

Align When They Want, Complement When They Need! Human-Centered Ensembles for Adaptive Human-AI Collaboration

Hasan Amin, Ming Yin, Rajiv Khanna

Purdue University, USA
{hasanamin, mingyin, rajivak}@purdue.edu

Abstract

In human-AI decision making, designing AI that complements human expertise has been a natural strategy to enhance human-AI collaboration, yet it often comes at the cost of decreased AI performance in areas of human strengths. This can inadvertently erode human trust and cause them to ignore AI advice precisely when it is most needed. Conversely, an aligned AI fosters trust yet risks reinforcing suboptimal human behavior and lowering human-AI team performance. In this paper, we start by identifying this fundamental tension between performance-boosting (i.e., *complementarity*) and trust-building (i.e., *alignment*) as an inherent limitation of the traditional approach for training a single AI model to assist human decision making. To overcome this, we introduce a novel, *human-centered adaptive AI ensemble* that strategically toggles between two specialist AI models—the aligned model and the complementary model—based on contextual cues, using an elegantly simple yet provably near-optimal *Rational Routing Shortcut* mechanism. Comprehensive theoretical analyses elucidate why the adaptive AI ensemble is effective and when it yields maximum benefits. Moreover, experiments on both simulated and real-world data show that when humans are assisted by the adaptive AI ensemble in decision making, they can achieve significantly higher performance than when they are assisted by single AI models that are trained to either optimize for its independent performance or even the human-AI team performance.

Supplementary material —

<https://github.com/shasanamin/aaai26-adaptive-ai>

Introduction

AI systems are increasingly being developed to support human decision making in various domains, yet the grand promise of effective human-AI collaboration remains largely unfulfilled. Contrary to intuition, human-AI teams routinely underperform compared to their best constituent, i.e., the human or AI alone (Dell’Acqua, Kogut, and Perkowski 2023; Agarwal et al. 2023; Vaccaro, Almaatouq, and Malone 2024; Goh et al. 2024). This recurring failure signals a fundamental gap in our current understanding and raises a deeper question at the core of collaborative AI design: *what makes an AI a truly effective teammate?*

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

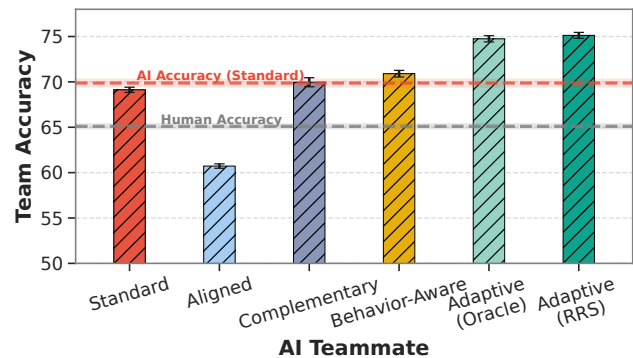


Figure 1: Summary comparison of human-AI team performance across AI design paradigms. Standard AI optimizes independent accuracy; aligned and complementary AIs specialize in trust-building and error-correction, respectively; behavior-aware AI optimizes team loss under our human interaction model. Our proposed adaptive AI ensemble, including the more practical one using Rational Routing Shortcut (RRS), strategically toggles between aligned and complementary AI, and achieves the highest team accuracy. Results shown here correspond to the behavior-grounded evaluation on WoofNette data.

A dominant paradigm in human-AI collaboration emphasizes *complementarity*: optimizing AI to be correct on instances where human judgment is weak (Madras, Pitassi, and Zemel 2018; Lai et al. 2022; Steyvers et al. 2022; Zhang, Lee, and Carter 2022; Holstein and Alevin 2021). While this approach can elevate team performance in principle, it often neglects a crucial psychological and behavioral reality: human trust in AI is not merely a function of AI’s objective accuracy, but hinges disproportionately on AI’s alignment with human judgment, especially on instances where humans feel more confident (Lu and Yin 2021; Chong et al. 2022). When such complementary AI diverges from human intuition on instances where humans are confident, even if wrongly so, it can degrade trust and reliance, resulting in users paradoxically ignoring AI advice precisely when it would have been most valuable (Lu and Yin 2021; Lu, Wang, and Yin 2024). Conversely, an AI that focuses on *alignment* with human judgment may foster trust but also reinforce suboptimal

human decisions, squandering the performance gains that complementary AI can deliver (Grgić-Hlača, Castelluccia, and Gummadi 2022; Askell et al. 2021; Ouyang et al. 2022; Guo et al. 2024). Recognizing the importance of human-AI team performance, some recent work has attempted to optimize for team decision making directly, taking human-AI interaction behavior into account (Bansal et al. 2021; Mahmood, Lu, and Yin 2024). Such behavior-aware AI has been a promising proof-of-concept, already leading to some gains in team performance compared to human or AI alone, despite strong assumptions like rationality (Bansal et al. 2021) or confidence thresholding (Mahmood, Lu, and Yin 2024).

In this work, we start by extending the behavior-aware AI framework to incorporate more realistic human behavior, capturing the dual need to foster trust through alignment and boost performance through complementarity. We introduce a behaviorally grounded *Confidence-Gated Probabilistic Reliance (CGPR)* model that formally links human self-confidence and trust to the likelihood of accepting AI recommendations. Building on this, we derive the optimal single-model, behavior-aware AI that directly optimizes team decision making performance. In doing so, we uncover an intriguing phenomenon: the “*complementarity-alignment tradeoff*.” That is, even the best single-model, behavior-aware AI cannot simultaneously optimize for both trust (through alignment) and team performance (through complementarity) across all regions of the decision space.

In light of these fundamental limitations of the single-model paradigms, we supplement the behavior-aware AI with a new direction: *human-centered adaptive AI ensembles* that dynamically toggle between alignment and complementarity—preserving trust where it matters most, and providing support where it helps most. In particular, we propose a practicable instantiation of the adaptive AI framework that learns separate specialist AI models for alignment and complementarity, and then dynamically selects between them at inference time using *Rational Routing Shortcut (RRS)*, a simple and theoretically principled mechanism that approximates human-aware, optimal routing without requiring access to private or hard-to-measure human states. We further provide rigorous theoretical analyses quantifying when and why adaptive AI yields substantial team performance gains over any single-model alternative. Finally, we showcase the generalizability and robustness of our theoretical insights through simulations and a behavior-grounded image decision making benchmark constructed from real human accuracy and confidence data. Notably, our adaptive AI ensemble achieves remarkable empirical improvements in enhancing human-AI team accuracy in decision making in practice: up to 9% increase in team accuracy over standard AI optimized solely for independent accuracy, and 6% over behavior-aware AI optimized explicitly for team performance (see Figure 1)—remarkably, while relying on specialized AI components individually less accurate than the standard AI model.

While empirical work on human-AI collaboration is rapidly growing, a formal understanding of the interplay between team performance, complementarity and alignment is severely lacking. Our work provides one of the first foun-

dational frameworks to characterize and address this gap, and highlights a shift in human-AI collaboration: moving from static, single-model approaches to *adaptive, behaviorally grounded, and theoretically grounded* systems that aim to “align when desired, and complement when needed.” Our main contributions are:

1. **Principled human behavior modeling and complementarity-alignment tradeoff characterization:** We introduce CGPR as a realistic, behaviorally-grounded model of human-AI interaction in AI-assisted decision making. Through this, we provide the first rigorous characterization of a fundamental complementarity-alignment tradeoff, proving why single-model approaches are inherently limited. Our models and analyses incorporate important behavioral factors, such as confidence and trust, into the core of the machine learning, rather than treating them as secondary considerations or empirical observations.
2. **Practical adaptive AI ensemble framework:** We propose a human-centered adaptive AI ensemble to overcome tradeoffs like above. We offer a practical instantiation that trains specialist models for alignment and complementarity, and then efficiently toggles between them at test time through RRS, without requiring access to private or hard-to-measure human data. The benefits of our framework are validated through comprehensive experimentation. Our approach is flexible and extensible to a broad range of models and applications.
3. **Extensive theoretical analyses and insights:** We provide extensive theoretical guarantees that bring a new level of rigor to human-centered machine learning. This includes precisely bounding the complementarity-alignment tradeoff (Theorem 2) as well as the performance gains of adaptive ensemble over optimal single model (Theorem 4), offering fresh perspective on when and why adaptive AI is most useful in practice. These results establish a new class of human-centered theoretical tools.

Other related work. Distinct from mainstream research on human-AI collaboration—which largely optimizes for complementarity or alignment, often overlooking nuanced behavioral dynamics—our work builds on and advances “behavior-aware” frameworks (Bansal et al. 2021; Mahmood, Lu, and Yin 2024) by introducing more realistic cognitive modeling and rigorous theory. Parallel efforts on “learning to defer” (Madras, Pitassi, and Zemel 2018; Wilder, Horvitz, and Kamar 2020; Bondi et al. 2022; Dvijotham et al. 2023) investigate the *division of labor* between humans and AI, typically assuming the AI can act as the final decision maker; in contrast, we focus on *AI-assisted decision making* where the human always retains control and AI provides recommendations. Our approach uniquely formalizes the interplay between trust, confidence, and team performance, moving beyond simplified reliance models and establishing theoretically grounded methods for adaptive, human-centered teaming.

The Complementarity-Alignment Tradeoff in Human-AI Collaboration

Problem Setup

In human-AI joint decision making, each instance of a decision making case is represented by features $\mathbf{x} \in \mathcal{X}$, and the objective is to predict an outcome $y \in \mathcal{Y}$. In our *AI-assisted decision making* setup, an AI model m provides a recommendation $y_m = m(\mathbf{x}; \theta_m)$ to a human decision maker (DM) h who has an independent judgment $y_h = h(\mathbf{x}; \theta_h)$. The final decision $d \in \mathcal{Y}$ is made by the human DM, based on their own initial judgment y_h and the AI’s recommendation y_m .

Typically, AI models are trained independently to minimize empirical risk over a dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$, by optimizing $\theta_m = \arg \min_{\theta_m} \frac{1}{N} \sum_{i=1}^N \ell(m(\mathbf{x}_i; \theta_m), y_i)$, where $\ell(\cdot)$ is an appropriate loss function. However, by optimizing for AI’s independent accuracy instead of *team performance*, this approach fails to account for how the human integrates the AI’s recommendations into their final decision. To model this human-AI interaction, we define a team decision making model $f(\cdot)$, which depends on the data instance, the AI recommendation, and the human DM’s own judgment. The team loss optimization is then formulated as:

$$\theta_m = \arg \min_{\theta_m} \frac{1}{N} \sum_{i=1}^N \ell(f(\mathbf{x}_i, m(\mathbf{x}_i; \theta_m), h(\mathbf{x}_i; \theta_h)), y_i)$$

Modeling Human Behavior

Optimizing human-AI team performance critically depends on accurately modeling how human DMs factor AI recommendations into their decisions. Prior work introduced a deterministic, confidence-thresholded reliance rule (Mahmood, Lu, and Yin 2024), where a human with confidence $\mathcal{C}^h(\mathbf{x})$ accepts the AI’s prediction y_m only when their own confidence falls below a threshold τ . Formally, this *Confidence-Gated Reliance (CGR)* model takes the form:

$$f_{\text{CGR}}(\mathbf{x}, m, h) = \begin{cases} y_h & \text{if } \mathcal{C}^h(\mathbf{x}) > \tau, \\ y_m & \text{otherwise.} \end{cases} \quad (1)$$

While analytically convenient, accumulating empirical evidence shows that human reliance on AI advice is often more nuanced and probabilistic, influenced significantly by human trust and observed AI behavior (Chong et al. 2022; Schemmer et al. 2023; Lu and Yin 2021; Chiang and Yin 2021; Narayanan et al. 2023).

To address this gap, we propose the *Confidence-Gated Probabilistic Reliance (CGPR)* model. Let $\mathcal{D}_a = \{\mathbf{x} : \mathcal{C}^h(\mathbf{x}) > \tau\}$ and $\mathcal{D}_c = \{\mathbf{x} : \mathcal{C}^h(\mathbf{x}) \leq \tau\}$ denote the “alignment” (high-confidence) and “complementarity” (low-confidence) regions, respectively. In \mathcal{D}_a , the human follows their own judgment, as in CGR. In \mathcal{D}_c , however, the DM relies on the AI prediction y_m *probabilistically*. Recent evidence suggests that trust is disproportionately influenced by how often the AI contradicts the human when the human feels confident (Lu and Yin 2021; Grgić-Hlača, Castelluccia, and Gummadi 2022; Wang, Lu, and Yin 2022). We thus model the reliance probability r to be governed

by the perceived alignment between the human and the AI in the high-confidence region: $r = 1 - L_h(\mathcal{D}_a, m)$, with $L_h(\mathcal{D}_a, m) = \mathbb{E}_{\mathbf{x} \in \mathcal{D}_a} [\ell(m(\mathbf{x}), h(\mathbf{x}))]$ quantifying disagreement between the AI and the human. We take $\ell(\cdot, \cdot)$ to be 0-1 loss here, so that $r \in [0, 1]$ represents a probability. Putting these components together, the CGPR team decision rule is:

$$f(\mathbf{x}, m, h) = \begin{cases} y_h & \text{if } \mathcal{C}^h(\mathbf{x}) > \tau, \\ y_m & \text{with prob. } r \text{ if } \mathcal{C}^h(\mathbf{x}) \leq \tau, \\ y_h & \text{otherwise.} \end{cases} \quad (2)$$

Here, we interpret trust as the human’s willingness to follow the AI’s recommendation, i.e., the reliance probability r in CGPR, and thus use “trust” and “reliance” interchangeably. The CGPR model uniquely integrates human confidence, probabilistic reliance, and alignment-driven trust into a coherent and optimizable framework, significantly extending prior approaches and providing a principled foundation for (adaptive) human-centered AI.

Optimizing Behavior-Aware AI

Building on the behavior-aware AI framework, we derive the optimal single-model AI model that directly optimizes team performance, accounting for the probabilistic and confidence-dependent nature of human reliance. Using Eq. 2, the team loss can be decomposed as:

$$L_{\text{team}}(\mathcal{D}, m, h) = L(\mathcal{D}_a, h) + L(\mathcal{D}_c, m) + [L(\mathcal{D}_c, h) - L(\mathcal{D}_c, m)] L_h(\mathcal{D}_a, m) \quad (3)$$

Since we cannot control a *given* human DM h , we consider the human-only components $L_{a,h} := L(\mathcal{D}_a, h)$ and $L_{c,h} := L(\mathcal{D}_c, h)$ to be constant for optimization purpose. Therefore, the optimal, behavior-aware AI parameters are obtained by solving the following optimization problem:

$$\theta_{ba}^* = \arg \min_{\theta_{ba}} L(\mathcal{D}_c, m) + [L_{c,h} - L(\mathcal{D}_c, m)] L_h(\mathcal{D}_a, m) \quad (4)$$

Multiobjective optimization and the complementarity-alignment tradeoff. The behavior-aware optimization objective surfaces a tension between two desirable but conflicting goals: complementarity, achieved by minimizing $L(\mathcal{D}_c, m)$, and alignment, achieved by minimizing $L_h(\mathcal{D}_a, m)$. Worse, their coupling in Eq. 4 creates a compounding effect: as the model becomes better at complementing ($L(\mathcal{D}_c, m)$ decreases), the influence of alignment loss grows (via the term $[L_{c,h} - L(\mathcal{D}_c, m)]$), potentially increasing total team loss despite localized gains. One could attempt to address this through standard multi-objective optimization, e.g., by minimizing a weighted combination $w \cdot L(\mathcal{D}_c, m) + (1-w) \cdot L_h(\mathcal{D}_a, m)$, thereby choosing a point on the Pareto frontier. However, this strategy merely shifts the compromise instead of resolving it. The core issue is more fundamental: improving complementarity often entails diverging from human judgment patterns (in low-confidence regions), while improving alignment necessitates mimicking human behavior (in high-confidence regions). These objectives could potentially be diametrically opposite. This reveals a critical insight that optimizing a single model to balance both alignment and complementarity is not only inefficient but also often inherently suboptimal.

Theoretical Analysis: One Model is Not Enough

Next, we formalize the conflict between the two objectives—complementarity and alignment—as a direct tradeoff, demonstrating that under the single-model paradigm, incremental improvements in alignment can lead to disproportionate degradation in complementarity, and vice versa. For illustrative simplicity, our theoretical analysis is grounded in the geometry of the logistic loss with ℓ_2 -regularization, enabling fully data-dependent guarantees. Detailed discussion and proofs are deferred to Appendix.

Disentangling human and model factors. To effectively incorporate the contribution of human factors to the tradeoff, we first derive how model disagreement with human judgments in the alignment region is linked to the model’s ground-truth prediction loss within the same region.

Lemma 1 (Alignment Loss Sensitivity). *The alignment loss $L_h(\mathcal{D}_a, m)$ can be decomposed as:*

$L_h(\mathcal{D}_a, m) = \alpha \cdot L(\mathcal{D}_a, m) + (1 - \alpha) [1 - L(\mathcal{D}_a, m)]$, where α is human accuracy in the alignment region \mathcal{D}_a . Therefore, the sensitivity of alignment loss to changes in model’s prediction loss can be quantified as:

$$\frac{\partial L_h(\mathcal{D}_a, m)}{\partial L(\mathcal{D}_a, m)} = 2\alpha - 1. \quad (5)$$

This means attempting to improve model at mimicking human judgments effectively scales the required changes in the model’s ability to predict ground truth by $(2\alpha - 1)^{-1}$. For small α , the alignment loss becomes largely insensitive to prediction loss, making it difficult to improve alignment by improving underlying ground-truth accuracy, and vice-versa. Note that this decomposition is obtained under conditional independence assumption in binary classification setting, with a comprehensive analysis deferred to Appendix.

Quantifying the complementarity-alignment tradeoff. Let $L_a(\theta) := L_h(\mathcal{D}_a, m(\cdot; \theta))$ and $L_c(\theta) := L(\mathcal{D}_c, m(\cdot; \theta))$ be the alignment and complementarity objectives, each a logistic loss with ℓ_2 -penalty. Denote by $\theta_{m_a}^*$ and $\theta_{m_c}^*$ the corresponding minimizers, with θ denoting parameters for an arbitrary third model—including possibly θ_{ba}^* from Eq. 4—in the region of interest. We consider a closed neighborhood \mathcal{R} that contains all three points, and let λ_{\max}^a and λ_{\min}^c be the maximum and minimum eigenvalues of the Hessians of L_a and L_c over \mathcal{R} . Define the *effective curvature ratio* $\lambda_r := \lambda_{\min}^c / \lambda_{\max}^a$, which captures the ratio between how quickly alignment can improve and how slowly complementarity can degrade in the worst case. For notational convenience, we also define the human-accuracy factor as $\kappa := 2\alpha - 1$, and quantify the distance from the specialist models to a given model as $d_c := \|\theta - \theta_{m_c}^*\|$ and $d_a := \|\theta - \theta_{m_a}^*\|$.

Theorem 2 (Complementarity-Alignment Tradeoff). *For any $\theta \in \mathcal{R}$, define the local unit tradeoff*

$$\begin{aligned} \mathcal{T}(\theta) &:= \lim_{\varepsilon \downarrow 0} \frac{L_c(\theta - \varepsilon \nabla L_a(\theta)) - L_c(\theta)}{(L_a(\theta - \varepsilon \nabla L_a(\theta)) - L_a(\theta))} \\ &= - \frac{\nabla L_c(\theta)^\top \nabla L_a(\theta)}{\|\nabla L_a(\theta)\|^2}, \end{aligned}$$

i.e. the instantaneous increase in complementarity loss per unit decrease in alignment loss when taking the steepest descent step for alignment. For every $\theta \in \mathcal{R}$,

$$\mathcal{T}(\theta) \geq \frac{\lambda_r}{\kappa} \frac{d_c}{d_a} (-\cos\phi(\theta)), \quad (6)$$

where $\phi(\theta)$ is the angle between $\nabla L_a(\theta)$ and $\nabla L_c(\theta)$. Moreover, if the gradients remain sufficiently opposed in the descent direction in a neighborhood of $\theta_{m_a}^$, i.e., $-\cos\phi(\theta) \geq c_0 > 0$ for some constant c_0 , then $\lim_{\theta \rightarrow \theta_{m_a}^*} \mathcal{T}(\theta) = +\infty$ at rate $\Omega\left(\frac{1}{d_a}\right)$.*

Theorem 2 starkly illustrates why a single model cannot satisfy both the objectives simultaneously. In fact, whenever the specialist solutions are distinct and human accuracy is imperfect, it becomes mathematically impossible to jointly achieve optimal alignment and complementarity. In the extreme case where the human’s high confidence decisions are essentially guesses ($\alpha \rightarrow 0.5$), the tradeoff becomes unbounded. Under realistic conditions, *one model is not enough* for truly effective human-AI collaboration.

Human-Centered Adaptive AI Ensembles

The fundamental tradeoff established in last section suggests that we must move beyond the single model paradigm. Fortunately, unlike classic tradeoffs that are uniform across the input space, adequate human behavior modeling guides us that the demands here vary *systematically* across regions: trust is essential where humans are highly self-confident, while performance is essential where they are not. Thus to effectively navigate the tension, we propose an adaptive ensemble of human-centered AI models. The key idea is to leverage two specialist models, each dedicated to its own specific objective of alignment or complementarity, and to dynamically route decisions between them based on context, thereby achieving the best of both worlds.

Adaptive AI for Complementarity and Alignment

Rather than forcing a single AI model onto the complementarity-alignment Pareto frontier, we can do better by training separate specialist AI models. The two specialists in our case naturally stem from Eq. 4 and the identified complementarity-alignment tradeoff. We thus train a *complementary AI* m_c parameterized by θ_{m_c} , which minimizes prediction loss explicitly in the complementarity region ($\theta_{m_c}^* = \arg \min L(\mathcal{D}_c, m_c)$), and an *aligned AI* m_a parameterized by θ_{m_a} , which minimizes the disagreement with human judgments in the alignment region ($\theta_{m_a}^* = \arg \min L_h(\mathcal{D}_a, m_a)$).

Having trained specialized AI models, our main challenge is to decide, for each instance, which model’s advice to present to the human DM. An ideal solution, termed *oracle routing*, with access to human confidence and threshold used within CGPR model (Eq. 2) “knows” whether a test instance belongs to the alignment or complementarity region, and routes accordingly:

$$m_{\text{oracle}}(\mathbf{x}) = \begin{cases} m_a(\mathbf{x}) & \text{if } \mathbf{x} \in \mathcal{D}_a \text{ (i.e., } C^h(\mathbf{x}) > \tau), \\ m_c(\mathbf{x}) & \text{otherwise.} \end{cases} \quad (7)$$

Rational Routing Shortcut. The oracle routing requires directly observing the human DM’s confidence and comparing it against their threshold, yet accurately estimating such human internal states is difficult, noisy, and context-dependent. To bypass this complexity, we propose the *Rational Routing Shortcut (RRS)* mechanism, which makes instance-wise routing decisions solely on the confidence estimates of the specialized models themselves. Specifically, let $\mathcal{C}^c(\mathbf{x})$ and $\mathcal{C}^a(\mathbf{x})$ be the respective confidence of the complementary and aligned specialists in their predictions. RRS routes to the specialist with greater confidence:

$$m_{\text{RRS}}(\mathbf{x}) = \begin{cases} m_a(\mathbf{x}) & \text{if } \mathcal{C}^a(\mathbf{x}) \geq \mathcal{C}^c(\mathbf{x}), \\ m_c(\mathbf{x}) & \text{otherwise.} \end{cases} \quad (8)$$

Why RRS intuitively works? RRS implicitly capitalizes on the intuition that higher model confidence is indicative of an instance falling within that model’s designated region of expertise, using confidence as an *implicit signal for region membership*. Hence, RRS intelligently approximates oracle-like routing without ever directly observing human confidence. Remarkably, this simple shortcut turns out to not only be intuitively appealing but also provably near-optimal for team decision making:

Theorem 3 (Near-Oracle Guarantee for RRS). *Let \mathcal{D}_a and \mathcal{D}_c denote the alignment and complementarity regions, and let m_a, m_c be the corresponding aligned and complementary specialists. Suppose we are given real-valued confidence-estimation functions $\mathcal{C}^a, \mathcal{C}^c : \mathcal{X} \rightarrow [0, 1]$. Fix $\varepsilon \in [0, 1]$ and suppose that for every instance \mathbf{x} :*

- (i) (**Calibration**) $|\mathcal{C}^c(\mathbf{x}) - \mathbb{P}[m_c(\mathbf{x}) = y \mid \mathbf{x}]| \leq \varepsilon$ and $|\mathcal{C}^a(\mathbf{x}) - \mathbb{P}[m_a(\mathbf{x}) = h(\mathbf{x}) \mid \mathbf{x}]| \leq \varepsilon$.
- (ii) (**Estimator dominance in alignment region**) If $\mathbf{x} \in \mathcal{D}_a$ then $\mathcal{C}^a(\mathbf{x}) \geq \mathcal{C}^c(\mathbf{x})$.
- (iii) (**Bounded sub-optimality outside alignment region**) If $\mathbf{x} \in \mathcal{D}_c$ and $\mathcal{C}^a(\mathbf{x}) \geq \mathcal{C}^c(\mathbf{x})$, then $\mathbb{P}[m_a(\mathbf{x}) = y] \geq \mathbb{P}[m_c(\mathbf{x}) = y] - \varepsilon$.

Then the expected CGPR team accuracy achieved by the Rational Routing Shortcut satisfies

$$\text{Accuracy}_{\text{RRS}} \geq \text{Accuracy}_{\text{Oracle}} - \varepsilon.$$

Theoretical Analysis: The Benefit of Two Models

To understand when and why our adaptive ensemble is provably advantageous, we analyze the performance gap between the (optimal) single-model solution and an oracle-style adaptive ensemble that uses the alignment specialist on \mathcal{D}_a and the complementarity specialist on \mathcal{D}_c . This yields a quantitative characterization of when adaptive combination of two models is fundamentally better than one model.

Let $L_a(\theta)$ and $L_c(\theta)$ denote the losses restricted to the alignment and complementarity regions, respectively, with $\theta_{m_a}^*$ and $\theta_{m_c}^*$ being their corresponding minimizers. Define distance between specialist models as $D := \|\theta_{m_a}^* - \theta_{m_c}^*\|$, and human accuracy factor $\kappa := 2\alpha - 1$.

Theorem 4 (Adaptive AI Performance Gain). *Assume the underlying prediction losses in the alignment and complementarity regions are μ -strongly convex in θ . Let $g_w(\theta) :=$*

$wL_a(\theta) + (1-w)L_c(\theta)$ denote a weighted single-model surrogate objective. In particular, for $w = p := \mathbb{P}[\mathbf{x} \in \mathcal{D}_a]$, let $L_{\text{single}}^ := \min_{\theta} g_p(\theta)$ be the loss of the best single model over the full population. The adaptive ensemble achieves $L_{\text{adapt}} = pL_a(\theta_{m_a}^*) + (1-p)L_c(\theta_{m_c}^*)$. Then, the advantage of the adaptive ensemble can be lower bounded as:*

$$\Gamma_{\text{team}} := L_{\text{single}}^* - L_{\text{adapt}} \geq \frac{\kappa \cdot \mu \cdot p(1-p) \cdot D^2}{2}.$$

This shows how the benefit of using two models is not merely conceptual but can be precisely quantified in terms of human reliability, loss surface geometry, task mixture balance and specialist divergence. Specifically, the gap increases when: (i) humans are more reliable in the aligned region (κ large), (ii) the loss landscape is “well-conditioned” or sharply curved (μ large), (iii) the two regions occur with comparable frequency ($p(1-p)$ large), and (iv) the alignment and complementarity optima are far apart (D^2 large).

Applicability to the behavior-aware team objective. Although the behavior-aware objective in Eq. 3 is non-convex due to the multiplicative term $(1 - L_h(\mathcal{D}_a, m))L(\mathcal{D}_c, m)$, the insight and bound of Theorem 4 remain indicative for the behavior-aware objective under mild regularity. When human reliance $L_h(\mathcal{D}_a, m)$ varies smoothly with model confidence and increases with alignment performance, the objective locally behaves like a convex combination of two strongly convex region-specific objectives (up to bounded perturbations). In this regime, the optimization landscape near the specialist solutions retains strong-convexity curvature, and the adaptive–single gap continues to scale quadratically with specialist separation. This approximation is consistent with our empirical results (Fig. 3).

Navigating Real-World Uncertainties

In practice, it is difficult to tell whether a given instance belongs to the complementarity or alignment region. We now (i) show how to train specialists under imperfect region knowledge, and (ii) quantify how this uncertainty propagates to test-time routing and resulting adaptive AI gains.

Impact of uncertainty at training time. In real-world settings, humans’ self-confidence threshold τ required to evaluate whether an instance belongs to the complementarity or alignment region may not only be unknown to the AI model developer, but also vary across different DMs and across time. To generalize to more realistic scenarios, we move beyond static human self-confidence threshold τ assumption, and instead model it as being drawn per-decision from a known (or estimated) distribution f_T . With τ stochastic, we introduce the notion of *probabilistic region membership* and expected region loss, which transforms our optimization problems into instance-weighted empirical risk minimization objectives. In fact, we can neatly integrate this into our analysis by simply generalizing loss definitions. For example, the *expected* complementarity loss becomes:

$$L(\mathcal{D}_c, m) = \mathbb{E}_{\tau} \left[\frac{1}{|\mathcal{D}|} \sum_{\mathbf{x}_i \in \mathcal{D}} \mathbb{I}\{\mathcal{C}_i^h \leq \tau\} \ell(m(\mathbf{x}_i), y_i) \right] = \frac{1}{|\mathcal{D}|} \sum_{\mathbf{x}_i \in \mathcal{D}} \underbrace{\mathbb{E}[(\mathbf{x}_i, y_i) \in \mathcal{D}_c]}_{w_i^c} \ell(m(\mathbf{x}_i), y_i).$$

Proposition 5. *When human confidence thresholds are drawn from a known distribution with CDF $F_T(\cdot)$, the optimal instance weighting in the complementarity region is given by $w_i^c = 1 - F_T(\mathcal{C}_i^h)$.*

A symmetric case can be made for expected alignment loss, with $w_i^a = F_T(\mathcal{C}_i^h)$. This weighting scheme naturally extends to both adaptive and behavior-aware AI training.

In line with insights from prior work, this weighting is principled and flexible: in fact, assuming uniform τ yields $w_i^c = 1 - \mathcal{C}_i^h$, recovering previously effective heuristics as a special case. Additionally, if $F_T(\cdot)$ or individual confidence values are unknown, uniform distribution assumption and confidence-prediction models (learned from a small pilot) can provide robust proxies, enabling practical instance-weighted training. Finally, while confidence calibration does influence the attainable gains, incorporating confidence into the AI optimization process leads to behavior-aware models that substantially outperform non-behavior-aware baselines (Mahmood, Lu, and Yin 2024).

Impact of uncertainty at test time. At deployment, uncertainty about region membership can lead to misrouting, i.e., instances assigned to the wrong specialist. Let $\bar{\rho}$ be the misrouting probability (expected fraction of misrouted instances), and \mathcal{H} the average region entropy, measuring uncertainty. The expected adaptive team loss is

$$L_{\text{adaptive}} = (1 - \bar{\rho})[L(\mathcal{D}_a, m_a) + L(\mathcal{D}_c, m_c)] + \bar{\rho}[L(\mathcal{D}_a, m_c) + L(\mathcal{D}_c, m_a)]. \quad (9)$$

As we show in Appendix, the misrouting probability can actually be bounded by the entropy: $\bar{\rho} \leq \frac{\mathcal{H}}{2 \log 2}$. This, coupled with Theorem 4, helps us precisely quantify the impact of region uncertainty.

Corollary 6 (Adaptive AI Performance Gain Under Uncertainty). *Let Γ_{team} denote the performance gain of adaptive AI over the best single model, and define $p := \mathbb{P}[\mathbf{x} \in \mathcal{D}_a]$. Then, under uncertainty level \mathcal{H} ,*

$$\Gamma_{\text{team}} \geq \left(1 - \frac{\mathcal{H}}{2 \log 2}\right) \frac{\kappa \cdot \mu \cdot p(1-p) \cdot D^2}{2}, \quad (10)$$

This shows that the benefit of adaptive ensembles degrades gracefully as uncertainty increases, and even moderate certainty suffices for substantial gains. In summary, adaptive AI is robust to real-world uncertainty in both training and deployment, and principled weighting and entropy-based analyses enable reliable human-AI collaboration even when perfect information is unattainable.

Evaluation on Simulated Data

To rigorously evaluate our theoretical framework and substantiate the practical benefits of our proposed adaptive AI ensemble, we conducted controlled simulations on College Admissions data that we synthesize. This setting enables precise measurement of the complementarity-alignment tradeoff, which helps validate Theorem 2, and direct quantification of adaptive AI ensemble gains under varying conditions, which helps validate Theorem 4 and Corollary 6.

Experimental Setup

We mimic a committee making a binary admission decision ($\mathcal{Y} = \{+1, -1\}$) based on an applicant’s Grade Point Average (GPA) and standardized test score. Consistent with some real settings, we assume applicants come from two subpopulations that differ in how well these features ($\mathbf{x} = \{x_{\text{GPA}}, x_{\text{Score}}\}$) indicate academic potential:

1. *Privileged Group:* We model the applicants from this group as residing within the feature space of alignment region (\mathcal{D}_a), where human evaluators are confident and typically accurate. Due to access to better schools and test prep, an applicant’s test score is assumed to be the most predictive feature.
2. *Underprivileged Group:* We model these applicants as residing within the complementarity region (\mathcal{D}_c), where evaluators are less confident. Here, an applicant’s GPA, reflecting long-term performance in their specific context, is assumed to be more predictive.

Ground truth generation. For each instance, x_{GPA} and x_{Score} are sampled uniformly from $[0, 1]$. The instance is assigned to the alignment region \mathcal{D}_a with probability p (default $p = 0.5$). The ground truth label y is then generated based on the region-specific predictive feature:

$$\begin{aligned} \mathbf{x} \in \mathcal{D}_a : y &= \mathbb{I}[(0.5 + \delta)x_{\text{Score}} + (0.5 - \delta)x_{\text{GPA}} \geq 0.5] \\ \mathbf{x} \in \mathcal{D}_c : y &= \mathbb{I}[(0.5 + \delta)x_{\text{GPA}} + (0.5 - \delta)x_{\text{Score}} \geq 0.5] \end{aligned}$$

Here, $\delta \in [0, 0.5]$ (default $\delta = 0.25$) controls the feature importance and thus the divergence between the two regions’ optimal decision boundaries, and consequently the specialist model divergence D . This scenario provides an intuitive analog for the tension between alignment and complementarity, while offering precise control over key parameters.

Human DM behavior. We simulate human judgments $h(\mathbf{x})$ to be correct with probability α in \mathcal{D}_a (default $\alpha = 1$); human accuracy in \mathcal{D}_c does not affect the CGPR dynamics and is set to a random baseline value when needed. Human confidence $\mathcal{C}^h(\mathbf{x})$ is then drawn uniformly from ± 0.1 around that region’s accuracy. This couples confidence and accuracy in a noisy, realistic manner. Unless otherwise noted, group membership is treated as known for routing (i.e., oracle routing) to isolate the other theoretical factors.

AI training. We trained the following logistic regression-based AI models: (i) *Single AI*, trained to minimize the weighted objective $g_p(\theta)$ from Theorem 4, which collapses to standard AI trained for empirical risk minimization under our default settings; (ii) *Complementary AI*, trained on \mathcal{D}_c using ground-truth labels y and weighted loss $(1 - \mathcal{C}^h(\mathbf{x}))$; (iii) *Aligned AI*, trained on \mathcal{D}_a using human predictions $h(\mathbf{x})$ as pseudo-labels and weighted loss $\mathcal{C}^h(\mathbf{x})$; and (iv) *Adaptive AI* ensemble that uses the *pre-trained* specialists m_a and m_c and, dynamically selects the correct specialist based on the instance’s known group membership ($G \in \{\mathcal{D}_a, \mathcal{D}_c\}$).

Evaluating Complementarity-Alignment Tradeoff

To probe the fundamental limitation of static models, we empirically measure the complementarity-alignment tradeoff

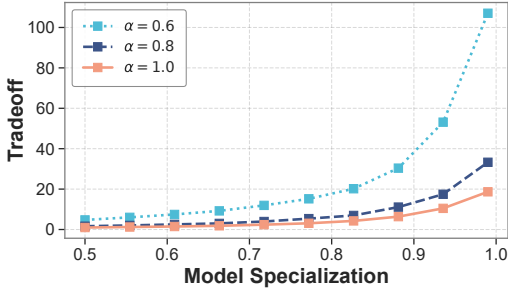


Figure 2: Complementarity-alignment tradeoff increases sharply with specialization, exacerbated by imperfect human DM in alignment region (α), illustrating that a single model cannot simultaneously optimize for trust and performance.

defined in Theorem 2. Specifically, we assess how a small perturbation in model parameters—in the direction of the alignment specialist—impacts the model’s loss in the complementarity region. Formally, for a single-model optimum θ_{single}^* , a perturbation direction v and small ϵ , the tradeoff ratio is measured as: $\text{Tradeoff} = \frac{\Delta L(\mathcal{D}_c, \theta_{\text{single}}^* + \epsilon v)}{-\Delta L(\mathcal{D}_a, \theta_{\text{single}}^* + \epsilon v)}$, where $\Delta L(\mathcal{D}, \theta)$ denotes the change in loss on region \mathcal{D} as the model parameters are nudged from θ_{single}^* .

To vary the degree of specialization, we increase the aligned group fraction p from 0.5 to 0.99, thereby moving the optimal single model closer to the aligned specialist. Figure 2 shows that the empirical tradeoff grows sharply as θ_{single}^* approaches the aligned specialist, and that the blow-up is faster when humans are less accurate in the aligned region (lower α), confirming the $(1/\kappa)$ scaling predicted by theory. These findings complement Theorem 2, and empirically establish that no single static model can simultaneously optimize alignment and complementarity when the regions are structurally distinct and human accuracy is imperfect.

Evaluating Adaptive AI Gains

We also empirically validate the adaptive AI performance gain predicted by our theoretical analysis (Theorem 4 and Corollary 6). We systematically vary each of the four key factors in our bound—specialist divergence (D), human reliability (κ), task mixture balance ($p(1-p)$), and group uncertainty (\mathcal{H})—and measure the resulting gain of the adaptive ensemble over the best single model. The results in Figure 3 confirm all four theoretical dependencies.

Impact of specialist divergence. Here, we vary $\delta \in [0, 0.5]$, which directly controls the geometric divergence between the aligned and complementary decision boundaries and hence the parameter distance D . Figure 3(i) shows that even moderate divergence yields meaningful improvements to adaptive ensemble’s accuracy gain, while large δ produces substantial adaptive advantages, confirming that *specialist heterogeneity is a primary driver of adaptive benefit*.

Impact of human accuracy. Next, we vary the human accuracy in the alignment region, $\alpha \in [0.5, 1]$, which modulates the reliability factor κ appearing in our theoretical bounds. As shown in Figure 3(ii), the adaptive advantage

increases monotonically with α . Higher human reliability strengthens the signal that the aligned specialist should follow, improving routing quality and magnifying the benefit of deferring to human expertise where appropriate. These results validate that *gains scale with human reliability*.

Impact of task-mixture balance. Then, we vary the mixture proportion $p = \mathbb{P}[\mathbf{x} \in \mathcal{D}_a]$, which controls the relative size of the two subpopulations and determines the $p(1-p)$ scaling in our adaptive–single gap bound. Figure 3(iii) confirms the predicted unimodal behavior: adaptive gains peak when the two regions are balanced ($p \approx 0.5$) and diminish when one region dominates the population. When one task is rare, a single model suffers little penalty from specializing toward the larger region, reducing the relative value of adaptivity. This supports the theoretical insight that *adaptive AI is most valuable when the environment contains multiple substantial and distinct regimes*.

Impact of region certainty. Finally, we vary the correctness of group labels, which determines how accurately the ensemble can assign each instance to the relevant specialist. Figure 3(iv) illustrates that while perfect group information maximizes the adaptive AI advantage, substantial gains persist even when group assignments are noisy. This aligns with our theoretical observation that the benefit degrades smoothly with misrouting, and the *ensemble remains robust as long as group identification is moderately reliable*.

Evaluation on Real-World Data

The controlled simulations above allowed us to isolate the structural factors predicted by our theory. To complement those findings, we examine whether the same qualitative patterns appear in a more complex perceptual domain where human competence, confidence, and uncertainty are empirically grounded. Instead of full-fledged real-world deployment, our goal is to assess whether the core mechanisms behind adaptive teaming continue to hold under realistic human variability. To this end, we evaluate on the WoofNette benchmark (Mahmood, Lu, and Yin 2024), which was explicitly built to study human–AI collaboration under heterogeneous and instance-specific uncertainty.

Task and dataset. WoofNette is a 10-class classification benchmark built from ImageNet, containing five everyday objects (e.g., Gas Pump, Parachute) and five dog breeds (e.g., Australian Terrier, Dingo), with 9,446 training and 4,054 test images. This structure naturally produces two regions: objects, where humans tend to be accurate and confident (approximate alignment region), and dog breeds, where human accuracy drops and confusion rises (approximate complementarity region). Importantly, instance-level variation within each class introduces additional uncertainty, making this a strong testbed for our adaptive framework.

Human DM behavior. Human predictions and self-reported confidence are estimated from a pilot study of 500 images and 4,644 annotations. We use inter-annotator agreement as a proxy for human confidence and train a lightweight ResNet-152 to generalize these confidence estimates to the full dataset, yielding $\hat{C}^h(\mathbf{x})$. Human predictions

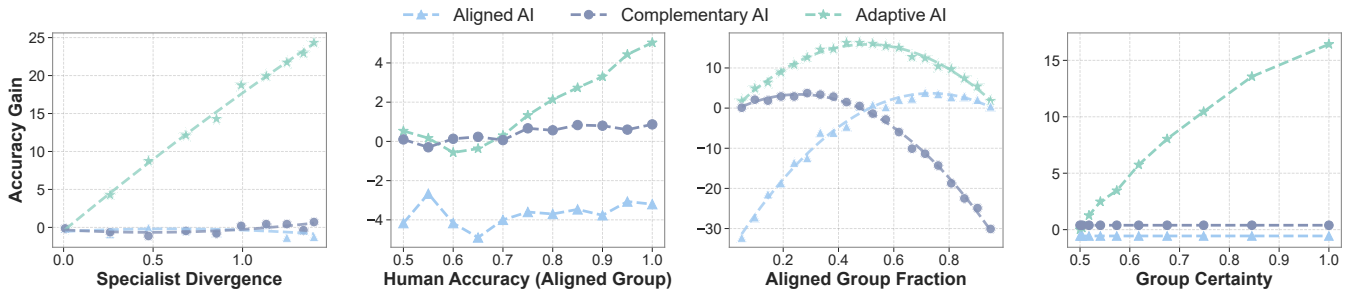


Figure 3: Accuracy gain (Γ_{team}) of adaptive AI ensemble over single AI on the College Admissions data. The plots (left to right) empirically validate that: (i) gain increases with specialist divergence ($\|\theta_{m_a}^* - \theta_{m_c}^*\|$); (ii) gain scales with human accuracy (α) in alignment region (i.e., the aligned group’s feature regime), reflecting the reliability factor κ ; (iii) gain peaks when the task mixture is balanced (i.e., aligned group fraction $p \approx 0.5$), matching the concave $p(1 - p)$ dependence; and (iv) gain increases linearly with group certainty ($1 - \mathcal{H}/2 \log 2$), demonstrating robustness to routing uncertainty.

Paradigm	AI Accuracy	Team Accuracy
Standard AI	69.87 \pm 0.44	69.13 \pm 0.28
Aligned AI	61.71 \pm 0.56	60.73 \pm 0.24
Complementary AI	61.01 \pm 0.77	69.96 \pm 0.50
Behavior-Aware AI	64.99 \pm 0.97	70.90 \pm 0.36
Adaptive AI (Oracle)	80.37 \pm 0.31	74.75 \pm 0.34
Adaptive AI (RRS)	82.64 \pm 0.35	75.13 \pm 0.32

Table 1: Accuracy (%) of independent AI and Human-AI team across different AI paradigms on the WoofNette benchmark. Human accuracy is 65.10 \pm 0.27.

are simulated using empirical confusion matrices. Because region boundaries are not directly observable in natural perceptual tasks, each human’s decision threshold τ is drawn from $\mathcal{U}[0.6, 0.8]$, modeling individual variability and the bimodal confidence distribution observed in pilot data.

AI training. All AI models share a ResNet-50 backbone, initialized from ImageNet. To preserve room for human contribution, we cap training early when AI reaches human-level accuracy ($\approx 65\%$). We train the following AI paradigms: (i) *Standard AI*, trained on all data, without regard for human strengths or weaknesses; (ii) *Aligned / Complementary AI*, trained with instance-weighted losses using $\hat{\mathcal{C}}^h(\mathbf{x})$ to reflect the specialist regions; (iii) *Behavior-Aware AI*, trained to optimize team performance under the CGPR model (Eq. 4); and (iv) *Adaptive AI (Oracle, RRS)*, which combines previously trained specialists using either the oracle confidence labels (Adaptive (Oracle); Eq. 7) or the Rational Routing Shortcut (Adaptive (RRS); Eq. 8).

Evaluation and Results

For each test image, we simulate a team decision by drawing a threshold τ , using the CGPR model to determine if the human defers to the AI, and recording the team accuracy over the entire test set (see Table 1).

Adaptive AI ensemble outperforms baselines. The adaptive AI ensemble consistently achieves the highest team

accuracy, outperforming both standard and behavior-aware baselines by clear margins. Notably, both Adaptive (Oracle) and Adaptive (RRS) achieve substantial improvements, with RRS variant performing virtually as well as the Oracle, despite having no access to human confidence at test time.

Gains realized despite weaker base AI models. The specialists used in the adaptive ensemble have lower individual accuracy than the standard AI, yet the adaptive team significantly outperforms the standard model. This reinforces the central insight of our theory and synthetic experiments, and underscores the value of intelligent combination and context-aware delegation over raw classifier performance.

Robustness to real-world constraints. This behavior-grounded experiment highlights that the adaptive AI framework, including its practical RRS instantiation, is robust to the uncertainties of a real-world task. Teaming gains persist even when region boundaries are instance-specific, human confidence is noisy, and AI predictions are imperfect.

Conclusion

This work advances the science of human-AI collaboration by introducing a principled and practical framework for adaptive teaming—one that is deeply informed by human behavioral realities, rigorous theory, and robust empirical validation. We formalize the complementarity-alignment tradeoff and show that the tension between performance-boosting and trust-building is an inherent limitation of single-model approaches, and provide the first tight theoretical characterization of this tradeoff. Our adaptive ensemble paradigm, powered by Confidence-Gated Probabilistic Reliance (CGPR) team decision making model and Rational Routing Shortcut (RRS)-based specialist selection, consistently achieves highest team performance in both synthetic and real-world experiments. Crucially, our results demonstrate that trustworthy, high-performing human-AI teams are possible only when AI systems dynamically align with human strengths and complement them where it matters most. We believe these advances contribute to the design of reliable, human-centered AI, enabling future systems that not only perform, but truly partner with, their human users.

Acknowledgments

We thank the support of the National Science Foundation under grant IIS-2229876 and IIS-2340209 on this work. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors alone.

References

- Agarwal, N.; Moehring, A.; Rajpurkar, P.; and Salz, T. 2023. Combining human expertise with artificial intelligence: Experimental evidence from radiology. Technical report, National Bureau of Economic Research.
- Askell, A.; Bai, Y.; Chen, A.; Drain, D.; Ganguli, D.; Henighan, T.; Jones, A.; Joseph, N.; Mann, B.; DasSarma, N.; et al. 2021. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.
- Bansal, G.; Nushi, B.; Kamar, E.; Horvitz, E.; and Weld, D. S. 2021. Is the most accurate ai the best teammate? optimizing ai for teamwork. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 11405–11414.
- Bondi, E.; Koster, R.; Sheahan, H.; Chadwick, M.; Bachrach, Y.; Cemgil, T.; Paquet, U.; and Dvijotham, K. 2022. Role of human-AI interaction in selective prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 5286–5294.
- Chiang, C.-W.; and Yin, M. 2021. You’d better stop! Understanding human reliance on machine learning models under covariate shift. In *Proceedings of the 13th ACM Web Science Conference 2021*, 120–129.
- Chong, L.; Zhang, G.; Goucher-Lambert, K.; Kotovsky, K.; and Cagan, J. 2022. Human confidence in artificial intelligence and in themselves: The evolution and impact of confidence on adoption of AI advice. *Computers in Human Behavior*, 127: 107018.
- Dell’Acqua, F.; Kogut, B.; and Perkowski, P. 2023. Super mario meets ai: Experimental effects of automation and skills on team performance and coordination. *Review of Economics and Statistics*, 1–47.
- 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.
- Goh, E.; Gallo, R.; Hom, J.; Strong, E.; Weng, Y.; Kerman, H.; Cool, J. A.; Kanjee, Z.; Parsons, A. S.; Ahuja, N.; et al. 2024. Large language model influence on diagnostic reasoning: a randomized clinical trial. *JAMA Network Open*, 7(10): e2440969–e2440969.
- Grgić-Hlača, N.; Castelluccia, C.; and Gummadi, K. P. 2022. Taking advice from (dis) similar machines: the impact of human-machine similarity on machine-assisted decision-making. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 10, 74–88.
- Guo, Y.; Cui, G.; Yuan, L.; Ding, N.; Sun, Z.; Sun, B.; Chen, H.; Xie, R.; Zhou, J.; Lin, Y.; et al. 2024. Controllable Preference Optimization: Toward Controllable Multi-Objective Alignment. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 1437–1454.
- Holstein, K.; and Aleven, V. 2021. Designing for human-AI complementarity in K-12 education. *arXiv preprint arXiv:2104.01266*.
- Lai, V.; Carton, S.; Bhatnagar, R.; Liao, Q. V.; Zhang, Y.; and Tan, C. 2022. Human-ai collaboration via conditional delegation: A case study of content moderation. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–18.
- Lu, Z.; Wang, D.; and Yin, M. 2024. Does more advice help? The effects of second opinions in AI-assisted decision making. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW1): 1–31.
- Lu, Z.; and Yin, M. 2021. Human Reliance on Machine Learning Models When Performance Feedback is Limited: Heuristics and Risks. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–16.
- Madras, D.; Pitassi, T.; and Zemel, R. 2018. Predict responsibly: improving fairness and accuracy by learning to defer. *Advances in Neural Information Processing Systems*, 31.
- Mahmood, S. H. A.; Lu, Z.; and Yin, M. 2024. Designing Behavior-Aware AI to Improve the Human-AI Team Performance in AI-Assisted Decision Making. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 3106–3114.
- Narayanan, S.; Yu, G.; Ho, C.-J.; and Yin, M. 2023. How does value similarity affect human reliance in AI-assisted ethical decision making? In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 49–57.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- 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*, 410–422.
- Steyvers, M.; Tejada, H.; Kerrigan, G.; and Smyth, P. 2022. Bayesian modeling of human-AI complementarity. *Proceedings of the National Academy of Sciences*, 119(11): e2111547119.
- 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*, 1–11.
- Wang, X.; Lu, Z.; and Yin, M. 2022. Will You Accept the AI Recommendation? Predicting Human Behavior in AI-Assisted Decision Making. In *Proceedings of the ACM Web Conference 2022*, 1697–1708.
- Wilder, B.; Horvitz, E.; and Kamar, E. 2020. Learning to complement humans. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, 1526–1533.

Zhang, Q.; Lee, M. L.; and Carter, S. 2022. You Complete Me: Human-AI Teams and Complementary Expertise. In *CHI Conference on Human Factors in Computing Systems*, 1–28.