985353>.
- [102] D. H. McKnight, V. Choudhury, and C. J. Kacmar, "Developing and validating trust measures for e-commerce: An integrative typology," *Inf. Syst. Res.*, vol. 13, no. 3, 2002.
- [103] D. H. McKnight, V. Choudhury, and C. J. Kacmar, "The impact of initial consumer trust on intentions to transact with a web site: A trust building model," *J. Strateg. Inf. Syst.*, vol. 11, no. 3-4, 2002.
- [104] C. McLeod, "Trust," in *The Stanford Encyclopedia of Philosophy*, E. N. Zalta and U. Nodelman, Eds., Fall 2023, Metaphysics Research Lab, 2023.
- [105] S. Mehrotra, C. Deguchi, O. Vereschak, C. M. Jonker, and M. L. Tielman, "A systematic review on fostering appropriate trust in human-ai interaction: Trends, opportunities and challenges," *ACM J. Responsib. Comput.*, vol. 1, no. 4, Nov. 2024.
- [106] B. M. Muir, "Trust between humans and machines, and the design of decision aids," *Int. J. Man Mach. Stud.*, vol. 27, no. 5-6, 1987.
- [107] A. Murillo, A. Elizondo, S. D'Angelo, et al., "Understanding and designing for trust in ai-powered developer tooling," *IEEE Softw.*, vol. 41, no. 6, 2024.
- [108] B. A. Myers, J. D. Hollan, and I. F. Cruz, "Strategic directions in human-computer interaction," *ACM Comput. Surv.*, vol. 28, no. 4, 1996.
- [109] C. Nass, Y. Moon, B. J. Fogg, B. Reeves, and D. C. Dryer, "Can computer personalities be human personalities?" *Int. J. Hum. Comput. Stud.*, vol. 43, no. 2, 1995.
- [110] P. J. Nickel, "Trust in medical artificial intelligence: A discretionary account," *Ethics Inf. Technol.*, vol. 24, no. 1, p. 7, 2022.
- [111] Y. Noller, R. Shariffdeen, X. Gao, and A. Roychoudhury, "Trust enhancement issues in program repair," in *ICSE 2022*, 2022.
- [112] F. Nothdurft, G. Bertrand, H. Lang, and W. Minker, "Adaptive explanation architecture for maintaining human-computer trust," in *IEEE COMPSAC 2012*, X. Bai, F. Belli, E. Bertino, et al., Eds., 2012.

- [113] S. Oh, K. Lee, S. Park, D. Kim, and H. Kim, “Poisoned ChatGPT finds work for idle hands: Exploring developers’ coding practices with insecure suggestions from poisoned AI models,” in *IEEE SP 2024*, 2024, pp. 1141–1159.
- [114] J.-G. Park and J. Lee, “Knowledge sharing in information systems development projects: Explicating the role of dependence and trust,” *Int. J. Proj. Manag.*, vol. 32, no. 1, 2014.
- [115] H. Pearce, B. Ahmad, B. Tan, B. Dolan-Gavitt, and R. Karri, “Asleep at the keyboard? assessing the security of github copilot’s code contributions,” in *SP 2022*, 2022, pp. 754–768.
- [116] N. Perry, M. Srivastava, D. Kumar, and D. Boneh, “Do users write more insecure code with AI assistants?” In *ACM SIGSAC CCS 2023*, 2023.
- [117] F. T. Piller, M. Srour, and T. J. Marion, “Generative ai, innovation, and trust,” *J. Appl. Behav. Sci.*, vol. 60, no. 4, 2024.
- [118] S. Pink, E. Quilty, J. Grundy, and R. Hoda, “Trust, artificial intelligence and software practitioners: An interdisciplinary agenda,” *AI & SOCIETY*, 2024.
- [119] A. Pinto, S. Sousa, A. Simões, and J. Santos, “A trust scale for human-robot interaction: Translation, adaptation, and validation of a human computer trust scale,” *Hum. Behav. Emerg. Tech.*, vol. 2022, no. 1, 2022.
- [120] F. Poursabzi-Sangdeh, D. G. Goldstein, J. M. Hofman, J. W. Vaughan, and H. M. Wallach, “Manipulating and measuring model interpretability,” in *ACM CHI 2021*, 2021.
- [121] J. K. Rempel, J. G. Holmes, and M. P. Zanna, “Trust in close relationships.,” *Journal of Personality and Social Psychology*, vol. 49, no. 1, 1985.
- [122] J. Riegelsberger, M. A. Sasse, and J. D. McCarthy, “The mechanics of trust: A framework for research and design,” *Int. J. Hum. Comput. Stud.*, vol. 62, no. 3, 2005.
- [123] D. M. Rousseau, S. B. Sitkin, R. S. Burt, and C. Camerer, “Not so different after all: A cross-discipline view of trust,” *Academy of management review*, vol. 23, no. 3, pp. 393–404, 1998.
- [124] M. Ryan, “In ai we trust: Ethics, artificial intelligence, and reliability,” *Sci. Eng. Ethics*, vol. 26, no. 5, 2020.
- [125] S. Sabouri, P. Eibl, X. Zhou, et al., “Trust dynamics in ai-assisted development: Definitions, factors, and implications,” in *ICSE 2025*, 2025.
- [126] A. Safdari, “Toward an empathy-based trust in human-otheroid relations,” *AI & SOCIETY*, 2024.
- [127] K. Saffarizadeh, M. Keil, and L. Maruping, “Relationship between trust in the ai creator and trust in ai systems: The crucial role of ai alignment and steerability,” *J. Manag. Inf. Syst.*, vol. 41, no. 3, 2024.
- [128] A. Salles, K. Evers, and M. F. and, “Anthropomorphism in ai,” *AJOB Neuroscience*, vol. 11, no. 2, 2020.
- [129] G. Sandoval, H. Pearce, T. Nys, R. Karri, S. Garg, and B. Dolan-Gavitt, “Lost at C: A user study on the security implications of large language model code assistants,” in *USENIX Security 2023*, J. A. Calandrino and C. Troncoso, Eds., 2023, pp. 2205–2222.
- [130] N. Scheman, “Trust and trustworthiness,” in *The Routledge Handbook of Trust and Philosophy*, Routledge, 2020.
- [131] P. Schmidt and F. Bießmann, *Quantifying interpretability and trust in machine learning systems*, 2019. arXiv: 1901.08558.
- [132] P. Schmidt, F. Bießmann, and T. Teubner, “Transparency and trust in artificial intelligence systems,” *J. Decis. Syst.*, vol. 29, no. 4, 2020.
- [133] J. R. Schoenherr, R. Abbas, K. Michael, P. Rivas, and T. D. Anderson, “Designing ai using a human-centered approach: Explainability and accuracy toward trustworthiness,” *IEEE Trans. Technol. Soc.*, vol. 4, no. 1, 2023.
- [134] D. D. Scholz, J. Kraus, and L. Miller, “Measuring the propensity to trust in automated technology: Examining similarities to dispositional trust in other humans and validation of the ptt-a scale,” *Int. J. Hum-Comput. Int.*, vol. 41, no. 2, 2025.
- [135] F. D. Schoorman, R. C. Mayer, and J. H. Davis, “Empowerment in veterinary clinics: The role of trust in delegation (1996),” *J. Trust Res.*, vol. 6, no. 1, 2016.
- [136] C. D. Schultz, “A trust framework model for situational contexts,” in *ACM PST 2006*, 2006.
- [137] R. Serafini, A. Yardim, and A. Naiakshina, “Exploring the impact of intervention methods on developers’ security behavior in a manipulated chatgpt study,” in *CHI 2025*, N. Yamashita, V. Evers, K. Yatani, et al., Eds., 2025, 520:1–520:26.
- [138] E. Shein, “The impact of AI on computer science education,” *Commun. ACM*, vol. 67, no. 9, 2024.
- [139] B. Shneiderman, “Bridging the gap between ethics and practice: Guidelines for reliable, safe, and trustworthy human-centered ai systems,” *ACM Trans. Interact. Intell. Syst.*, vol. 10, no. 4, 2020.
- [140] B. Shneiderman, “Human-centered artificial intelligence: Reliable, safe & trustworthy,” *Int. J. Hum. Comput. Interact.*, vol. 36, no. 6, 2020.
- [141] B. Shneiderman, “Human-centered artificial intelligence: Three fresh ideas,” *AIS Transactions on Human-Computer Interaction*, vol. 12, no. 3, pp. 109–124, 2020.
- [142] A. Silva, M. Schrum, E. Hedlund-Botti, N. Gopalan, and M. C. Gombolay, “Explainable artificial intelligence: Evaluating the objective and subjective impacts of xai on human-agent interaction,” *Int. J. Hum. Comput. Interact.*, vol. 39, no. 7, 2023.
- [143] T. W. Simpson, “What is trust?” *Pacific Philosophical Quarterly*, vol. 93, no. 4, 2012.
- [144] M. Söllner, “What we really know about antecedents of trust: A critical review of the empirical information systems literature on trust,” *Psychology of Trust: New Research*, 2013.
- [145] M. Söllner, A. Hoffmann, and J. M. Leimeister, “Why different trust relationships matter for information systems users,” *Eur. J. Inf. Syst.*, vol. 25, no. 3, 2016.

- [146] S. C. Sousa, D. Lamas, J. Cravino, and P. Martins, "Human-centered trustworthy framework: A human-computer interaction perspective," *IEEE Computer*, vol. 57, no. 3, 2024.
- [147] G. Sridhara, R. H. G., and S. Mazumdar, *Chatgpt: A study on its utility for ubiquitous software engineering tasks*, 2023. arXiv: 2305.16837.
- [148] K. J. Stewart, "How hypertext links influence consumer perceptions to build and degrade trust online," *J. Manag. Inf. Syst.*, vol. 23, no. 1, 2006.
- [149] H. Suen and K. Hung, "Building trust in automatic video interviews using various AI interfaces: Tangibility, immediacy, and transparency," *Comput. Hum. Behav.*, vol. 143, 2023.
- [150] C. Tam, A. Loureiro, and T. Oliveira, "The individual performance outcome behind e-commerce," *Internet Research*, vol. 30, no. 2, 2020.
- [151] C. Tansley and S. Newell, "Project social capital, leadership and trust," *Journal of Managerial Psychology*, vol. 22, no. 4, 2007.
- [152] J. B. Thatcher, D. H. McKnight, E. W. Baker, R. E. Arsal, and N. H. Roberts, "The role of trust in post-adoption IT exploration: An empirical examination of knowledge management systems," *IEEE Trans. Engineering Management*, vol. 58, no. 1, 2011.
- [153] S. Thiebes, S. Lins, and A. Sunyaev, "Trustworthy artificial intelligence," *Electron. Mark.*, vol. 31, no. 2, 2021.
- [154] K. Thompson, "Reflections on trusting trust," *Commun. ACM*, vol. 27, no. 8, 1984.
- [155] C. Tomitzka, M. Schaschek, L. Straub, and A. Winkelmann, "What is the minimum to trust ai? - A requirement analysis for (generative) ai-based texts," in *Wirtschaftsinformatik (WI)*, 2023, p. 35.
- [156] E. H. Trainer, B. Al-Ani, and D. F. Redmiles, "Impact of collaborative traces on trustworthiness," in *CHASE 2011*, 2011, pp. 40–47.
- [157] I. Troshani, S. R. Hill, C. Sherman, and D. Arthur, "Do we trust in ai? role of anthropomorphism and intelligence," *J. Comput. Inf. Syst.*, vol. 61, no. 5, 2021.
- [158] J. Viehoff, "Making trust safe for ai? non-agential trust as a conceptual engineering problem," *Philosophy & Technology*, vol. 36, no. 4, 2023.
- [159] M. Virvou, G. A. Tsihrintzis, and E. Tsichrintzi, "VIRTSI: A novel trust dynamics model enhancing artificial intelligence collaboration with human users - insights from a chatgpt evaluation study," *Inf. Sci.*, vol. 675, p. 120759, 2024. DOI: 10.1016/J.INS.2024.120759. [Online]. Available: <https://doi.org/10.1016/j.ins.2024.120759>.
- [160] N. Wang, X. Shen, and Y. Sun, "Transition of electronic word-of-mouth services from web to mobile context: A trust transfer perspective," *Decis. Support Syst.*, vol. 54, no. 3, 2013.
- [161] R. Wang, R. Cheng, D. Ford, and T. Zimmermann, "Investigating and designing for trust in ai-powered code generation tools," in *ACM FAccT 2024*, 2024.
- [162] Y. Wang and D. F. Redmiles, "Cheap talk, cooperation, and trust in global software engineering - an evolutionary game theory model with empirical support," *Empir. Softw. Eng.*, vol. 21, no. 6, 2016.
- [163] D. G. Widder, L. Dabbish, J. D. Herbsleb, A. Holloway, and S. Davidoff, "Trust in collaborative automation in high stakes software engineering work: A case study at NASA," in *ACM CHI 2021*, 2021.
- [164] A. F. Winfield and M. Jiroka, "Ethical governance is essential to building trust in robotics and artificial intelligence systems," *Philos. Trans. R. Soc. A*, vol. 376, no. 2133, 2018.
- [165] E. Winter, D. Bowes, S. Counsell, et al., "How do developers really feel about bug fixing? directions for automatic program repair," *IEEE Trans. Software Eng.*, vol. 49, no. 4, 2023.
- [166] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén, *Experimentation in Software Engineering*. Springer, 2012.
- [167] W. Xu, M. J. Dainoff, L. Ge, and Z. Gao, "Transitioning to human interaction with AI systems: New challenges and opportunities for HCI professionals to enable human-centered AI," *Int. J. Hum. Comput. Interact.*, vol. 39, no. 3, 2023.
- [168] T. Yamagishi and M. Yamagishi, "Trust and commitment in the united states and japan.," *Motiv. Emot.*, vol. 18, no. 2, 1994.
- [169] L. Yan, S. Greiff, Z. Teuber, and D. Gašević, "Promises and challenges of generative artificial intelligence for human learning," *Nature Human Behaviour*, vol. 8, no. 10, 2024.
- [170] R. Yang and S. Wibowo, "User trust in artificial intelligence: A comprehensive conceptual framework," *Electron. Mark.*, vol. 32, no. 4, 2022.
- [171] I. Yaniv, "Receiving other people's advice: Influence and benefit," *Organizational Behavior and Human Decision Processes*, vol. 93, no. 1, 2004.
- [172] R. Yokoi, Y. Eguchi, T. Fujita, and K. Nakayachi, "Artificial intelligence is trusted less than a doctor in medical treatment decisions: Influence of perceived care and value similarity," *Int. J. Hum. Comput. Interact.*, vol. 37, no. 10, 2021.
- [173] G. Zanotti, M. Petrolo, D. Chiffi, and V. Schiaffonati, "Keep trusting! a plea for the notion of trustworthy ai," *AI & SOCIETY*, 2023.
- [174] M. Zhang, X. Zhao, and M. Lyles, "Effects of absorptive capacity, trust and information systems on product innovation," *Int. J. Oper. Prod. Manag.*, vol. 38, no. 2, 2018.
- [175] P. Zhang, S. Ding, and Q. Zhao, "Exploiting blockchain to make AI trustworthy: A software development lifecycle view," *ACM Comput. Surv.*, vol. 56, no. 7, 2024.
- [176] A. Ziegler, E. Kalliamvakou, X. A. Li, et al., "Measuring github copilot's impact on productivity," *Commun. ACM*, vol. 67, no. 3, 2024.