

# Collective Intelligence in Human-AI Teams: A Bayesian Theory of Mind Approach

Samuel Westby<sup>1</sup>, Christoph Riedl<sup>2</sup>,

<sup>1</sup> Network Science Institute, Northeastern University, Boston, MA

<sup>2</sup> Khoury College of Computer Sciences, Northeastern University, Boston, MA  
westby.s@northeastern.edu, c.riedl@northeastern.edu

## Abstract

We develop a network of Bayesian agents that collectively model the mental states of teammates from the observed communication. Using a generative computational approach to cognition, we make two contributions. First, we show that our agent could generate interventions that improve the collective intelligence of a human-AI team beyond what humans alone would achieve. Second, we develop a real-time measure of human’s theory of mind ability and test theories about human cognition. We use data collected from an online experiment in which 145 individuals in 29 human-only teams of five communicate through a chat-based system to solve a cognitive task. We find that humans (a) struggle to fully integrate information from teammates into their decisions, especially when communication load is high, and (b) have cognitive biases which lead them to underweight certain useful, but ambiguous, information. Our theory of mind ability measure predicts both individual- and team-level performance. Observing teams’ first 25% of messages explains about 8% of the variation in final team performance, a 170% improvement compared to the current state of the art.

## Introduction

The reliance on teamwork in organizations (Wuchty, Jones, and Uzzi 2007), coupled with remarkable recent progress in artificial intelligence, have supercharged the vision to develop collaborative Human-AI teams (Malone and Bernstein 2015; O’Neill et al. 2020). Human-AI teams promise to overcome human biases and information processing limitations, reaching performance higher than human-only teams could (Brynjolfsson, Rock, and Syverson 2018). Despite some recent advances (e.g., Bansal et al. 2019b; Pynadath et al. 2022; Seraj et al. 2022) there remain significant difficulties in developing agents that interact with multiple, heterogeneous humans working on cognitive tasks engaged in cooperative communication in an ad-hoc team. Here, we draw on research of cognitive processes to develop Human-AI teams and explain collaborative decision making.

To communicate efficiently, humans infer the beliefs, opinions, knowledge, and related states of mind of other people (Nickerson 1999; Call and Tomasello 2008). This is referred to as social perceptiveness or theory of mind (ToM;

Premack and Woodruff 1978). Recent research on collective intelligence, has provided a wide range of empirical evidence suggesting that ToM (and related processes governing collective memory, attention, and reasoning) is a significant predictor of human collective intelligence (Woolley et al. 2010; Riedl et al. 2021; Woolley et al. 2022; Engel et al. 2014). Indeed, ToM is especially beneficial for interdependent cognitive tasks that benefit when teams leverage their members’ expertise (Lewis 2003). Some work suggests that ToM is the mechanism that allows collectives to use more performance-relevant information from their environment than a single individual without such connections could, for example, by facilitating a balance between diversity and cognitive efficiency (Riedl and Woolley 2017; Hong and Page 2004). As human civilization shifts further toward knowledge work (Autor 2014) where the most value is realized if members fully use and integrate their unique expertise, this ability is increasingly important.

Recent work has started to develop a formal account of collective intelligence to explain the relationship between individual interaction and collective performance using (approximate or variational) Bayesian inference (or free energy; Friston 2010, 2013; Heins et al. 2022). The free energy principle is a mathematical framework for multiscale behavioral processes that suggests a system of self-similar agents self-organizes by minimizing variational free energy in its exchanges with the environment (Fig. 1a; Friston, Kilner, and Harrison 2006). Recent extensions have applied the framework to explain human communication (Vasil et al. 2020). A key advantage of this approach is that free energy minimization can be translated into a generative, agent-based process theory (Friston et al. 2017; Kaufmann, Gupta, and Taylor 2021). This generative theory provides a computational approach to cognition (Tenenbaum et al. 2011; Griffiths 2015) that allows us to simultaneously (a) build agents for Human-AI teams that are highly explainable but also (b) test theories about human cognitive processes and measure human theory of mind *ability* in real time. This promises to advance our understanding of a key process of human collective intelligence. The current state of the art to measure theory of mind—the Reading the Mind in the Eyes test (Baron-Cohen et al. 2001)—is a static, indirect, survey-based instrument, which typically explains about 3% of the variation (Riedl et al. 2021).

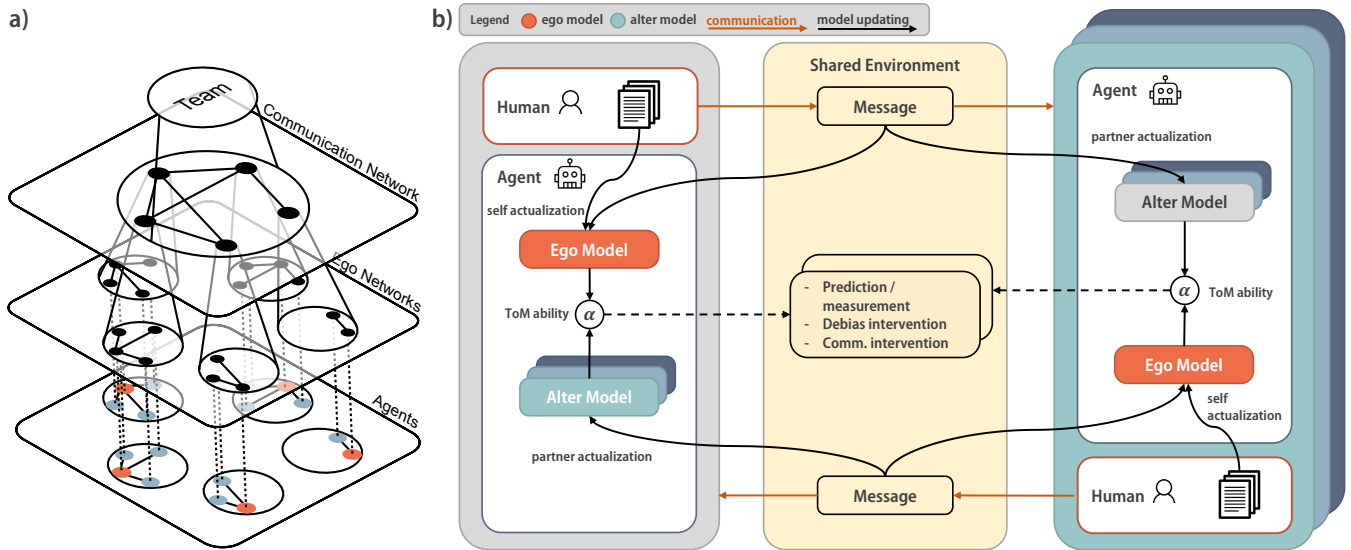


Figure 1: *Framework of human-AI teaming with Theory of Mind (ToM)*. a) Nested layers of ToM agents. Agents model ego networks of the individual they shadow. b) Every human team member is paired with an AI agent. Humans send messages to others through a shared environment. The ToM agent infers beliefs for both own ideas (*Ego Model*), and ideas of others (one *Alter Model* per network neighbor). *Ego Model* is updated with initial information and new knowledge generated by the human through self actualization. *Alter Models* are updated based on incoming messages from teammates through partner actualization. Agents combine information from the ego and alter models with weighting determined  $\alpha$  denoting ToM ability.

In this paper, we develop a Bayesian theory of mind agent that can form ad-hoc mental models about its teammates, based exclusively on observations drawn from human communication (Fig. 1b). We use data collected from a large, IRB approved human-subject experiment in which 145 individuals in 29 teams of five, randomly assigned to different social networks controlling the team’s communication topology, communicate through a chat-based system to solve a Hidden Profile task (Stasser and Titus 1985). We then simulate artificial teams in which each human is shadowed by an agent. The agent observes the same incoming and outgoing messages as the human did in the experiment. Modeling human behavior with our generative AI model allows us to test whether people do indeed form mental models of what their teammates know, how effectively they do so, and whether this ability is predictive of team performance. In a last step, we perform a counterfactual simulation to demonstrate how our Bayesian agent could trigger effective interventions that would increase Human-AI team performance over the observed human-only teams.

Our work provides a framework that expands theory of mind (and collective intelligence more broadly) from a static construct to a dynamical one that may vary according to situational factors, for example, due to changes in arousal, anxiety, and motivation with dynamically changing task requirements, time pressure, and recognition (Qin et al. 2022; Baliatti and Riedl 2021). We contribute to a body of research that has so far mostly used toy models—often using only a single agent—with an application to real data from multi-human communication (Vasil et al. 2020; Kaufmann, Gupta, and Taylor 2021; Albarracin et al. 2022; Heins et al. 2022).

Our work generates important cognitive insights into how humans communicate and reason to uncover hidden profiles. In summary, we make four main contributions.

1. We develop a networked Bayesian agent that models beliefs using human-human communication. We apply the agent to data from a human team experiment and demonstrate how the agent can monitor theory of mind in real time, predict both correct and human answers, and intervene to raise human performance.
2. We find the model accurately captures the decisions made by humans, varying in predictable ways with experimental stimuli like network position and task difficulty. Comparing model fits with simpler “lesioned” ToM models shows the value contributed by each component.
3. We develop two real-time measures for human theory of mind *ability*. The first, based on observed human communication and decisions, explains 51% variation in final team performance. The second, based on communication alone, explains 8% variation in *final* team performance after observing just the first quarter of communication, a 170% improvement compared to the current state of the art, the Reading the Mind in the Eyes test. Simulations of artificial human-AI teams suggest a significant 4% performance increase from AI triggered interventions.
4. We contribute to cognitive theory by presenting empirical evidence that cognitive biases explain the shortfall of human performance, such as a tendency to under-weight ambiguous information and failure to fully integrate information provided by others. We explain temporal patterns showing that high functioning teams send the most

useful information early before converging on common beliefs.

## Related Work

**Human-Agent teaming.** A long history of Human-Agent teaming has evolved alongside technological developments. Early examples such as Vannevar Bush’s Memex (Bush et al. 1945) demonstrate a longstanding fascination with augmenting human performance. Recently, work has specialized in many sub-fields including understanding mental models and related constructs like team situational awareness (Chen and Barnes 2014; Converse, Cannon-Bowers, and Salas 1993; Glikson and Woolley 2020). For example, effective Human-AI interaction has been shown to rely critically on the ability to form mental models about what AI teammates are doing (Bansal et al. 2019b; Paleja et al. 2021; Bansal et al. 2019a; Gero et al. 2020; Alipour et al. 2021).

Significantly less work has focused on how AI can form “mental models” of humans. Fügener et al. (2022) highlight this disparity by identifying situations where humans having mental models of the AI are not helpful, while AI having “mental models” of humans are. Given the challenges of designing multi-agent systems, human-AI teaming work has often focused on studying pairs of one agent and one human (e.g., Bansal et al. 2019a,b; Baker, Saxe, and Tenenbaum 2011; Fügener et al. 2022; Alipour et al. 2021). Furthermore, past work has often side-stepped challenges posed by language-based communication by constraining the scope to spatial or highly stylized tasks (Kaufmann, Gupta, and Taylor 2021; Baker et al. 2017; Khalvati et al. 2019). Others use Wizard of Oz techniques (Schelble et al. 2022; Hohenstein et al. 2022) to facilitate communication-based human-AI teaming interaction. To build an autonomous AI teammate that improves human-only team performance, one must build agents that overcome these obstacles. This becomes more challenging in a team of unfamiliar teammates, without a priori knowledge, while learning dynamically from language-based communication (Stone et al. 2010).

**Multi-agent Bayesian models** Multi-agent Bayesian models have been used to study coordination (Khalvati et al. 2019; Wu et al. 2021), opinion dynamics (Albarracin et al. 2022), efficient information fusion (Pavlin et al. 2010), and theory of mind (Baker et al. 2017). This can be modeled as a partially observable Markov decision process (POMDP) for each agent where the states are the set of other agents’ beliefs and observations are dependent on other agents’ actions (Smith, Friston, and Whyte 2022).

## Hidden Profile & Human Subject Data

A primary advantage of teams over lone individuals when solving complex problems is their ability to expand the pool of available information, thereby enabling teams to reach higher quality solutions (Mesmer-Magnus and DeChurch 2009). The Hidden Profile task (Stasser and Titus 1985) is a research task designed to mimic this decision-making scenario in which individuals hold private knowledge (Stone et al. 2010). In the task, some information is commonly held among all team members while each individual is also

endowed with unique private information. Subjects do not know what information is shared or private. Information sharing is the central process through which teammates collectively solve the task (Mesmer-Magnus and DeChurch 2009); conversely, failing to share all available information causes them to come to incorrect conclusions. Despite the importance of information sharing for team performance, past research has shown teams often deviate from the optimal use of information (Stasser and Titus 1985). Discussions tend to reinforce information held in common, rather than share information held uniquely by one team member (Nickerson 1999). One reason for this is that individuals impute their own knowledge on others and hence assume that private information is already shared (Nickerson 1999).

This gives rise to the “hidden profile” and directly points to avenues in which AI may improve team performance: identifying which information is uniquely held by each teammate and encouraging them to share it. Specifically, an agent may detect if their own mental model diverges from the inferred mental model of another (“I know something that I believe you don’t know”) indicating a window of opportunity for an effective intervention. It also provides the basis for our measure of theory of mind *ability*. Individuals who form more precise mental models of their teammates (and who impute less of their own knowledge on others) will be more efficient communicators who share more useful information in a more targeted manner.

We use data from an IRB approved online experiment conducted on the Volunteer Science platform (Radford et al. 2016) in which 145 individuals in 29 human-only teams of five solved a Hidden Profile task (data and code available at <https://github.com/riedlc/HumanAITeamsAndCI>). The task is framed as a crime investigation: the team needs to pool clues to answer questions about the target, culprit, and time of an art heist. There are six clues for each question. When combined, the correct answer out of five possible options is obvious. One randomly selected clue is given to every teammate (the public clue) and each individual receives one of the remaining five clues (the private clue, also randomly selected). Teams were randomly assigned to a communication topology using all 21 possible connected five node graphs (e.g., star, ring, chain; see Fig. 1a for one example). Teams then communicate via text-based chat where each message is sent to all neighboring teammates. After a five minute discussion phase, each individual submits answers for the culprit’s identity, the target, and the day of the art heist. Subjects were recruited from Amazon Mechanical Turk (Paolacci, Chandler, and Ipeirotis 2010) and paid a \$0.75 flat fee for participation as well as a \$0.25 performance bonus for each correct answer. The entire task took about seven minutes to complete. Subjects are blind to the network topology they have been assigned to as well as the total number of individuals on their team. For simplicity of exposition and analysis, we rely only on the *culprit* dimension of the task. We compute individual performance as  $\{0, 1\}$  depending on whether the culprit guess is correct and team performance as the average individual performance (majority voting does not substantively change results). On average, individuals received 3.1 (SD 1.9) chat messages from each partner.

To make initial clues and communication machine interpretable, we manually code the content as *strong no* (SN), *maybe no* (MN), *maybe yes* (MY), and *strong yes* (SY) for each of the five answer options (the inferred states). This creates a set of 20 possible observations. We translate messages into likelihoods for the inferred states using fixed values, either estimated from the data using maximum likelihood estimation (MLE) from a grid search or using untrained intuitive values. For example, the message, “it might be #4”, would be coded as *maybe yes* with a likelihood of 1.4 for #3, leaving the likelihoods for the states not mentioned in the message unaffected. Ambiguous statements and messages related to team coordination were coded as “neutral” and dropped from the analysis. Notice that agents form beliefs solely based on the observed human communication, even if humans make certain statements about wrong facts (e.g., “strong yes #4” when the correct answer is #3), or ambiguous statements about correct facts. Agents can thus form wrong beliefs (Albarracín et al. 2022).

### Bayesian Multi-Agent Model

We create a networked Bayesian agent for each individual. Each agent “shadows” one human, observing the same messages (both are inside the same Markov blanket; Fig. 1a) and infers human beliefs and derive the answer to the Hidden Profile task. That is, our Bayesian system is a model of beliefs about hidden states of their environment. Ideally, the state inferred after the communication phase is identical to the correct answer to the Hidden Profile task. The resulting model has five parameters: four information weights *SN*, *MN*, *MY*, *SY* determining the likelihood distribution of observations under inferred beliefs, and the theory of mind ability  $\alpha_D$  which modulates the relative weighting of the self vs. partner beliefs which we describe in more detail below.

**Mental Models.** We use Bayesian inference to infer a posterior distribution  $p(\mathbf{s} | \mathbf{o})$  over states  $\mathbf{s}$  (five answer options), from a set of observations  $\mathbf{o}$  (messages sent between players in the recorded chat). Since there are five discrete states, we can compute posteriors directly without the need to approximate them. More complicated environments may require the use of approximate inference methods like free energy/active inference (Friston, Kilner, and Harrison 2006).

Agents are comprised of one *Ego Model* and one *Alter Model* for each neighbor (Fig. 1b). That is, the model follows a multi-agent paradigm with independent mental models nested within an agent. All models hold a posterior distribution of inferred states, but differ in how they are initialized and updated. *Ego Models* are initialized with priors derived from the public and private clues assigned to the player (paper stack icons in Fig. 1b) and updated with *outgoing* messages from a player (*self actualization*). *Alter Models* are initialized with uniform priors and updated with *incoming* messages of the corresponding (*partner actualization*).

Mental models are updated by accounting for the surprise of the observation. Surprise-weighting thus encodes a preference for posteriors with low surprise. That is, the effect of new observations is diminished relative to the *Ego* or *Alter* model’s existing posterior using

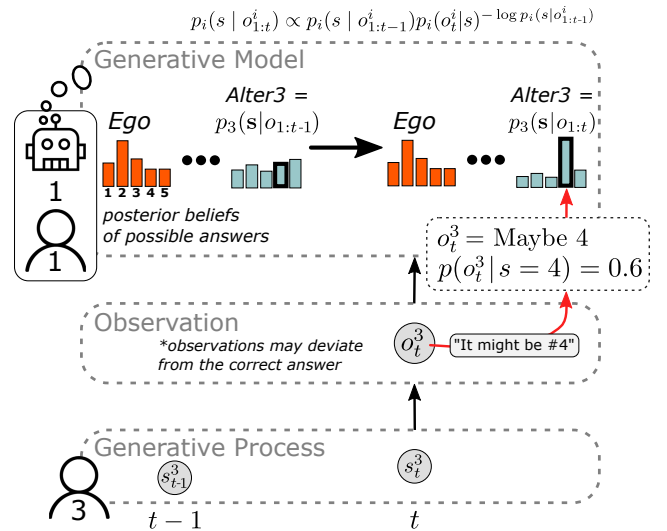


Figure 2: *Bayesian model of Theory of Mind*. Agent 1, shadowing Player 1, models all teammates in Player 1’s ego network (Markov blanket). In Agent 1’s generative model, *Alter3* corresponds to Agent 1’s beliefs of Player 3’s beliefs. At time  $t$  Player 3 says, “It might be #4”, which is coded as a *MY* for answer 4. Given this new observation, Agent 1 uses Equation 1 to update its beliefs about Player 3’s beliefs.

$$p_i(s | o_{1:t}^i) \propto p_i(s | o_{1:t-1}^i) p_i(o_t^i | s)^{\overbrace{-\log p_i(s | o_{1:t-1}^i)}^{\text{surprise}}} \quad (1)$$

where  $s$  is a state and  $o_t^i$  is an observation (message) sent by player  $i$  at time  $t$ . The likelihood is raised to the negative log of the previous time step’s posterior.

**Agent.** The agent is the hierarchical coordinator and aggregator of its mental models. The ToM ability parameter  $\alpha_D$  modulates the relative weight with which the agent combines its *Ego* and *Alter* models. We conceptualize  $\alpha_D$  as the ability to accurately infer beliefs of other agents and pay attention to them (Apperly and Butterfill 2009). It represents the relative weighting between the agent’s own *Ego Model* and its *Alter Models*. When  $\alpha_D = 0$ , the *Alter* posterior is uniform and has no effect and the final prediction is based only on the *Ego Model*. When  $\alpha_D = 1$ , the *Alter Models* are weighted equally to the *Ego Model*. Agent  $i$  aggregates its mental models into a final posterior distribution using

$$p_i(s | \mathbf{M}_i) \propto p_i(s | \mathbf{o}^i) \prod_{\substack{m \in \mathbf{M}_i \\ m \neq i}} p_m(s | \mathbf{o}^m)^{\alpha_D} \quad (2)$$

where  $\mathbf{M}_i$  is Agent  $i$ ’s set of mental models and  $p_m(s | \mathbf{o}^m)$  is posterior of state  $s$  for the mental model of player  $m$  over  $\mathbf{o}^m$ , the set of  $i$ ’s observations of  $m$ .

**Alternative Models.** To test whether the full representational capacity of theory of mind with both *self-actualization* and *partner-actualization* loops are necessary to understand human mental states, we formulate two alternative models that “lesion” one or both of the updating loops. This allows

	Performance (% Correct)	Human-Agent Agreement	LogLik	Model Comparison (Likelihood Ratio test)	$\alpha_D$
<b>Human</b>	<b>66.2%</b>				
Random	19.6 $\pm$ 3.1%	20.0 $\pm$ 3.0%	-231.759		
Prior only	48.8 $\pm$ 3.1%	46.6 $\pm$ 2.9%	-215.906	vs. Random $p < 0.0001$	
ToM (self-actualization only, MLE)	63.6 $\pm$ 1.8%	73.0 $\pm$ 1.9%	-146.372	vs. Prior-only $p < 0.0001$	0
ToM (partner-actualization only, MLE)	76.7 $\pm$ 1.0%	66.3 $\pm$ 1.0%	-133.653	vs. Self-only $p < 0.0001$	1
ToM (MLE)	72.8 $\pm$ 0.8%	75.1 $\pm$ 1.0%	<b>-106.640</b>	vs. Partner-only $p < 0.0001$	0.95
ToM (max performance)	<b>77.2 <math>\pm</math> 0.8%</b>	71.3 $\pm$ 0.9%	-109.826	vs. MLE $p = 0.012$	0.95
ToM (max agreement)	71.8 $\pm$ 0.7%	<b>79.7 <math>\pm</math> 0.8%</b>	-118.464	vs. MLE $p < 0.0001$	0.45
With random intervention	79.0 $\pm$ 1.8%	70.8 $\pm$ 1.6%			0.95
With intervention	<b>82.1 <math>\pm</math> 0.7%</b>	70.0 $\pm$ 0.6%		vs. Rand. int. $p < 0.0001$	0.95

Table 1: *Model evaluation results and comparison with human behavior.* P-values based on likelihood ratio test. We calculate standard deviations over 100 trials. Information weights ( $SN, MN, SY, MY$ ) are learned from the data through a grid search: Prior (0.05, 0.05, 2, 2), self-act. (0.05, 0.05, 1.5, 2), partner-act. (0.15, 1, 1.55, 2), MLE (0.1, 1, 1.45, 2), max perf. (0.35, 0.85, 1.95, 2), max agg. (0.05, 0.75, 1.25, 1.95). Interventions use the same parameters as max performance and  $t$ -test for comparison.

us to test whether it is possible to explain human inferences about their teammates without appealing to a fully developed theory of mind. We compute  $p$ -values from likelihood ratio tests comparing the models.

## Results

**Model Evaluation.** We find strong support for the hypothesis that humans use Bayesian inference to model the minds of their teammates and communicate and make decisions according to those models (Table 1). Compared to a model using only prior information (the clues distributed in the experiment), a model capturing humans’ ability to update their own beliefs (self-actualization only) fits the data significantly better. A model allowing humans to update beliefs about their teammates (partner-actualization only) fits significantly better still. Finally, a model including the capability to update both own and partner beliefs has the highest fit. Higher values for  $\alpha_D$  generally lead to more peaked posterior distributions. This explains why the parameter values that produce the highest likelihood differ slightly from those of the highest accuracy ( $\alpha_D^{MLE} = 0.95$  vs.  $\alpha_D^{maxacc} = 0.45$ ). In summary, the comparative fit analysis provides reliable evidence for the mutually inferred alignment of attention (cf., mental states) among teammates.

Our model accurately captures the judgments of human participants, varying in predictable ways with random experimental manipulation of task difficulty and the number of communication partners. We measure the task difficulty faced by each individual based on how much information the individual can draw about the correct answer from the two clues they initially received. This captures how difficult it is for an individual to guess the correct answer before communicating with other players (this is a somewhat noisy measure as it ignores team-level effects of the clue distributions). Not surprisingly, human performance decreases with task difficulty, suggesting that humans suffer from cognitive overload (Fig. 3a). Our agent achieves high accuracy predicting human’s incorrect answers under high task difficulty

(high true-negative rate).

Human performance varies with the number of communication partners (Fig. 3b). Given the nature of the task, access to more communication partners should be beneficial as this guarantees access to more information. Humans, however, perform worse with more communication partners while our ToM agent achieves its highest performance when placed in the most central network position (agent is 20% better than human with four partners). This suggests that humans struggle to integrate information when communicating with many teammates. This picture becomes even clearer when contrasting this with ToM ability  $\alpha_D$  (Fig. 3c). Higher levels of ToM ability  $\alpha_D$  have the highest benefit on performance in central network positions, yet  $\alpha_D$  hardly matters when connected to just a single teammate.

**Analysis of Human Decision Biases.** The ToM model predicts with high accuracy instances in which humans provide the correct answer as well as those in which they provide the wrong answer (48% true-negative accuracy). Comparing the information weighting parameters for optimal performance with those for the highest model fit with data from the human subject experiment from the MLE estimates, we can directly see why human performance falls short. Humans pay not enough attention to information ruling out alternatives (optimal information weighting for *strong no* 0.25 vs. MLE fit 0.05). The difference is even more pronounced for ambiguous information (optimal information weighting for *maybe no* 0.9 vs. MLE fit 0.05): humans undervalue information that is ambiguous, yet crucial in arriving at the correct answer. Because this information is ambiguous, humans may attempt to make sense of it by imputing their own understanding (i.e., resorting to their own prior) instead of updating their beliefs in the direction of the ambiguous message. A similar weighting difference for *maybe yes* statements suggests that humans communicate *strong yes* information in vague ways (maybe-ing their statements) and could significantly improve their performance by placing higher weight on such statements (or communicat-






Figure 3: *Human performance varies with task difficulty and number of communication partners.* From left to right: a) Human performance decreases with task difficulty and is outperformed by AI agent in most cases. b) Agent improves over human performance especially when communicating with many teammates. c) Agents with many communication partners benefit most from high ToM ability  $\alpha_D$ . (parameters:  $\alpha_D = 0.95$ ,  $SN = 0.35$ ;  $MN = 0.85$ ;  $MY = 1.95$ ;  $SY = 2$ ).

ing them more forcefully).

**Measuring Theory of Mind.** We propose two measures of human theory of mind ability:  $\alpha_D$  and  $\alpha_C$ . The first,  $\alpha_D$ , is based on an individual’s ability to form and integrate accurate mental models of others when making *decisions* and corresponds directly to our model parameter that governs the relative weighting of the *Ego* vs. *Alter Models*. The second,  $\alpha_C$ , captures an individual’s ability to *communicate* the most useful information. We perform maximum likelihood estimate using a grid search over the relevant parameter space (Baliotti, Klein, and Riedl 2021). Then, we fix the maximum likelihood estimate of the nuisance parameters for information weighting ( $SN$ ,  $MN$ ,  $MY$ ,  $SY$ ) but consider the marginal of all values of  $\alpha_D$ . Instead of then picking the global best fitting value for the entire data set, we pick the maximum likelihood estimate of  $\alpha_D$  separately for each individual. That is, we use the model’s inner ToM working to estimate which value of individual  $i$ ’s  $\alpha_D$  produces the highest likelihood of the observed decision. For the second measure  $\alpha_C$ , we consider outgoing messages sent by each individual and compute the expected surprise that this message should produce for the recipient, relative to ego’s *Alter Model* of the recipient. Notice that we compute this internally within the Markov blanket of an agent. We do not use information about how surprising the message is for the recipient but rather how useful the sender thinks it should be relative to what they think the recipient knows. Intuitively, individuals who possess a high theory of mind ability, will be better at sending the right message to the right person compared to those with lower ToM ability. Both measures capture social perceptiveness: how much attention an individual pays to what others in the team know.

We find that individual-level ToM ability  $\alpha_D$  is a strong predictor of individual-level performance ( $\beta = 0.59$ ;  $p < 0.001$ ;  $R^2 = 0.26$ ). Aggregating to the team level, we find that average ToM ability  $\alpha_D^{team}$  is a strong predictor of final team performance (Fig. 4a). We find that the effect of ToM ability is moderated by average betweenness centrality suggesting team performance increases most when high-ToM ability  $\alpha_D$  individuals occupy high betweenness network positions ( $\beta = 0.39$ ;  $p = 0.04$ ). The amount of com-

munication sent within a team, notably, is *not* a significant predictor of team performance ( $\beta = -0.00$ ;  $p = 0.265$ ).

Turning to our analysis of theory of mind communication ability  $\alpha_C^{team}$ , we find that it is a strong predictor of team-level performance ( $\beta = 0.47$ ;  $p = 0.019$ ). Given that we can measure  $\alpha_C$  on the message level, it can serve as a real-time measure of theory of mind. We find that after observing only the *first 25%* of a team’s messages,  $\alpha_C^{team}$  is a significant predictor of *final* team performance (Fig. 4b). We analyze the temporal pattern in which high- vs. low-performing teams communicate (Fig. 4c). High-performing teams send messages with high information content (high surprise) early during the team task but then send consolidating, low information content messages at the end to facilitate convergence (low surprise). We illustrate this in the example below (\*’s indicate high surprise messages with novel content).

Human 1: *It will not happen on Tuesday\**  
 Human 3: *No Wednesday or Friday\**  
 Human 2: *Monday or Thursday\**  
 Human 2: *Did you get a no Thursday info?*  
 Human 5: *Yeah I got no Thursday\**  
 Human 3: *I got no Thursday*  
 Human 5: *So it must be Monday*

This suggests that team cognition is not static but instead emerges dynamically and that high-performing teams have the collective capacity to modulate shared cognition dynamically to achieve both efficient information transfer and convergence during different periods of the task. Low-performing teams on the other hand fail to send high information content messages and also fail to achieve convergence sending more surprising messages late during the task. This pattern illustrates that high information content alone is not desirable. Instead, convergence and joint attention to team consensus are crucial (Woolley et al. 2022).

The current standard to predict social perceptiveness is the Reading the Mind in the Eyes (RME) test (Baron-Cohen et al. 2001; Almaatouq et al. 2021). Using data from a large meta analysis (Riedl et al. 2021) of 5,279 individuals in 1,356 groups, we find RME explains between 0% and 3% of the variation in team performance (depending on the task). Our  $\alpha_C^{team}$  measure explains 8% after observing only 25% of




Figure 4: *Theory of Mind ability predicts team performance.* a) Team average ToM  $\alpha_D^{team}$  is a strong predictor of the final team performance. b) Communication ToM  $\alpha_C^{team}$  serves as a real-time measure of collective intelligence. Only about the first 25% of team messages are necessary to make significant predictions of *final* team performance. c) High- and low-performing teams have markedly different temporal patterns of ToM  $\alpha_C$ .

team communication, an improvement of about 170%. Our proposed measure captures social perceptiveness passively and in real time which can be used to interpret a team’s current status and determine opportunities for interventions. Furthermore, RME captures social perceptiveness of a single individual, while our measure is group based. Our work also extends previous measures of “information diversity” (Riedl and Woolley 2017). It thus captures aspects of collective attention and memory (Gupta and Woolley 2020).

**Human-Agent Team Performance.** So far, our AI agent was only passively shadowing its assigned human, reasoning about the mental states of that human and its connected teammates. In this section, we extend this reasoning to allow the agent to trigger interventions that could be deployed in human-AI teams and quantify what performance improvement this might yield. We perform a counterfactual simulation in which we allow each AI agent to identify and send one message to each network neighbor. Each agent compares its *Ego Model* against its *Alter Models* to identify divergence in inferred beliefs. Each agent then draws from the set of messages it received during the team discussion and chooses a message to send to each neighbor. To do this, the agent calculates the effect of sharing one of its available messages on the *Alter Models* and shares the message that results in the lowest KL divergence, defined as  $D_{KL}(Q \parallel P) = \sum_i Q(i) \ln \frac{Q(i)}{P(i)}$ , between the *Ego* and *Alter<sub>i</sub>* posteriors over all five possible answer options. If no message lowers the KL divergence, the agent shares no message. This is summarized as taking the action  $a_{ij}$  for each agent  $i$  and neighbor  $j$  where

$$a_{ij} = \arg \min_{m \in \mathbf{o}} D_{KL}(p_{ego_i}(\mathbf{s} \mid \mathbf{o}) \parallel p_{alter_j}(\mathbf{s} \mid \mathbf{o}', m)) \quad (3)$$

$m$  is selected from the set of messages  $\mathbf{o}$  that agent  $i$  sent or received,  $\mathbf{s}$  is a vector of the five possible answers, and  $\mathbf{o}'$  is the set of messages agent  $i$  received from agent  $j$ . To establish a baseline intervention, we let  $a_{ij}$  be a random message in  $\mathbf{o} \cup \{\text{no message}\}$ . Here, performance improves to  $79.0 \pm 1.8\%$  averaged over 100 trials. For the targeted intervention, performance improves 4.9% to  $82.1 \pm 0.7\%$  which is significantly higher ( $t$ -test  $p < 0.0001$ ) than the random intervention. Notice that this intervention would not be possible

without our ToM-based multi-agent model. Without it, we could not determine which message to send to which alter.

## Discussion

We develop a framework that combines theory of mind, Bayesian inference, and collective intelligence into a generative computational model. Our model accurately captures the decisions made by human participants in a cognitive decision-making experiment, varying in predictable ways with experimental manipulation of task difficulty and network position. Our results suggest that humans use Bayesian inference and Theory of Mind to model their own beliefs and those of their teammates and communicate and make decisions according to those models. We provide empirical evidence that humans do not do this perfectly but suffer from cognitive biases. Nonetheless, our Bayesian agent is robust and achieves high performance even when fed biased and incorrect information, providing a pathway to implement high-performing human-AI teams. Notably, our agent works in ad hoc teams with heterogeneous partners without any pretraining. In such human-AI teams, our AI could augment humans’ limited cognitive memory, attention, and reasoning abilities to increase collective intelligence.

We show empirical evidence that the collective dynamics of Bayesian agents updating probabilities of hypotheses using observations, collectively predict the performance at the team level. This provides the basis for a real-time measure of theory of mind ability, and maybe even collective intelligence more broadly (Heins et al. 2022). The better the mental models of the team members align—the less surprising observations drawn from communication become—the higher the team’s collective intelligence. Our implementation of direct surprise weighting could be extended with a fuller implementation of the free energy principle that would allow agents to learn asymmetric beliefs about the reliability of their partners’ signals. Taken together, this is a framework to capture the emergence of collective memory, attention, and reasoning in real time (Luria 1973; Gupta and Woolley 2020).

## Acknowledgements

This work was supported by the Army Research Laboratory [Grant W911NF-19-2-0135].

## References

- Albarracín, M.; Demekas, D.; Ramstead, M. J.; and Heins, C. 2022. Epistemic communities under active inference. *Entropy*, 24(4): 476.
- Alipour, K.; Ray, A.; Lin, X.; Cogswell, M.; Schulze, J. P.; Yao, Y.; and Burachas, G. T. 2021. Improving users' mental model with attention-directed counterfactual edits. *Applied AI Letters*, 2(4): e47.
- Almaatouq, A.; Alsobay, M.; Yin, M.; and Watts, D. J. 2021. Task complexity moderates group synergy. *Proceedings of the National Academy of Sciences*, 118(36): e2101062118.
- Apperly, I. A.; and Butterfill, S. A. 2009. Do humans have two systems to track beliefs and belief-like states? *Psychological Review*, 116(4): 953.
- Autor, D. H. 2014. Skills, education, and the rise of earnings inequality among the “other 99 percent”. *Science*, 344(6186): 843–851.
- Baker, C.; Saxe, R.; and Tenenbaum, J. 2011. Bayesian theory of mind: Modeling joint belief-desire attribution. In *Proceedings of the Thirty-Third Annual Meeting of the Cognitive Science Society*, 2469–2474.
- Baker, C. L.; Jara-Ettinger, J.; Saxe, R.; and Tenenbaum, J. B. 2017. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4): 1–10.
- Baliatti, S.; Klein, B.; and Riedl, C. 2021. Optimal design of experiments to identify latent behavioral types. *Experimental Economics*, 24(3): 772–799.
- Baliatti, S.; and Riedl, C. 2021. Incentives, competition, and inequality in markets for creative production. *Research Policy*, 50(4): 104212.
- Bansal, G.; Nushi, B.; Kamar, E.; Lasecki, W. S.; Weld, D. S.; and Horvitz, E. 2019a. Beyond accuracy: The role of mental models in human-AI team performance. In *Proceedings of the Seventh AAAI Conference on Human Computation and Crowdsourcing*, 2–11.
- Bansal, G.; Nushi, B.; Kamar, E.; Weld, D. S.; Lasecki, W. S.; and Horvitz, E. 2019b. Updates in human-AI teams: Understanding and addressing the performance/compatibility tradeoff. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, 2429–2437.
- Baron-Cohen, S.; Wheelwright, S.; Hill, J.; Raste, Y.; and Plumb, I. 2001. The “Reading the Mind in the Eyes” Test revised version: a study with normal adults, and adults with Asperger syndrome or high-functioning autism. *The Journal of Child Psychology and Psychiatry and Allied Disciplines*, 42(2): 241–251.
- Brynjolfsson, E.; Rock, D.; and Syverson, C. 2018. Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In *The Economics of Artificial Intelligence: An Agenda*, 23–57. Chicago, IL: University of Chicago Press.
- Bush, V.; et al. 1945. As we may think. *The Atlantic Monthly*, 176(1): 101–108.
- Call, J.; and Tomasello, M. 2008. Does the chimpanzee have a theory of mind? 30 years later. *Trends in Cognitive Sciences*, 12(5): 187–192.
- Chen, J. Y.; and Barnes, M. J. 2014. Human-agent teaming for multirobot control: A review of human factors issues. *IEEE Transactions on Human-Machine Systems*, 44(1): 13–29.
- Converse, S.; Cannon-Bowers, J.; and Salas, E. 1993. Shared mental models in expert team decision making. *Individual and group decision making: Current issues*, 221: 221–46.
- Engel, D.; Woolley, A. W.; Jing, L. X.; Chabris, C. F.; and Malone, T. W. 2014. Reading the mind in the eyes or reading between the lines? Theory of mind predicts collective intelligence equally well online and face-to-face. *PLoS One*, 9(12): e115212.
- Friston, K. 2010. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2): 127–138.
- Friston, K. 2013. Life as we know it. *Journal of the Royal Society Interface*, 10(86): 20130475.
- Friston, K.; FitzGerald, T.; Rigoli, F.; Schwartenbeck, P.; and Pezzulo, G. 2017. Active inference: a process theory. *Neural Computation*, 29(1): 1–49.
- Friston, K.; Kilner, J.; and Harrison, L. 2006. A free energy principle for the brain. *Journal of Physiology*, 100(1-3): 70–87.
- Fügener, A.; Grahl, J.; Gupta, A.; and Ketter, W. 2022. Cognitive challenges in human-artificial intelligence collaboration: investigating the path toward productive delegation. *Information Systems Research*, 33(2): 678–696.
- Gero, K. I.; Ashktorab, Z.; Dugan, C.; Pan, Q.; Johnson, J.; Geyer, W.; Ruiz, M.; Miller, S.; Millen, D. R.; Campbell, M.; et al. 2020. Mental models of AI agents in a cooperative game setting. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–12.
- Glikson, E.; and Woolley, A. W. 2020. Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2): 627–660.
- Griffiths, T. L. 2015. Manifesto for a new (computational) cognitive revolution. *Cognition*, 135: 21–23.
- Gupta, P.; and Woolley, A. W. 2020. The emergence of collective intelligence behavior. In *Proceedings of the Paper presented at the 8th ACM Collective Intelligence (CI) Conference, Virtual Event, Zurich, Switzerland*.
- Heins, C.; Klein, B.; Demekas, D.; Aguilera, M.; and Buckley, C. 2022. Spin glass systems as collective active inference. *arXiv preprint arXiv:2207.06970*.
- Hohenstein, J.; Larson, L. E.; Hou, Y. T.-Y.; Harris, A. M.; Schechter, A.; Dechurch, L.; Contractor, N.; and Jung, M. F. 2022. Vero: A Method for Remotely Studying Human-AI Collaboration. In *Proceedings of the 55th Hawaii International Conference on System Sciences*.



- Hong, L.; and Page, S. E. 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*, 101(46): 16385–16389.
- Kaufmann, R.; Gupta, P.; and Taylor, J. 2021. An active inference model of collective intelligence. *Entropy*, 23(7): 830.
- Khalvati, K.; Park, S. A.; Mirbagheri, S.; Philippe, R.; Sestito, M.; Dreher, J.-C.; and Rao, R. P. 2019. Modeling other minds: Bayesian inference explains human choices in group decision-making. *Science Advances*, 5(11): eaax8783.
- Lewis, K. 2003. Measuring transactive memory systems in the field: scale development and validation. *Journal of Applied Psychology*, 88(4): 587.
- Luria, A. R. A. R. 1973. *The working brain; an Introduction to Neuropsychology*. New York, NY: Basic Books.
- Malone, T. W.; and Bernstein, M. S. 2015. *Handbook of Collective Intelligence*. MIT Press.
- Mesmer-Magnus, J. R.; and DeChurch, L. A. 2009. Information sharing and team performance: a meta-analysis. *Journal of Applied Psychology*, 94(2): 535.
- Nickerson, R. S. 1999. How we know—and sometimes misjudge—what others know: Imputing one’s own knowledge to others. *Psychological Bulletin*, 125(6): 737.
- O’Neill, T.; McNeese, N.; Barron, A.; and Schelble, B. 2020. Human–autonomy teaming: A review and analysis of the empirical literature. *Human Factors*, 64: 904–938.
- Paleja, R.; Ghuy, M.; Ranawaka Arachchige, N.; Jensen, R.; and Gombolay, M. 2021. The utility of explainable ai in ad hoc human-machine teaming. *Advances in Neural Information Processing Systems (NeurIPS)*, 34: 610–623.
- Paolacci, G.; Chandler, J.; and Ipeirotis, P. G. 2010. Running experiments on amazon mechanical turk. *Judgment and Decision making*, 5(5): 411–419.
- Pavlin, G.; de Oude, P.; Maris, M.; Nunnink, J.; and Hood, T. 2010. A multi-agent systems approach to distributed bayesian information fusion. *Information Fusion*, 11(3): 267–282.
- Premack, D.; and Woodruff, G. 1978. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(4): 515–526.
- Pynadath, D. V.; Dilkina, B.; Jeong, D. C.; John, R. S.; Marsella, S. C.; Merchant, C.; Miller, L. C.; and Read, S. J. 2022. Disaster world. *Computational and Mathematical Organization Theory*, in press.
- Qin, Y.; Zhang, W.; Lee, R.; Sun, X.; and Sajda, P. 2022. Predictive Power of Pupil Dynamics in a Team Based Virtual Reality Task. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops*, 592–593. IEEE.
- Radford, J.; Pilny, A.; Reichelmann, A.; Keegan, B.; Welles, B. F.; Hoye, J.; Ognyanova, K.; Meleis, W.; and Lazer, D. 2016. Volunteer science: An online laboratory for experiments in social psychology. *Social Psychology Quarterly*, 79(4): 376–396.
- Riedl, C.; Kim, Y. J.; Gupta, P.; Malone, T. W.; and Woolley, A. W. 2021. Quantifying collective intelligence in human groups. *Proceedings of the National Academy of Sciences*, 118(21): e2005737118.
- Riedl, C.; and Woolley, A. W. 2017. Teams vs. crowds: A field test of the relative contribution of incentives, member ability, and emergent collaboration to crowd-based problem solving performance. *Academy of Management Discoveries*, 3(4): 382–403.
- Schelble, B. G.; Flathmann, C.; McNeese, N. J.; Freeman, G.; and Mallick, R. 2022. Let’s Think Together! Assessing Shared Mental Models, Performance, and Trust in Human-Agent Teams. *Proceedings of the ACM on Human-Computer Interaction*, 6(GROUP): 1–29.
- Seraj, E.; Wang, Z.; Paleja, R.; Martin, D.; Sklar, M.; Patel, A.; and Gombolay, M. 2022. Learning efficient diverse communication for cooperative heterogeneous teaming. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 1173–1182.
- Smith, R.; Friston, K. J.; and Whyte, C. J. 2022. A step-by-step tutorial on active inference and its application to empirical data. *Journal of Mathematical Psychology*, 107: 102632.
- Stasser, G.; and Titus, W. 1985. Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology*, 48(6): 1467.
- Stone, P.; Kaminka, G. A.; Kraus, S.; and Rosenschein, J. S. 2010. Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination. In *Proceedings of the Twenty-Fourth Conference on Artificial Intelligence*.
- Tenenbaum, J. B.; Kemp, C.; Griffiths, T. L.; and Goodman, N. D. 2011. How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022): 1279–1285.
- Vasil, J.; Badcock, P. B.; Constant, A.; Friston, K.; and Ramstead, M. J. 2020. A world unto itself: human communication as active inference. *Frontiers in Psychology*, 11: 417.
- Woolley, A. W.; Chabris, C. F.; Pentland, A.; Hashmi, N.; and Malone, T. W. 2010. Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004): 686–688.
- Woolley, A. W.; Chow, R. M.; Mayo, A. T.; Riedl, C.; and Chang, J. W. 2022. Collective Attention and Collective Intelligence: The Role of Hierarchy and Team Gender Composition. *Organization Science*, in press.
- Wu, S. A.; Wang, R. E.; Evans, J. A.; Tenenbaum, J. B.; Parkes, D. C.; and Kleiman-Weiner, M. 2021. Too Many Cooks: Bayesian Inference for Coordinating Multi-Agent Collaboration. *Topics in Cognitive Science*, 13(2): 414–432.
- Wuchty, S.; Jones, B. F.; and Uzzi, B. 2007. The increasing dominance of teams in production of knowledge. *Science*, 316(5827): 1036–1039.