

Epistemology Gives a Future to Complementarity in Human-AI Interactions

Andrea Ferrario^{1,2,3*}, Alessandro Facchini^{2,4}, Juan M. Durán⁵

¹Institute of Biomedical Ethics and History of Medicine, University of Zürich, Zürich, Switzerland

²SUPSI, Dalle Molle Institute for Artificial Intelligence (IDSIA), Lugano, Switzerland

³ETH Zürich, Zürich, Switzerland

⁴Management in Networked and Digital Societies (MINDS) Department, Kozminski University, Warsaw, Poland

⁵TU Delft, Delft, The Netherlands

aferrario@ethz.ch

Abstract

Human-AI complementarity is the claim that a human supported by an AI system can outperform either alone in a decision-making process. Since its introduction in the human-AI interaction literature, it has gained traction by generalizing the reliance paradigm and by offering a more practical alternative to the contested construct of trust in AI. Yet complementarity faces key theoretical challenges: it lacks precise theoretical anchoring, it is formalized only as a post hoc indicator of relative predictive accuracy, it remains silent about other desiderata of human-AI interactions, and it abstracts away from the magnitude-cost profile of its performance gain. As a result, complementarity is difficult to obtain in empirical settings. In this work, we leverage epistemology to address these challenges by reframing complementarity within the discourse on *justificatory AI*. Drawing on computational reliabilism, we argue that historical instances of complementarity function as evidence that a given human-AI interaction is a reliable epistemic process for a given predictive task. Together with other reliability indicators assessing the alignment of the human-AI team with the epistemic standards and socio-technical practices, complementarity contributes to the degree of reliability of human-AI teams when generating predictions. This theoretical repositioning supports the practical reasoning of those affected by these outputs—patients, managers, regulators, and others. Our approach suggests that the role and value of complementarity lie not in providing a stand-alone measure of relative predictive accuracy, but in helping calibrate decision-making to the reliability of AI-supported processes. We conclude by translating this epistemological repositioning into design- and governance-oriented recommendations, including a minimal reporting checklist for justificatory human-AI interactions and measures of efficient complementarity.

Extended version — <https://arxiv.org/abs/2601.09871>

1 Introduction

AI-assisted decision-making is increasingly used in high-stakes domains such as healthcare, education, and public administration. In such settings, a central question is whether

interaction between a human and an AI system actually improves decision quality relative to either acting alone. Research on human-AI interaction captures this expectation under the label *complementarity*: in prediction tasks, a human assisted by an AI system, often referred to as a *human-AI team*, should outperform both the human and the AI alone (Bansal et al. 2021b,a). Since Bansal et al. (2021b), complementarity has attracted considerable attention in the human-AI interaction community (Bansal et al. 2021a,b; Hemmer et al. 2021, 2025; Vaccaro, Almaatouq, and Malone 2024; Gonzalez et al. 2026). A growing body of work shows that, while humans and AI models often make different kinds of errors (Geirhos et al. 2021), human-AI teams can outperform humans alone (Alufaisan et al. 2021; Inkpen et al. 2023; Vaccaro, Almaatouq, and Malone 2024), and, in some settings, they can even outperform AI systems (Dvijotham et al. 2023; Ma et al. 2023). Here, complementarity serves as a design standard that motivates interaction protocols extending the reliance paradigm, in which the output of the human-AI interaction is constrained to equal either the human’s prediction (self-reliance) or the AI’s prediction (AI reliance) (Schemmer et al. 2023; Zhang, Liao, and Bellamy 2020; Bansal et al. 2021b), by admitting interaction outputs that differ from the human’s and the AI’s predictions as the result, for instance, of repeated exchange of information.

Despite its intuitive appeal, complementarity faces four key theoretical challenges. First, complementarity is defined in terms of a performance measure called *complementarity team performance* (CTP) and its theoretical anchoring remains limited (Hemmer et al. 2021, 2025; Donahue, Chouldechova, and Kenthapadi 2022). Second, like other performance metrics in supervised learning, CTP is an *ex post*, ground-truth-dependent measure: it is not available at decision time and its stability under AI system updates and human learning is poorly understood (Hemmer et al. 2025). Third, complementarity is defined solely over predictive and *relative* accuracy; in many contexts, however, accuracy is only one among several decision-relevant desiderata, such as fairness, robustness, explainability, and resource constraints, among others (Vaccaro, Almaatouq, and Malone 2024). Finally, the literature on complementarity currently abstracts from discussing the magnitude of the performance gain achieved by the human-AI team, and the cost of interaction required to obtain it. Taken together, these ob-

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

servations suggest that, despite the growing body of work in human-AI interaction, complementarity is not suited to function as a stand-alone, *relative* accuracy-focused gold standard for human-AI collaboration. Not only do we lack a clear theoretical stance on complementarity, but empirically we still struggle to achieve it (Vaccaro, Almaatouq, and Malone 2024). So, **what future lies ahead for complementarity in human-AI interactions?**

In this paper, we argue that by reconsidering complementarity through an epistemological lens, we can still assign it a meaningful role in human-AI interaction research despite these limitations. Our central claim is that complementarity should not be treated as a stand-alone gold standard for human-AI collaboration, but as historically grounded evidence that a prediction-task human-AI interaction may constitute a reliable epistemic process for a given task. More specifically, we show that complementarity can contribute to the epistemic justification of outputs produced by human-AI teams. To do so, we reinterpret it as one reliability indicator within *computational reliabilism* (Durán and Formanek 2018; Durán 2025; Durán 2026), a process-reliabilist framework that connects the historical performance and surrounding practices of computational processes to the justification of beliefs formed on their basis. On this view, complementarity contributes to what we call *justificatory AI* (Ferrario 2024; Alvarado 2023a,b) by helping answer a practical question for affected third parties: *when is it reasonable for a patient, student, manager, or regulator to accept the output of a human-AI team as epistemically adequate for a given task?* Our contributions are as follows:

- **We synthesize the main theoretical challenges to complementarity** identified in the human-AI interaction literature, and discuss how they constrain its empirical effectiveness.
- **We reconceptualize complementarity within computational reliabilism**, treating prediction-oriented human-AI interactions as computational processes whose reliability comes in degrees and is supported by heterogeneous markers, with historical instances of complementarity serving as one important indicator. This shows how complementarity can contribute to the epistemic justification of outputs produced by human-AI teams. Furthermore, CR allows clarifying what should be measured and reported beyond relative accuracy, e.g., stability under updates, task validity, and governance and competence scaffolding, to assess epistemic adequacy of human-AI teams.
- **We argue that researchers should target *efficient complementarity***: if predictive gains require sustained monitoring, prolonged deliberation, or extensive retraining, the design may merely shift epistemic burden onto users and erode the practical value of the improvement. To operationalize this perspective, we provide a minimal reporting checklist for *justificatory human-AI interactions* and measures of efficient complementarity.

2 The Basics: Human-AI Interactions, Reliance, and Complementarity

We begin with a minimal formal setup for human-AI interactions, reliance, and complementarity.

Human-AI Interactions

We call *human-AI interactions* (HAIs) goal-oriented processes in which a human and an AI system contribute informational inputs that are combined into a single output, e.g., a prediction or AI-generated content, possibly after a update.¹ In prediction settings, such as a clinician consulting a medical AI to classify skin lesions, the informational inputs are typically the agents' per-instance predictions: the AI is used as a recommendation-driven decision-support system (Miller 2023).² In this work, interactions whose goal is to approximate ground truth for labeled instances play a key role.

Prediction-Task-Oriented HAIs and the Self-Reliance vs. AI Reliance Paradigm

A prediction task τ is the datum of an input space $\mathcal{X}$, an output/label space $\mathcal{Y}$, and a pointwise loss ℓ used to evaluate predictions against ground truth. Given τ , a labeled dataset is a finite sample $D := \{(x_i, y_i)\}_{i=1}^n$ with $x_i \in \mathcal{X}$ and $y_i \in \mathcal{Y}$. Typically, datasets are used to estimate performance for τ of human-AI interactions over time. More precisely, a *prediction-task human-AI interaction* (PT-HAI)—in symbols: $\text{PT-HAI}^\tau(\Pi)$ —for τ is an HAI whose output $\hat{y}_i^{\text{HAI}}$ is the approximation (or prediction) of the ground truth y_i of any instance (x_i, y_i) in a dataset D . A $\text{PT-HAI}^\tau(\Pi)$ is typically instantiated over time for the same task τ , under a fixed interaction protocol Π , i.e., the specification of how the human consults and integrates the AI, but on (possibly different) datasets $D_1, \dots, D_k$, where each D_j is a finite sample from $\mathcal{X} \times \mathcal{Y}$. When unambiguous, we omit the explicit reference to τ and Π .³

If the AI is used as a recommendation-driven decision support, then the inputs of the PT-HAI are the human's and the AI's predictions $\hat{y}_i^H$ and $\hat{y}_i^{AI}$ of y_i for a given task τ . This is the case for a prominent class of PT-HAIs called *self-reliance vs. AI reliance*—in short: *reliance*—where no

¹This definition is deliberately minimal: it characterizes HAIs at the level of information exchange and output formation, without committing to any particular interaction protocol, degree of agency, or normative relationship between the agents.

²Miller uses the term *recommendation* for what we call a *prediction* (Miller 2023). More generally, however, these AI inputs may include any informational content supplied to the interaction protocol; for instance the AI system may provide information to help the human refine their prediction *without* sharing a prediction, i.e., as a hypothesis-driven decision support (Miller 2023).

³We treat PT-HAIs as processes that can be instantiated many times over time, on different cases and datasets, and potentially by different qualified humans, while keeping the task, protocol, and the AI fixed. We separate the process from its tokens, i.e., the interaction instances. This matters later, because our reliability claims introduced in Section 3 concern the PT-HAI, not any single token.

interaction between the predictions of the human and AI occurs, and the output given by the selector of the two single-agent predictions: if $\hat{y}_i^{HAI} = \hat{y}_i^H$ we have *self-reliance*; if $\hat{y}_i^{HAI} = \hat{y}_i^{AI}$ we have *AI reliance* (Schemmer et al. 2023; Zhang, Liao, and Bellamy 2020; Ferrario 2025) instead. We use “reliance” in the narrow selector sense, to keep it contrastive with interactions where information exchange between the human and the AI allows outputs outside $\{\hat{y}_i^H, \hat{y}_i^{AI}\}$ —see Section 2 below. If the human is initially wrong, the AI is right, and the human follows the AI or the human is initially right and ignores incorrect AI advice, then *appropriate reliance* emerges. This is the target of reliance relations and it is actively discussed in the human-AI interaction and philosophy of AI literature (Schemmer et al. 2023; Zhang, Liao, and Bellamy 2020; Ferrario 2025).

Complementarity

Bansal et al. (2021a) observe that, in many contexts where humans use AI systems in decision-making, we are interested in the performance of the human-AI team rather than that of either component alone. The reason is that, in many applications, AI systems provide *advice* that humans incorporate into decision-making. (They therefore exclude cases of full automation.) In these scenarios, the AI system assists humans by sharing its outputs, while human agents retain accountability for the final decisions (Bansal et al. 2021b): these decisions are taken by the human assisted by the AI or, hereafter, the human-AI team.⁴ Paradigmatic cases happen in high-stakes domains, such as medicine and the judicial system, where experts, e.g., doctors and forensic experts, use AI systems to assist their decision-making, e.g., diagnoses, as well as facial and voice detection assessments. Following (Hemmer et al. 2025; Donahue, Chouldechova, and Kenthapadi 2022), we introduce complementarity as follows:

Definition 1 (Complementarity) Let $PT\text{-}HAI^\tau(\Pi)$ be a prediction-task human-AI interaction with protocol Π , relative to a task τ and dataset $D := \{(x_i, y_i)\}_{i=1}^n$, with human and AI predictions $\hat{y}_i^H$ and $\hat{y}_i^{AI}$ for each y_i , and a team output $\hat{y}_i^{HAI}$. Let L_H , L_{AI} , and L_{HAI} be empirical losses⁵ measuring the error incurred by the human, the AI system, and the human-AI team, respectively, in approximating the ground truths y_i in D .

Define the Complementarity Team Performance for task τ relative to D by:

$$CTP_\tau(D) := \begin{cases} 1, & \text{if } L_{HAI}(D) < \min\{L_H(D), L_{AI}(D)\}, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

We say that $PT\text{-}HAI^\tau(\Pi)$ achieves complementarity on D whenever $CTP_\tau(D) = 1$.

⁴In this work, human-AI team refers to the human agent responsible for producing a prediction as a result of their interaction with an AI system. Our terminology is aligned with what Hemmer et al. (2025) define as “human-AI collaboration.”

⁵These aggregate the loss ℓ computed at each prediction. They depend on the prediction task.

Complementarity is a binary property of a PT-HAI: a human-AI interaction achieves complementarity,⁶ i.e., $CTP_\tau(D) = 1$, as per eq. (1), when the human and the AI use their capabilities in such a way that their respective errors are mitigated or corrected (Bansal et al. 2021a).⁷ Importantly, as remarked by Hemmer et al. (2025), the human-AI interaction output $\hat{y}_i^{HAI}$ may be different from $\hat{y}_i^H$ and $\hat{y}_i^{AI}$. Two remarks. First, Definition 1 can be generalized in terms of objectives, such that utilities, payoffs, or policies, which may differ from empirical losses (Rastogi et al. 2023; Donahue, Chouldechova, and Kenthapadi 2022), but this does not affect our core argument. Similarly, the binary measure $CTP_\tau(\cdot)$ can be relaxed and replaced by continuous ones; also this choice does not affect our core argument. Furthermore, complementarity extends reliance by requiring active interaction that entails the exchange of information between the human and the AI system, resulting in a human-AI team’s prediction that does not need to match either the human’s or the AI’s. For instance, in regression tasks, the human-AI team may average or otherwise combine the two predictions as the result of this informational exchange. In multiclass classification, the human-AI team may select a class that neither the human nor the AI initially (and independently) proposed. As a result, this collaboration augments the space of outputs of the interaction.⁸

Complementarity remains attractive in human-AI research because it shifts attention from optimizing AI systems in isolation to understanding when collaboration between humans and AI systems yields genuine epistemic gain in practice (Miller 2023; Hemmer et al. 2025). Furthermore, $CTP_\tau(\cdot)$ provides an easy-to-understand quantitative measure of human-AI collaboration, helping shift attention away from contested constructs such as trust in AI and calibrating trust in trustworthiness, which have been widely criticized in both philosophy and human-AI interaction (Ryan 2020; Durán and Pozzi 2025; Benk et al. 2025; Ferrario 2025). That said, as currently formulated, complementarity faces several key theoretical challenges, which we discuss below.

Complementarity: Key Theoretical Challenges

We list four theoretical challenges affecting complementarity that will motivate our approach in Section 3.

I. What theoretical anchoring for complementarity?

Complementarity still lacks a clear theoretical positioning. Hemmer et al. (2025) suggest that empirical findings showing humans and AI systems making different types of errors

⁶Note that Vaccaro, Almaatouq, and Malone (2024) call complementarity “human-AI synergy,” that is “where the human-AI group performs better than both the human alone and the AI alone” (Vaccaro, Almaatouq, and Malone 2024, pag. 2294). We refer to (Hemmer et al. 2025), for a review of concepts similar to complementarity that commonly appear in the literature.

⁷In this sense, $CTP_\tau(D)$ is a descriptive quantity: it records whether a PT-HAI achieved complementarity on D . By contrast, treating complementarity as an interaction target is normative.

⁸PT-HAIs where the AI is used purely as an information source to help the human refine their initial prediction can still be analyzed through the lens of complementarity, provided the AI is in principle capable of computing predictions.

can be framed within an Input-Processing-Output model of decision-making (D’Arcy et al. 2014), which rests on two key ideas: information and capability asymmetry. However, the model is highly general, used across domains and technological artifacts, and does not distinguish between different kinds of human-AI interactions, such as reliance versus complementarity. It highlights differences in information and its processing, but not the concrete mechanisms by which inputs are compared, reconciled, and integrated into a human-AI team decision, see Fig. 1 in (Hemmer et al. 2025). A more recent and notable attempt is the framework by Gonzalez et al. (2026), which offers a socio-technical vocabulary for thinking about human-AI teaming across dimensions such as reasoning, memory, attention, and governance. Yet this framework also remains high-level and programmatic: it is neither formal nor prescriptive, and its empirical basis is still partly limited by the early-stage and heterogeneous nature of current evidence.

II. Limited applicability of complementarity at decision time. CTP and other complementarity-related metrics, see Hemmer et al. (2025), depend on ground truth labels and summarize historical performance over a set of cases, epistemically analogous to accuracy or mean squared error in supervised learning.⁹ They therefore do not tell a decision-maker at a given time, in the absence of ground truth, whether to rely on the AI, on themselves, or on an alternative information source (and if ground truth were available at decision time, AI predictions would be redundant). In other words, we have no direct way to measure the emergence (or absence) of complementarity as decisions are being made: *complementarity is an ex post construct to evaluate human-AI interactions*.¹⁰ Finally, current formulations of complementarity are silent on how robust CTP must be across temporal changes, e.g., data distribution shifts and human epistemic process updates, and under what conditions one may expect complementarity to persist (Hemmer et al. 2025; Vaccaro, Almaatouq, and Malone 2024).

III. More than relative predictive accuracy is at stake in human-AI interactions. Complementarity is a *relative* accuracy criterion: a human-AI team can satisfy CTP even if its absolute performance is very low, as long as it is better than each component alone: even very unskilled humans and highly inaccurate AIs can still achieve complementarity—

⁹Practically, however, complementarity differs from standard supervised-learning evaluation as in many applied pipelines collecting labels to estimate accuracy-type metrics is easy to moderately difficult: labels can be obtained through routine outcomes, or periodic audits. By contrast, estimating complementarity additionally requires instrumenting and recording the human-AI interaction itself: one must capture not only ground truth y_i , but also the human’s and AI’s predictions (possibly multiple, updated over the interaction), the interaction protocol, and the team output $\hat{y}_i^{HAI}$ in a way that supports faithful reconstruction of the process. That said, a theory of complementarity outside supervised learning is an open avenue of research.

¹⁰Note that Alur, Raghavan, and Shah’s framework (2024) can guide design ex ante, but their “superiority relations” are still learned against post-hoc labeled data (see their Algorithm 1).

simply by being less bad together than each is on their own. However, many real-world decision settings impose constraints that are not captured by relative loss comparisons alone: robustness to distribution shift, validity of the target construct, transparency and fairness requirements, resource constraints, legal compliance, and institutional accountability, among others. Some of these desiderata are not purely epistemic, but in high-stakes contexts they function as practical constraints on what counts as an epistemically acceptable decision procedure. A human-AI team can display excellent aggregate predictive performance while simultaneously amplifying disparate error rates across subgroups (Donahue, Chouldechova, and Kenthapadi 2022). Because CTP is computed at the dataset level, subgroup-level complementarity failures can be masked by aggregate gains, a structural limitation that any stand-alone use of CTP inherits. More broadly, a team can satisfy CTP while increasing cognitive burden on the human agent through excessive monitoring of the AI system and its outputs, or while undermining accountability and patient trust. Thus, *complementarity is, at best, a conditional desideratum of human-AI interactions*: it is epistemically and practically attractive only when it is achieved subject to constraints on other epistemic or non-epistemic (e.g., ethical) values.

IV. Complementarity ignores the magnitude-cost profile of epistemic gain. Beyond being a relative accuracy criterion, complementarity abstracts away from two relevant dimensions: (1) the magnitude of the performance gain achieved by the human-AI team, and (2) the cost of interaction required to obtain it. In practice, the epistemic and practical significance of complementarity gains varies widely. In some settings, achieving CTP yields only negligible improvements over reliance; in others, it results in substantial error reduction. However, achieving complementarity may require non-trivial costs, including increased cognitive effort, longer decision times, additional training and monitoring, or institutional commitments. The current literature does not account for this gain-cost trade-off: as a result, *complementarity treats interactions with very different epistemic and practical profiles as equivalent, despite their divergent implications for real-world deployment*.

These theoretical challenges affect complementarity in empirical settings. In fact, research indicates that, although there is some partial support, complementarity does not emerge in practice systematically. Vaccaro et al.’s (Vaccaro, Almaatouq, and Malone 2024) meta-analysis of 100+ studies shows that, while human-AI teams frequently outperform humans alone, on average, human-AI teams do not perform better than either humans or AI alone. It also shows that task type moderates complementarity: for prediction tasks the pooled effect size for complementarity is significantly negative, whereas it is positive for creation tasks, such as AI-assisted generation of artistic images or text. Design interventions such as explanations or confidence displays do not reliably induce complementarity and hindrances include miscalibrated self-assessment, coordination frictions, and forms of error mirroring (Turel and Kalhan 2023; Vac-

caro, Almaatouq, and Malone 2024; Vaccaro and Waldo 2019; Zhang, Liao, and Bellamy 2020; Cabitza, Campagner, and Sconfienza 2021).

Given these theoretical limitations and empirical challenges, complementarity should be theoretically repositioned in AI-assisted decision-making. Our claim is not that complementarity lacks value. Rather, when treated as an *ex post*, ground-truth-dependent, relative-accuracy criterion that is difficult to obtain empirically, it is poorly suited to function as a stand-alone gold standard for human-AI interaction. What is needed is a theoretical repositioning that preserves what is informative about complementarity while embedding it in a broader account of epistemic adequacy. We defend this perspective by positioning complementarity within a framework for justifying the predictions of humans assisted by AI systems. Our approach is described next.

3 Our Approach: Complementarity and Computational Reliabilism

Here, our aim is to show that complementarity can play a meaningful role in the epistemic justification of outputs produced by human-AI teams, reinterpreting it as one reliability indicator within the framework of computational reliabilism (CR) (Durán and Formanek 2018; Durán et al. 2024; Durán 2025). To do so, we need to start with a primer on CR.

A Primer on Computational Reliabilism

Goldman’s *process reliabilism* holds that a belief is justified when it is produced by a reliable belief-forming process, that is, one that tends to yield true (or otherwise epistemically adequate) outputs in the relevant circumstances (Goldman 1986). Computational reliabilism applies and extends this idea to cases in which beliefs are formed through direct or indirect interaction with *computational processes*, including those realized through computer simulations, algorithms, and AI systems (Durán 2025; Durán and Formanek 2018; Durán 2026). Thus, on CR, the central unit of epistemic evaluation is the computational process that performs a task τ , e.g., generating a prediction, and the justificatory question is whether that process is *reliable* in a given context—we explain the notion of reliability below. Here, a computational process is an input-output procedure that transforms information according to some method, e.g., statistical or rule-based inference, or simulation. The computational process is embedded in a socio-techno-scientific context that shapes the design of the entity, e.g., an AI system, realizing the computation at different instantiations (or tokens) of the process over time (Durán 2025; Durán 2026). CR reads:

Definition 2 (Computational reliabilism) *Let X be an agent and A a computational process used for task τ . X is justified in accepting A ’s output as epistemically adequate for τ iff A is reliable for τ .*

In Def. 2, reliability is dispositional: a computational process is *reliable* when, across the kinds of cases in which it is used for τ , it *usually* yields epistemically adequate results (for instance, medically valid predictions in a medical context by an AI system) (Durán 2025; Durán 2026).

Here, *epistemically adequate* is a domain-relative success notion, see (Durán 2025; Durán 2026). In prediction tasks, adequacy is primarily predictive accuracy under the domain-relevant performance measure, but in high-stakes it includes also disciplined treatment of uncertainty and adherence to constitutive domain constraints, e.g., scope of use and acceptable error trade-offs. The reliability of the computational process provides X with defeasible reasons to treat outputs as admissible inputs to action and further deliberation. CR is compatible with both frequentist and propensity-style interpretations of *usually*; we will not (need to) take a stand on this debate here.¹¹

Crucially, under CR, the reliability of the computational process is given by families of properties, capabilities, and practices that provide defeasible evidence that the process is reliable in the context where it is performed. These properties, capabilities, and practices are collected into *reliability indicators*, understood as marks of methodological, cognitive, social, and epistemic competence. Thus, endorsing CR, the identification of reliability indicators is key to justifying X ’s belief that the output of a computational process is epistemically adequate. More precisely, CR distinguishes between categories of indicators (*type-RIs*), whose content depends on what entity realizes the computational process:

1. **the technical performance and operational behavior of the realizing entity** (*type-RI₁*);
2. **the proper operationalization of scientific concepts, models, and domain knowledge** (*type-RI₂*);
3. **the social construction of reliability** (*type-RI₃*).

As a brief example, consider a computational process realized by an AI system (e.g., fracture detection from medical images or sepsis prediction from physiological streams). *type-RI₁* captures the system’s engineering and operational performance, namely, how it is built and maintained, and how well it behaves in use, e.g., accuracy, error modes, robustness to shifts (Durán 2025; Durán 2026). *type-RI₂* captures epistemic and scientific fit: how domain concepts, models, and values from bioinformatics and medicine are translated into computational form, and how well that operationalization supports inquiry in everyday clinical practice (Durán et al. 2024; Durán 2025; Durán 2026). Finally, *type-RI₃* is concerned with how different communities, e.g., bioinformaticians, clinicians, nurses, and the public come to endorse or dismiss AI systems and their outputs, how well these systems realize their intended values and purposes, and how their overall epistemic standing is judged. *type-RI₃* indicators concern institutionalized practices that make reliability contestable, auditable, and corrigible over time, e.g., public and expert debate, stress-testing, and other forms of critical, collective scrutiny (Durán et al. 2024; Durán 2025; Durán 2026).

Two remarks are due. First, while CR provides justification for predictions of computational processes, these predictions can occasionally be imprecise, inadequate for a specific purpose, or simply wrong. Over time, the accrual

¹¹For classic discussions, see Alston (1995).

of these errors would affect the propensity of the process to compute usually epistemically adequate predictions, ultimately making it unreliable. Furthermore, although conceptually distinct, the reliability indicators interlock in practice and jointly support the process reliability. In fact, the reliability of a computational process comes in degrees and the weight a particular type of reliability indicator carries for a process's reliability depends on the socio-technical context in which it unfolds. The relative importance of a reliability indicator type is shaped by the epistemic and practical aims in play, as well as by the values of the relevant epistemic community, namely, their habits and norms shaping the management of computational processes. In fact, no single reliability indicator can secure reliability across all processes and contexts (and to the same degree). Even when a given indicator is well-suited to a particular process at one point in time, subsequent developments may undermine its weight: indicators that once seemed compelling can lose their justificatory power as new methods, standards, or evidence emerge (Durán et al. 2024; Durán 2025; Durán 2026).

The Role of Complementarity in CR

Finally, we show how complementarity fits within CR. The key move is to treat PT-HAIs as computational processes, and any human-AI team in a given use episode as a token (instantiation) of that process—Section 2. On this view, the relevant computation is not performed by the AI system alone as in Section 3: it is jointly realized by a human and an AI system acting under a fixed interaction protocol. Two philosophical considerations support this process view. First, Extended Mind Theory and distributed cognition approaches hold that cognitive work can be functionally distributed across people and cognitive artifacts (Hutchins 1995; Hollan, Hutchins, and Kirsh 2000; Clark and Chalmers 1998; Heersmink 2015). Second, following Alvarado (2023a), AI systems are epistemic technologies: they can enhance human inquiry by contributing information and capabilities that extend a human user's cognitive and epistemic capabilities. Taken together, these perspectives motivate treating the PT-HAI as a legitimate unit of epistemic assessment in CR and human-AI teams as its token realizations over time. Then:

Definition 3 (Computational reliabilism for PT-HAIs)

Let X be a human agent, A an AI system, and let $\hat{y}^{HAI}$ denote the output of a prediction-task interaction between them for task τ and relative to a dataset D . A third party Z (e.g., a patient, student, manager, or regulator) is justified in believing that $\hat{y}^{HAI}$ is epistemically adequate for τ iff $\hat{y}^{HAI}$ is produced by a reliable prediction-task human-AI interaction for τ .

Following Section 3, also PT-HAI reliability in Def. 3 comes in degrees and it is granted by reliability indicators.

type-RI₁ and complementarity. type-RI₁ indicators capture properties of the PT-HAI as an epistemic process, rather than of the human or the AI in isolation. **Within type-RI₁, we treat complementarity—as captured by CTP_τ(·)—as a central reliability indicator of the in-**

teraction: it tells us whether, in a given prediction task, the human-AI team performs better than either component alone. Other indicators in type-RI₁ include systematic patterns of error mitigation or correction, e.g., humans appropriately contest or override AI outputs, escalation and review behavior in difficult cases, stability of joint performance, and the extent to which interaction protocols prevent over- or under-reliance. It also includes disaggregated performance evidence, namely the distribution of joint accuracy and error rates across operationally relevant subgroups, since aggregate complementarity can mask subgroup-level failures that affect the reliability profile of the PT-HAI. Epistemically, *historical complementarity* supports a counterfactual claim: that had either component acted alone, the PT-HAI would have been *less reliable*. Complementarity does not merely track performance at a given instantiation of a PT-HAI but evidences a division of epistemic labor whose relevance accrues over time.

That said, *is complementarity necessary for PT-HAI reliability?* In general, the answer is no due to degenerate cases. For instance, suppose the AI system is perfect on task τ for multiple datasets $D_1, \dots, D_m$, i.e., $L_{AI}(D_k) = 0$, $k = 1, \dots, m$. Then no PT-HAI can strictly outperform the AI, so $CTP_\tau(D_k) = 1$ is impossible; yet a trivial interaction in which the human simply defers to the AI can still be—to an extent that depends, in particular, on the context—reliable for τ . Similar considerations hold in the case of a human with perfect performance on certain datasets of a task τ .¹² Set these degenerate cases aside, what complementarity does provide is a particularly strong type-RI₁ signal that the interaction protocol enables genuine *epistemic gain* beyond either component alone.

Accordingly, we treat complementarity as a central indicator of PT-HAI reliability, but neither as a necessary condition nor as a sufficient one. Thus, on our account, complementarity is *evidential rather than constitutive*: it strengthens the case that a prediction-task human-AI interaction is reliable, but it does not by itself define or secure that reliability.

type-RI₂ and epistemic standards. type-RI₂ indicators capture the extent to which the PT-HAI process is aligned with the epistemic and scientific standards that govern sound inference in the relevant domain. type-RI₂ asks whether the interaction is the right kind of epistemic procedure for task τ in context. Concretely, it includes indicators such as: (i) *task and target validity*, i.e., whether τ is well-defined and whether $\hat{y}^{HAI}$ targets a construct that is epistemically appropriate (rather than a poorly motivated proxy); (ii) *evidential fit*, i.e., whether the inputs and features used by the AI system and the human's background information are relevant and adequate for τ ; (iii) *validation and generalization standards*, i.e., whether there is adequate evidence that the process performs reliably under the intended conditions of

¹²One could avoid the degenerate cases by adopting a tolerance-based notion of complementarity, e.g., setting $CTP_\tau(D) = 1$ whenever $L_{HAI}(D) \leq \min\{L_H(D), L_{AI}(D)\}$ (or by requiring an improvement of at least $\varepsilon > 0$). In this work, we follow the strict definition introduced by Hemmer et al. (2025) and related work.

use, e.g., external validation, subgroup analyses where relevant, and stress tests for foreseeable distribution shifts; (iv) *uncertainty and calibration discipline*, i.e., whether uncertainty is represented and communicated in a manner that supports appropriate uptake, e.g., calibrated confidence estimates, abstention policies, or “do-not-use” triggers; (v) normative performance trade-offs as epistemic constraints, i.e., whether the chosen operating point, such as a sensitivity vs. specificity trade-off, reflects domain standards for acceptable error; and (vi) *fairness of the target construct and its operationalization*, namely whether the task definition, label choice, proxy variables, and feature set are adequate given the domain’s normative framework, including whether the chosen fairness criterion (e.g., predictive parity, equalized odds, calibration within groups) is epistemically appropriate for τ and whether fairness-accuracy trade-offs at the AI-alone level are coherently addressed at the PT-HAI level (Donahue, Chouldechova, and Kenthapadi 2022).

type-RI₃ and socio-technical practices. *type-RI₃* indicators capture the practices, institutional arrangements, and governance mechanisms that stabilize (or undermine) the reliability of the PT-HAI over time. Even where a human-AI team can in principle achieve strong *type-RI₁* performance and meet *type-RI₂* standards, reliability of the corresponding PT-HAI may fail in practice unless there are durable procedures for competent use, monitoring, and change management. Accordingly, *type-RI₃* includes indicators such as: (i) *training and competence scaffolding*, e.g., targeted education about failure modes, appropriate override behavior; (ii) *workflow integration*, e.g., clear role assignment for X and A , decision checkpoints, documentation of override; (iii) *monitoring and feedback*, e.g., post-deployment performance tracking, incident reporting, auditing of systematic errors; (iv) *accountability and incentives*, e.g., who is responsible for outcomes, whether institutional incentives encourage critical thinking; (v) *lifecycle and update governance*, e.g., versioning, communication of changes to users, and ensuring that the human’s expectations remain aligned with the system’s actual behavior; and (vi) *fairness audits and disparity-mitigation governance*, namely the institutional practices through which subgroup performance is monitored, disparities are surfaced and contested, remediation is triggered, and accountability for fairness outcomes is assigned. *type-RI₃* indicators make explicit that the reliability of the PT-HAI is a property of an evolving socio-technical practice, providing evidence that the PT-HAI remains a reliable computational process across time.

In summary, complementarity is an indicator of the reliability of PT-HAIs in the framework of computational reliabilism. It provides strong evidence for PT-HAI reliability when complementarity gains are non-trivial in magnitude (i.e., $\text{CTP}_\tau(D) = 1$ is practically meaningful for the domain), stable across time and robust to foreseeable distribution shifts and system updates, and efficient, i.e., not achieved by shifting disproportionate cognitive or organizational burden onto human operators. However, complementarity is not enough to secure high levels of reliability

by itself. In fact, as previously noted, complementarity is a purely relative performance criterion and a human-AI team can satisfy CTP even when the AI system is poorly reliable or the AI is highly reliable but the human contributor is not, e.g., an untrained user, so that the interaction remains fragile or unsafe despite occasional gains.¹³ **High PT-HAI reliability therefore requires complementarity together with additional indicators of epistemic adequacy.** Accordingly, this high reliability must also be supported by established validation practices, appropriate incentive structures, transparent communication and documentation, institutional accountability, and community-endorsed criteria of success.

A Brief Summary of the Two Justificatory Perspectives at Hand

We briefly clarify the two justificatory perspectives discussed so far. We summarize them graphically in Figure 1. In the first justificatory AI perspective—see Section 3—an agent X , e.g., a data scientist using a model, or a clinician using a medical AI is justified in accepting the output of an AI system A as epistemically adequate for a task τ iff A is reliable for τ as in Def. 2. This is the classical target of the justificatory AI literature (Durán and Formanek 2018; Ferrario 2024). Our second perspective, see Def. 3, shifts the unit of evaluation from the AI system to the PT-HAI. Here, a third party Z , e.g., a patient, a manager, or a regulator, is justified in accepting a human-AI team output $\hat{y}^{HAI}$ insofar as it is produced by a PT-HAI that is reliable for τ . Thus, what is epistemically justified is the output of a PT-HAI whose interaction protocol determines how the human and the AI contribute to $\hat{y}^{HAI}$. This second perspective is novel; it has been hinted at in the context of trust in human-AI dyads by Ferrario and Loi (2022), but has not yet been developed in detail in justificatory AI.

A final question follows naturally: *can a PT-HAI be reliable even if the AI system, considered on its own, is not reliable in the sense of Def. 2?* On our account, the answer can be yes. An AI system may be an unreliable stand-alone predictor for τ and still contribute to a reliable PT-HAI when the interaction protocol and context of use limit its influence on the final team output. For instance, low AI reliability can trigger scrutiny, enforce escalation and review that eventually result in complementarity instances, increasing the degree of reliability of that PT-HAI.

How Does CR Address the Key Challenges of Complementarity?

Finally, we can return to the four challenges in Section 2. Our claim is not that CR solves complementarity. Rather, it repositions it, explaining what complementarity is good for, why its empirical fragility is not fatal, and what else must be measured and governed for PT-HAIs in everyday life to be epistemically defensible.

¹³For instance, let $y_i^H = y_i - \varepsilon$, $y_i^{AI} = y_i + 0.5\varepsilon$, and $y_i^{HAI} = \frac{1}{2}(y_i^H + y_i^{AI}) = y_i - 0.25\varepsilon$ with $\varepsilon > 0$. Then $L_{HAI} < \min\{L_H, L_{AI}\}$ even though all predictors are arbitrarily inaccurate.




Figure 1: Two justificatory perspectives under CR. **Left:** justification targets the AI system as a computational process in isolation as in Def. 2. **Right:** justification targets the PT-HAI as a computational process, where reliability emerges from interaction protocols, user competence, and organizational setting—see Def. 3. The same families of reliability indicators (type-RI₁-type-RI₃) apply to different targets depending on the justificatory perspective.

I. What theoretical anchoring for complementarity?

Rather than seeking a theory of when complementarity should emerge from information-processing models that are specific neither to AI systems nor to different types of human-AI interactions, CR situates complementarity in an account of process reliability and epistemic justification. Within CR, CTP matters because it can function as a reliability indicator: historical instances of complementarity provide defeasible evidence that the human-AI interaction protocol yields epistemic gain beyond either component alone, and thus that the interaction is a reliable process for task τ in context. Furthermore, instead of searching for a single, domain-general mechanism of complementarity, CR encourages researchers to identify which interactional and socio-technical features strengthen the overall reliability profile *in a given domain*, treating complementarity as one (potentially context-sensitive) signal within a family of reliability indicators.

II. Limited applicability of complementarity at decision time.

Complementarity remains an ex post metric in CR, but it is treated as historical evidence for assessing the reliability of a PT-HAI. Accordingly, it is helpful at decision-time as justification grounded in the dispositional reliability of the process, supported by past performance and other reliability indicators, rather than by a live estimate of $CTP_\tau(\cdot)$ for the current instance. In CR, the value of complementarity lies in its stability over time and in what it reveals about the drivers of effective teaming, not in serving as a decision-time property of any single interaction instance. CR also treats update-sensitivity as part of the reliability profile of the PT-HAI: if model updates, distribution shifts, or human learning undermine stability, this weakens relevant type-RI₁ and type-RI₃ indicators, such as monitoring, versioning, and communication practices.

III. More than predictive accuracy is at stake in human-AI interactions.

By embedding complementarity inside a

framework that already includes non-relative-performance desiderata via different types of reliability indicators, CR prevents CTP from being treated as a (relative) accuracy-only gold standard. CR clearly states that complementarity is epistemically attractive *only conditionally*, namely when it is achieved under constraints expressed by other reliability indicators of system robustness, human expertise, and institutional fit. Conversely, CR makes room for partially reliable PT-HAIs with limited complementarity when other reliability indicators are strong, e.g., where the key epistemic work is done by accurate humans under stringent validation, careful scoping of use, and robust governance that prevents overclaiming and out-of-scope deployment.

IV. Complementarity ignores the magnitude-cost profile of epistemic gain.

CR does not treat all complementarity improvements as equally significant. Because reliability comes in degrees and depends on the interaction of multiple indicators, small performance gains achieved at high interactional or institutional cost may carry little justificatory weight, while larger or more robust gains may significantly strengthen the reliability profile of the process. On this view, complementarity is informative insofar as it contributes—relative to its costs and in combination with other indicators—to the overall reliability of the human-AI interaction. This cost-sensitive perspective in measuring efficient complementarity can be operationalized, as we will discuss in Section 5. The extended version of this paper provides the full formal treatment of efficient complementarity, including the complementarity-gain measures (Ferrario, Facchini, and Durán 2026).

Situating Our Approach within AI Epistemology

Our contribution is not a general theory of epistemic justification for AI-assisted decision-making. Instead, we adopt CR to reposition complementarity within justificatory AI, clarifying what it can do in providing reasons to believe in the adequacy of AI-assisted predictions in real-world set-

tings where humans interact with this technology. We believe CR offers a bridge between evidence on performance, epistemic standards, and institutional practices without collapsing epistemic justification into accuracy or treating complementarity as the sole locus of normative significance in human-AI interactions. Furthermore, it allows clarifying the relation between concepts, such as reliance, complementarity, and reliability, that permeate the human-AI interaction and justificatory AI literature, but still lack a unified, cross-domain perspective (Grote, Genin, and Sullivan 2024). That said, any such choice leaves open questions that we cannot fully address here. While several classic criticisms of CR have received rebuttal (Beebe 2004; Comesaña 2006; BonJour 1980; Goldman 1986), more recent challenges are worth mentioning (Alvarado 2026). For instance, worries about *warrant transmission and credit assignment* arise. In our account, we require that reliability claims be supported by ample families of contestable indicators that can be reported in checklists to warrant the transmission of reliability—see Section 5. The extended version provides the full checklist and further implementation details (Ferrario, Facchini, and Durán 2026). Furthermore, while Alvarado (2026) emphasizes the epistemic importance of *endogenous features* (internal properties of the computational system), our account does not deny their relevance *tout court*: it treats access to internal properties as one (often optional) contributor within a broader reliability profile alongside protocol design, user competence, and governance. Finally, in response to *error-related opacity* and the difficulty of diagnosing failures from the outside (Alvarado 2026), our account makes monitoring, escalation, and update governance *type-RI₃* indicators, so that failures are anticipated, detectable, and actionable over time.

4 Three Examples of Complementarity in Computational Reliabilism

We continue by elaborating on complementarity as reliability indicator in CR through three real-world examples.

The Technologist Dermatologist

Consider an experienced dermatologist who routinely uses an AI system that classifies dermoscopic images as benign or malignant in a modern clinic. The clinician is an active technologist: they consult the system's outputs and uncertainty cues, review disagreement cases, ask colleagues or domain experts when predictions conflict with clinical knowledge, and routinely attend courses on AI-assisted dermatology. The AI system is from a well-respected company that has published many empirical studies in leading journals and documents design, implementation, and maintenance protocols of the AI extensively. Furthermore, the clinic organizes routine training courses where the accuracy of AI is compared to that of its dermatologists and human-AI teams. Over time, errors and cases of CTP are discussed, analyzing the system's failure modes to calibrate the clinician's prediction accordingly. The clinic, in turn, runs routine training and audits comparing the performance of the AI, its dermatologists, and their PT-HAIs. In CR terms, all three reliabil-

ity indicator families carry relevant justificatory weight for any patient relying on the AI-assisted clinician's prediction: they are part of highly reliable PT-HAIs. Complementarity (*type-RI₁*) strengthens reliability only insofar as gains are non-trivial and achieved without excessive burden onto the clinician. Given the increasing expertise of the dermatologist, this burden decreases over time.

All Complementarity and No Credentials Makes Jack a Reliable Student

Jack, a mathematics student, uses a free large language model (LLM) to solve exercises, especially computation-heavy ones. The system is only sometimes helpful: it produces plausible-looking mistakes, skips assumptions, and occasionally invents steps. There is no strong task-specific validation, no institutional endorsement, and no credible basis for treating the model, on its own, as a reliable source of correct proofs or calculations. What can still make the PT-HAI somewhat reliable is complementarity. Jack routinely forces the model to share alternative computations and intermediate steps, tests edge cases, and cross-checks against textbooks and lecture notes. When the LLM is wrong, Jack often detects and repairs; when Jack is wrong, the LLM's solution sometimes triggers correction. In CR terms, justificatory weight is carried mainly by *type-RI₁* through a stable pattern of error detection, correction, and complementarity gains. By contrast, *type-RI₂* and *type-RI₃* remain weak, so the overall PT-HAI displays low reliability. A second student, Lloyd, has forgotten to do his homework and asks Jack for last-minute help. On our view, Lloyd has reasons to accept the output of the Jack-LLM interaction because Jack's past performance exhibits stable complementarity under a learning protocol. These reasons could be stronger if the LLM had been extensively tested for mathematical tasks, endorsed through use in academic contexts by lecturers and students alike. But the deadline is coming soon and Lloyd has little time left.

When Complementarity Does Not Pay

A forensic lab processes a high volume of cases using an automatic speaker recognition tool to compare short, noisy robbery calls with suspect interviews. The tool can output a likelihood ratio, but only within validated conditions, e.g., duration, language, no strong disguise. Under these operational conditions, the (computational process realized by the) tool shows high reliability. Clarice, a certified forensic examiner, applies a strict protocol by checking recording quality, detecting artifacts, and documenting her assessment for court. She achieves high predictive accuracy consistently. Here, complementarity gains are often marginal: Clarice and the tool usually agree in most cases that fall within validated conditions, and the human-AI team usually improves only slightly over the better component. However, the cost of this PT-HAI is important. Each extra iteration consumes Clarice's scarce time and must be documented to a standard that survives legal and technical cross-examination, and explained in a way that is defensible in court. Small gains usually translate into relevant operational

penalties, including delayed investigations, postponed hearings, and growing court backlogs. The forensic lab treats certified cases of Clarice-tool complementarity as welcome in those cases where operational costs are limited, but not a target to optimize. The reliability of the forensic lab investigation is secured mainly by `type-RI3` constraints—validated scope, abstention triggers, peer review, proficiency testing, version control, and reporting conventions—precisely because they keep the process stable under high throughput and expert scarcity; this indicator family gives reasons to judges and other legal professionals to endorse the lab’s outputs. Thus, on a CR view, this PT-HAI can be highly reliable even when complementarity gains are small, because what matters is not squeezing out marginal accuracy, but maintaining a defensible and court-robust epistemic practice.

5 Recommendations for Design, Governance and Future Research

Our conceptual repositioning of complementarity changes what complementarity is evidence for, how failures to achieve it should be interpreted, and what else must be measured and governed for AI-assisted decision-making to be epistemically defensible. We summarize implications of our approach for design and governance, operationalizing them through an assurance checklist and measures of efficient complementarity.

Recommendations for Human-AI Design

Designers should not optimize for complementarity alone. Instead, they should treat it as one diagnostic signal within a broader reliability profile and report, at minimum, whether complementarity occurs, how large the gain is, how stable it remains over time, and what interactional and organizational costs are required to sustain it. In practice, this means prioritizing interaction protocols that make epistemically relevant behavior observable and correctable—for instance, structured contestation, escalation pathways, and explicit disagreement handling—rather than relying primarily on explanations or confidence displays that do not reliably improve teaming (Vaccaro, Almaatouq, and Malone 2024). Finally, evaluation should report whether complementarity occurs and the magnitude, stability over time, and cost of the gain, since these dimensions determine its justificatory weight in CR. Practically, human-AI interaction design should use a minimal assurance checklist that makes PT-HAI reliability assessable by operators, governance and policy organizations, and affected third parties, e.g., patients, managers, regulators, who otherwise lack direct epistemic access to the PT-HAI protocol and its evidence base. The extended version provides a minimal reporting checklist that designers can treat as a deployment deliverable: it documents the interaction protocol, provides interaction-level performance evidence, and encodes lifecycle commitments, including monitoring and update governance that foster PT-HAI reliability over time (Ferrario, Facchini, and Durán 2026). Finally, designers should target and report *efficient complementarity* explicitly. Efficient complementarity obtains when the observed complementarity gain is worth its interactional and

organizational cost: $\Delta_{\tau}^{\text{net}}(D) = \Delta_{\tau}(D) - \lambda c(D) > 0$. Here, $c(D)$ captures costs such as monitoring, review, documentation, escalation, or reduced throughput, while λ is an institutionally justified threshold specifying the minimum gain per unit cost. Thus, complementarity should be reported together with its gain, cost, and threshold, because otherwise it may merely shift epistemic burden onto users. The extended version provides the full formal treatment of efficient complementarity (Ferrario, Facchini, and Durán 2026). Finally, longitudinal studies of PT-HAIs should measure the human side of the interaction, not only the AI side or aggregate team performance. Relevant indicators include cognitive burden, fatigue, and possible deskilling over time. These factors directly affect whether complementarity remains stable: they should therefore be treated as part of the `type-RI1/type-RI3` evidence base for PT-HAI reliability.

Recommendations for AI Governance and Policy

Regulators and institutions should avoid mandating complementarity as a compliance target, because doing so encourages optimization of a relative, binary metric while leaving the broader reliability of the socio-technical process unspecified. Governance should evaluate PT-HAIs as evolving socio-technical practices rather than as static performance snapshots. In high-stakes settings, this requires collecting evidence of reliability through lifecycle-sensitive obligations such as monitoring, incident reporting, update documentation, revalidation triggers, and competence maintenance, since distribution shift, model updates, and human learning can rapidly erode the historical basis for justification. This lifecycle-sensitive view on human-AI team reliability aligns with regulatory logics already visible in high-risk AI governance, for instance in the EU AI Act’s emphasis on post-deployment monitoring, technical documentation, and change management—see Article 3(20) and Annex IV, Section 9 (EU AI Act 2024). More generally, policy should enforce proportionality: epistemic gains attributable to human-AI interaction should be commensurate with the cognitive and institutional burdens needed to sustain them, and small or fragile gains should trigger tighter scoping and stronger oversight requirements.

6 Conclusions

Complementarity remains a valuable concept in human-AI interaction research, but not in the role it is currently asked to play. When treated as a stand-alone, ground-truth-dependent, relative-accuracy criterion, it is theoretically under-anchored and empirically fragile. When repositioned within computational reliabilism, it acquires a clearer role: it functions as historically grounded evidence about whether a prediction-task human-AI interaction is a reliable epistemic process for a task. This theoretical reframing turns complementarity from a binary success label into one component of a broader, graded account of epistemic justification that must also attend to validity, uncertainty discipline, competence scaffolding, and lifecycle governance if AI-assisted decision-making is to be epistemically defensible in high-stakes settings.

Acknowledgments

The work of A. Ferrario and A. Facchini was partly conducted within the framework of the EUonAIR Centre of Excellence in Responsible AI and Education. Their work was partially supported by a grant from Movetia, funded by the Swiss Confederation.

References

- Alston, W. P. 1995. How to think about reliability. *Philosophical Topics*, 23(1): 1–29.
- Alufaisan, Y.; Marusich, L. R.; Bakdash, J. Z.; Zhou, Y.; and Kantarcioglu, M. 2021. Does explainable artificial intelligence improve human decision-making? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 6618–6626.
- Alur, R.; Raghavan, M.; and Shah, D. 2024. Human expertise in algorithmic prediction. *Advances in Neural Information Processing Systems*, 37: 138088–138129.
- Alvarado, R. 2023a. AI as an epistemic technology. *Science and Engineering Ethics*, 29(5): 32.
- Alvarado, R. 2023b. What kind of trust does AI deserve, if any? *AI and Ethics*, 3(4): 1169–1183.
- Alvarado, R. 2026. Challenges for computational reliabilism in AI and other computational methods. In *Philosophy of Science for Machine Learning: Core Issues and New Perspectives*, 81–106. Springer.
- Bansal, G.; Nushi, B.; Kamar, E.; Horvitz, E.; and Weld, D. S. 2021a. Is the most accurate AI the best teammate? Optimizing AI for teamwork. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 11405–11414.
- Bansal, G.; Wu, T.; Zhou, J.; Fok, R.; Nushi, B.; Kamar, E.; Ribeiro, M. T.; and Weld, D. 2021b. Does the whole exceed its parts? The effect of AI explanations on complementary team performance. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, 1–16.
- Beebe, J. R. 2004. The generality problem, statistical relevance and the tri-level hypothesis. *Noûs*, 38(1): 177–195.
- Benk, M.; Kerstan, S.; von Wangenheim, F.; and Ferrario, A. 2025. Twenty-four years of empirical research on trust in AI: A bibliometric review of trends, overlooked issues, and future directions. *AI & Society*, 40(4): 2083–2106.
- BonJour, L. 1980. Externalist theories of empirical knowledge. *Midwest studies in philosophy*, 5(1): 53–74.
- Cabitza, F.; Campagner, A.; and Sconfienza, L. M. 2021. Studying human-AI collaboration protocols: The case of the Kasparov’s law in radiological double reading. *Health Information Science and Systems*, 9(1): 8.
- Clark, A.; and Chalmers, D. 1998. The extended mind. *Analysis*, 58(1): 7–19.
- Comesaña, J. 2006. A well-founded solution to the generality problem. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 129(1): 27–47.
- Donahue, K.; Chouldechova, A.; and Kenthapadi, K. 2022. Human-algorithm collaboration: Achieving complementarity and avoiding unfairness. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1639–1656.
- Durán, J. M. 2025. In defense of reliabilist epistemology of algorithms. *European Journal for Philosophy of Science*, 15(2): 37.
- Durán, J. M. 2026. Beyond transparency: Computational reliabilism as an externalist epistemology of algorithms. In *Philosophy of Science for Machine Learning: Core Issues and New Perspectives*, 50–78. Springer.
- Durán, J. M.; and Formanek, N. 2018. Grounds for trust: Essential epistemic opacity and computational reliabilism. *Minds and Machines*, 28(4): 645–666.
- Durán, J. M.; and Pozzi, G. 2025. Trust and Trustworthiness in AI. *Philosophy & Technology*, 38(1): 16.
- Durán, J. M.; van der Vloed, D.; Ruifrok, A.; and Ypma, R. J. 2024. From understanding to justifying: Computational reliabilism for AI-based forensic evidence evaluation. *Forensic Science International: Synergy*, 9: 100554.
- Dvijotham, K.; Winkens, J.; Barsbey, M.; Ghaisas, S.; Stanforth, R.; Pawlowski, N.; Strachan, P.; Ahmed, Z.; Azizi, S.; Bachrach, Y.; et al. 2023. Enhancing the reliability and accuracy of AI-enabled diagnosis via complementarity-driven deferral to clinicians. *Nature Medicine*, 29(7): 1814–1820.
- D’Arcy, J.; Gupta, A.; Tarafdar, M.; and Turel, O. 2014. Reflecting on the “dark side” of information technology use. *Communications of the Association for Information Systems*, 35(1): 5.
- EU AI Act. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence. Technical report, European Union.
- Ferrario, A. 2024. Justifying our credences in the trustworthiness of AI systems: A reliabilistic approach. *Science and Engineering Ethics*, 30(6): 55.
- Ferrario, A. 2025. Being pragmatic about reliance and trust in artificial intelligence. *Minds and Machines*, 36(1): 5.
- Ferrario, A.; Facchini, A.; and Durán, J. M. 2026. Epistemology Gives a Future to Complementarity in Human-AI Interactions. Extended version, arXiv:2601.09871.
- Ferrario, A.; and Loi, M. 2022. How explainability contributes to trust in AI. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1457–1466.
- Geirhos, R.; Narayanappa, K.; Mitzkus, B.; Thieringer, T.; Bethge, M.; Wichmann, F. A.; and Brendel, W. 2021. Partial success in closing the gap between human and machine vision. *Advances in Neural Information Processing Systems*, 34: 23885–23899.
- Goldman, A. I. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Gonzalez, C.; Donahue, K.; Goldstein, D. G.; Heidari, H.; Jalali, M. S.; Schelble, B.; Singh, A.; and Woolley, A. W. 2026. Toward a science of human–AI teaming for decision making: A complementarity framework. *PNAS nexus*, 5(3): pgag030.

- Grote, T.; Genin, K.; and Sullivan, E. 2024. Reliability in machine learning. *Philosophy Compass*, 19(5): e12974.
- Heersmink, R. 2015. Dimensions of integration in embedded and extended cognitive systems. *Phenomenology and the Cognitive Sciences*, 14(3): 577–598.
- Hemmer, P.; Schemmer, M.; Kühl, N.; Vössing, M.; and Satzger, G. 2025. Complementarity in human-AI collaboration: Concept, sources, and evidence. *European Journal of Information Systems*, 1–24.
- Hemmer, P.; Schemmer, M.; Vössing, M.; and Kühl, N. 2021. Human-AI complementarity in hybrid intelligence systems: A structured literature review. *PACIS*, 78: 118.
- Hollan, J.; Hutchins, E.; and Kirsh, D. 2000. Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(2): 174–196.
- Hutchins, E. 1995. *Cognition in the Wild*. MIT Press.
- Inkpen, K.; Chappidi, S.; Mallari, K.; Nushi, B.; Ramesh, D.; Michelucci, P.; Mandava, V.; Vepřek, L. H.; and Quinn, G. 2023. Advancing human-AI complementarity: The impact of user expertise and algorithmic tuning on joint decision making. *ACM Transactions on Computer-Human Interaction*, 30(5): 1–29.
- Ma, S.; Lei, Y.; Wang, X.; Zheng, C.; Shi, C.; Yin, M.; and Ma, X. 2023. Who should I trust: AI or myself? Leveraging human and AI correctness likelihood to promote appropriate trust in AI-assisted decision-making. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–19.
- Miller, T. 2023. Explainable AI is dead, long live explainable AI! Hypothesis-driven decision support using evaluative AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 333–342.
- Rastogi, C.; Leqi, L.; Holstein, K.; and Heidari, H. 2023. A taxonomy of human and ML strengths in decision-making to investigate human-ML complementarity. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 11, 127–139.
- Ryan, M. 2020. In AI we trust: Ethics, artificial intelligence, and reliability. *Science and Engineering Ethics*, 26(5): 2749–2767.
- Schemmer, M.; Kuehl, N.; Benz, C.; Bartos, A.; and Satzger, G. 2023. Appropriate reliance on AI advice: Conceptualization and the effect of explanations. In *Proceedings of the 28th International Conference on Intelligent User Interfaces, IUI '23*, 410–422. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701061.
- Turel, O.; and Kalhan, S. 2023. Prejudiced against the machine? Implicit associations and the transience of algorithm aversion. *MIS Quarterly*, 47(4): 1369–1394.
- Vaccaro, M.; Almaatouq, A.; and Malone, T. 2024. When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8(12): 2293–2303.
- Vaccaro, M.; and Waldo, J. 2019. The effects of mixing machine learning and human judgment. *Communications of the ACM*, 62(11): 104–110.
- Zhang, Y.; Liao, Q. V.; and Bellamy, R. K. 2020. Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 295–305.