

GOMA: Proactive Embodied Cooperative Communication via Goal-Oriented Mental Alignment

Lance Ying^{1,2}, Kunal Jha³, Shivam Aarya⁴, Joshua B. Tenenbaum², Antonio Torralba², Tianmin Shu⁴

Abstract—Verbal communication plays a crucial role in human cooperation, particularly when the partners only have incomplete information about the task, environment, and each other’s mental state. In this paper, we propose a novel cooperative communication framework, Goal-Oriented Mental Alignment (GOMA). GOMA formulates verbal communication as a planning problem that minimizes the misalignment between the parts of agents’ mental states that are relevant to the goals. This approach enables an embodied assistant to reason about when and how to proactively initialize communication with humans verbally using natural language to help achieve better cooperation. We evaluate our approach against strong baselines in two challenging environments, Overcooked (a multiplayer game) and VirtualHome (a household simulator). Our experimental results demonstrate that large language models struggle with generating meaningful communication that is grounded in the social and physical context. In contrast, our approach can successfully generate concise verbal communication for the embodied assistant to effectively boost the performance of the cooperation as well as human users’ perception of the assistant.

I. INTRODUCTION

Rich verbal communication naturally emerges from human cooperation when people only have partial information about the environments and/or about each other’s mental states [1]. It serves as a complementary source of information, in addition to the visual inputs, to help achieve better cooperation by aligning each other’s mental states (including goals, beliefs, and eventually plans [2], [3], [4]). Recent advances in large language models (LLM) and machine Theory of Mind (ToM) have sparked interest in building cooperative robots that can not only physically cooperate with humans but also verbally communicate with humans using natural language [5], [6]. However, it remains challenging to enable robots to actively *initiate* verbal communication that is both concise (only communicate when necessary) and consistent with the physical environment and the social context (e.g., what humans want to do, believe, know, and need to know).

Studies in psychology has shown that proactive verbal communication serves to align the mental states of agents [7]. Imagine you are going to get some groceries for your mom. As you put on your shoes and walk towards the door, your mom gets out of the kitchen and says “It’s going to rain, get your umbrella, and don’t forget about the avocados.” In

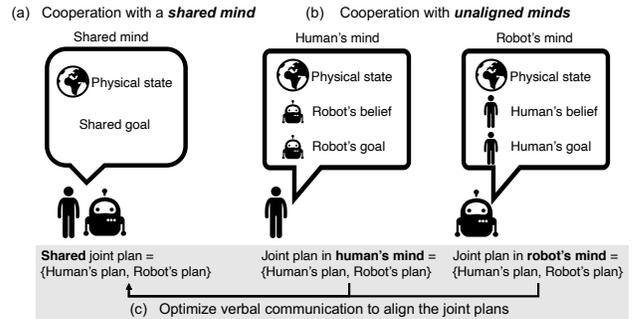


Fig. 1: Illustration of cooperation with a shared mind or misaligned minds and communication optimized via goal-oriented mental alignment. (a) When human and robot minds are perfectly signed (i.e., a shared mind), they share the same belief of the physical state and the same goal, which leads to the same joint plan shared by both agents, an ideal condition for human-robot cooperation. (b) However, in real-world tasks, human and robot minds are typically unaligned, leading to two different (often conflicting) joint plans in their minds. (c) To achieve a shared joint plan that optimizes cooperation, we optimize the robot’s verbal communication to actively align the joint plans in both agents’ minds.

this scenario, your mom decides to communicate with you because she is uncertain whether you have the same beliefs regarding the weather forecast and the required grocery items as you walk out the door.

When cooperating with one another, each agent not only needs to plan for itself but also has to imagine the plans of its partners. Such planning process is termed as *joint planning* [8], [9]. To achieve joint planning, prior works typically assumed that both agents have full observability and complete knowledge about the task. In other words, they have a shared mind, based on which they can derive the same joint plan (Fig. 1(a)). However, in real-world embodied cooperation, robot assistants only have partial observations and often do not know the true human goals (Fig. 1(b)).

The goal of cooperative communication is then to reach a shared mind (two agents’ are perfectly aligned) so that the resulting joint plans in both agents’ minds are the same (Fig. 1(c)). Once we reach such mental alignment condition, both agents know exactly what each other plans to do, and therefore achieve optimal cooperation. However, an agent belief can be about any part of a state. If the state is high dimensional, it is extremely difficult to make sure two beliefs are the same. Our key insight is that we *only* need to align

¹Harvard University, Cambridge, MA 02138, USA lanceying@seas.harvard.edu

²Massachusetts Institute of Technology, Cambridge, MA 01239, USA jbt@mit.edu, torralba@csail.mit.edu

³Dartmouth College, Hanover, NH 03755, USA kunal.a.jha.24@dartmouth.edu

⁴Johns Hopkins University, Baltimore, MD 21218, USA {saarya1, tianmin.shu}@mit.edu

the part of the belief that is relevant to reaching the goal.

Following this insight, we propose a novel cooperative communication framework, Goal-Oriented Mental Alignment (GOMA). In this framework, we aim to generate optimal communication in the belief space. That is, verbal communication, by exchanging information, can help reshape agents’ beliefs. In particular, GOMA first seeks to detect misalignment in agents’ goal-relevant beliefs via divergence between the joint plans based on an agent’s own belief and a simulated hypothetical mind after acquiring additional knowledge from another agent via communication. We then optimize the communication using a proxy reward derived from the divergence between the plans. The resulting communication can then minimize the difference between the joint plan in each agent’s mind and the true joint plan, which improves agents’ coordination.

We evaluate GOMA in two popular human-AI cooperation domains, Overcooked and VirtualHome. Our experimental results with a simulated human agent and real human participants show that our GOMA outperforms strong baselines (including a recent LLM-based baseline). The GOMA-enabled assistant also receives higher subjective ratings from human participants.

In sum, our contributions include (1) a novel embodied cooperative communication framework – GOMA, (2) extensive evaluation of strong baselines and GOMA in two challenging domains, and (3) a human user study that evaluates the task performance of AI assistants and humans’ perception of them.

II. RELATED WORK

A. Communication in Collaboration

Human communication is grounded in cooperative intentions. [7] argues that language communication is a joint activity that attempts to achieve mutual understanding. [1] proposes three communicative motives: requesting help or information, informing the other agents, and sharing feelings or attitudes. These communicative motives help to align the mental states of the agents. Through verbal communication, agents can assess others’ goals, knowledge, emotions, and beliefs, which they can then use to plan for the next actions.

However, verbal communication can also be costly, as it demands cognitive resources and distracts agents when performing actions [10]. Prior works on multi-agent teaming have formalized communication costs in collaborative settings [10], [11], [12], showing that excessive communication can degrade the performance of the team. Therefore, when designing communication policies, the AI assistant needs to communicate useful, concise, and relevant information yet not too frequently.

B. Collaborative and Communicative AI Agent

Communication between humans and robots has also been extensively studied. Most existing literature has focused on one-directional communication where the human instructs the robot [13], [14]. Some recent studies have proposed bi-directional communication. For example, Yuan et al. [15]

proposes a bi-directional human-robot collaborative communication framework that allows the robot to communicate decisions with explanations from human feedback. Unhelkar et al. [12] introduced CommPlann, a bi-directional communication framework that allows the robot to ask for human’s intent, share the robot’s intent, and give commands to humans. There have been recent works that use LLMs as a communication module in bi-directional human-robot communication, (e.g., [5], [16], [17], [18]). While these recent LLM-based agents can achieve certain success, the communication generated by LLMs is often redundant and/or not grounded in agents’ mental states, actions, and plans.

In addition, most human-robot communication frameworks, such as [12], [19], [20], assume full agent observability. The resulting communication is thus only restricted to informing and inquiring about goals and plans. Our work extends to scenarios where both agents have partial observability of the environment and allows the robot to communicate to resolve partial knowledge and false beliefs about the environment state. This requires agents to model and reason about each other’s mental states recursively (e.g., the robot thinks the human thinks the glass is in the fridge, but it knows that it is in the cabinet), which remains a challenge for LLMs today [21]. As a result, such cooperative communication capacity remains an open research question in embodied cooperation.

C. Theory of Mind for Cooperative Robot Planning

There have been many studies on inferring an agent’s goals and beliefs (e.g., [22], [23], [24], [25], [26], [2], [27]), commonly referred to as Theory of Mind reasoning, to better coordinate with humans in collaborative tasks. Previous studies have leveraged explicit mental reasoning to improve cooperative robot planning. This includes generating more expressive or explainable plans to improve humans’ understanding of robots’ plans [28], [29], [30], [31], [4], [32] or better understanding of humans’ cooperative actions [3], all via reasoning about humans’ mental models of the robot. There have also been works on developing a shared joint planner in two agents’ minds to reach optimal coordination by reasoning about both agents’ plans jointly. However, existing works do not allow verbal communication between humans and robots in addition to action planning. Our work aims to fill this gap by jointly planning for actions that change the physical state and verbal communication that changes the mental states of humans and robots.

III. PROBLEM FORMULATION

In this work, we consider two agents, a human user and a robot assistant. To successfully communicate and cooperate, the two agents must infer each other’s mind. We adopt the Interactive Partially Observable Markov Decision Process (I-POMDP) [33], [34] to formulate the mental reasoning between the human and the robot.

A. Background: I-POMDP

I-POMDP is a framework that enables an agent to recursively model other agents, which captures complex social

interactions between agents. Here, we consider the interactions between two agents, i and j , in which agent i infers agent j 's mental state recursively. In an I-POMDP, there are states s^t ; agents' observations, o_i^t and o_j^t , sampled from their conditional observation probabilities, $O_i(o_i^t|s^t)$ and $O_j(o_j^t|s^t)$; and agents' actions a_i^t and a_j^t . Agents have their beliefs, b_i^t and b_j^t , and goals, g_i and g_j . To model the recursive mental reasoning, we define interactive states for the agents, i.e., $is_{i,\ell}$ and $is_{j,\ell}$, at level- ℓ . From agent i 's perspective, we define its interactive state at each level as

- Level 0: $is_{i,0} = s$
- Level 1: $is_{i,1} = (s, b_{j,0}, g_j)$
- ...
- Level ℓ : $is_{i,\ell} = (s, b_{j,\ell-1}, g_j)$

The level- ℓ inference for agent i is to infer the belief $b_{i,\ell}^t = p(is_{i,\ell}^t | o_i^{1:t}, a_i^{1:t-1})$. Since the level- ℓ agent i 's interactive state, $is_{i,\ell}^t = (s^t, b_{j,\ell-1}^t, g_j)$, includes j 's belief at level $\ell - 1$ ($b_{j,\ell-1}^t$), the inference at level ℓ depends on inference at level $\ell - 1$ which depends on inference at level $\ell - 2$, and so on. This recursive inference terminates at level 0. That is, the belief at level-0 is only about the physical state, $b_{i,0} = p(s^t | o_i^{1:t}, a_i^{1:t-1})$. This becomes a standard POMDP [35] which does not model other agents.

B. Two-level Reasoning for Embodied Cooperative Communication

Theoretically, the level of agents' reasoning about other agents' minds can go to infinity (e.g. robot thinks human thinks robot thinks...) yet we cap the depth at two in our model, which is in line with most empirical evidence suggesting that humans rarely engage in greater than 2 levels of recursive Theory of Mind reasoning [36]. Therefore, we adopt a two-level I-POMDP for modeling the mental reasoning between a human user and a robot assistant in embodied cooperation. In particular, we define the mind of each agent as the belief of the level-1 interactive state of the agent. For the human user's mind, we have $m_H = (b(is_{H,1}), g_H) = ((b_{H,0}, b(b_{R,0}), b(g_R)), g_H)$, where $b_{R,0}$ is the robot's interactive state at level 0, i.e., its belief about the physical state; g_R is the robot's goal; and g_H is the human's goal. Similarly, for the robot assistant, we define its mind as $m_R = (b(is_{R,1}), g_R) = ((b_{R,0}, b(b_{H,0}), b(g_H)), g_R)$, where $b_{H,0}$ is the human's belief about the physical state. Intuitively, each mind models the agent's belief about (1) the physical state, (2) another agent's belief about the physical state, and (3) the goal of another agent. Due to the cooperative nature of our problem setting, we further constrain the goal inference to be either one of the following two conditions:

- Condition 1: Both agents share a known common goal;
- Condition 2: The robot's goal is the human goal inferred by the robot, and the human user knows that the robot is trying to help with the inferred human goal.

Condition 1 models human-robot teaming, in which the human and robot agents are teammates who work on the same task assigned to them a priori. Condition 2 models

robot assistance, in which the human's true goal is unknown to the robot a priori, thus the robot must infer the human's goal and provide assistance. In both cases, agents only have partial observability of the physical state, and thus they have to infer both the physical state and each other's belief about the physical state. It is worth noting that our formulation departs from most previous assistance-game setups, which either assume that the agents have full observability or that they share a known goal. As in collaborative tasks, agents often do not have perfect knowledge of the environment and thus need to represent other agents' beliefs differently from theirs and communicate and coordinate their actions, our formulation is more aligned with real-world embodied cooperation.

IV. GOAL-ORIENTED MENTAL ALIGNMENT

As Fig. 1 illustrates, when there is a shared mind, two agents will share the same joint plan. In our Goal-oriented Mental Alignment (GOMA) framework, we formulate communication optimization as the convergence of the current joint plan and the joint plan given a shared mind achieved by exchanging information through verbal communication. In particular, we consider two types of communication – sharing information and requesting information. These are two dominant types of verbal communication in human cooperation [1]. We hypothesize that these are also two types of communication that a robot assistant can *proactively* initiate to achieve joint plan alignment. To reason whether to communicate and what to communicate, we define a proxy reward for minimizing the divergence between plans before and after one type of communication. We summarize GOMA in Algorithm 1, which works with any off-the-shelf action planner. We introduce key components of the algorithm in the rest of the section.

A. Goal Inference and Joint Planning for the Robot

Unless the human goal is given to the robot a priori (i.e., condition 1 defined in Section III-B), the robot must infer the human goal. We adopt the approach introduced by [2], which leverages an LLM to conduct goal inference based on the observed human actions and messages (Line 8-9 in Algorithm 1). We then sample the possible goals of humans

The joint plan for the robot includes two components. First, the robot's policy given its goal and its belief, i.e., $\pi_R(a_R | b_{R,0}, g_R)$. Second, the expected human's policy inferred by the robot, i.e., $\mathbb{E}_{b(b_{H,0}), b(g_H)}[\pi_H(a_H | b_{H,0}, g_H)]$. In practice, we can estimate this expectation via sampling particles of possible human beliefs (i.e., $\{b_{H,0}^{(l)}(s^t)\}_{l=1}^L$ in Algorithm 1) and possible goals (Line 11 in Algorithm 1).

B. Agent Knowledge From Level-0 Belief

Recall that the level-0 belief of an agent $b_{i,0}$ represents the agent's belief of the physical state s . If we partition the state s into multiple sub-states such as states of all objects in the environment, then we can evaluate the uncertainty in the belief of each sub-states. We define the sub-states that have certain belief distributions as knowledge of an agent.

Algorithm 1 GOMA

```

1: Input:  $\text{Planner}()$ ,  $T_{\max}$ 
2: Initialization:  $b(g_H)$ ,  $b_{R,0}(s^0)$ , particles of sampled human beliefs:
    $\{b_{H,0}^{(l)}(s^0)\}_{l=1}^L$ 
3:  $t \leftarrow 1$ ,  $u_R^0 = \text{None}$ 
4: repeat
5:   Observe  $o_R^t$  and receive human message  $u_H^{t-1}$ 
6:   Update level-0 belief:  $b_{R,0}(s^t)$  based on both  $o_R^t$  and  $u_H^{t-1}$ 
7:   Robot knowledge:  $K_R^t = K_R(b_{R,0}(s^t))$ 
8:   Update human goal inference:
9:    $b(g_H) \propto P(a_H^{t-1}|g_H)P(u_H^{t-1}|g_H)b(g_H), \forall g_H \in \mathcal{G}$ 
10:  for all  $l = 1, \dots, L$  do
11:    Sample a human goal based on the goal inference:  $\hat{g}_H^{(l)} \sim b(g_H)$ 
12:    Set the robot goal as the inferred human goal:  $g_R^{(l)} \leftarrow \hat{g}_H^{(l)}$ 
13:    Sample an environment state  $s^t \sim b_{R,0}(s^t)$ 
14:    Sample inferred human observations  $\hat{o}_H^t \sim O_H(\hat{o}_H^t|s^t)$ 
15:    Update  $b_{H,0}^{(l)}(s^t)$  based on both  $\hat{o}_H^t$  and  $u_R^{t-1}$ 
16:    Human plan given the inferred human belief:
17:     $\pi_H(a_H^t|b_{H,0}^{(l)}, \hat{g}_H^{(l)}) \leftarrow \text{Planner}(b_{H,0}^{(l)}, \hat{g}_H^{(l)})$ 
18:    Human plans given the shared minds augmented by different sub-
19:    states in robot knowledge:
20:     $\{\pi_H(a_H^t|b_{H,0}^{(l)}, \hat{g}_H^{(l)}) \leftarrow \text{Planner}(b_{H,0}^{(l)}, \hat{g}_H^{(l)}); \forall s_n^t \in K_R^t\}$ 
21:    Robot plan given the robot belief:
22:     $\pi_R(a_R^t|b_{R,0}, g_R^{(l)}) \leftarrow \text{Planner}(b_{R,0}, g_R^{(l)})$ 
23:    Robot plans given the shared minds augmented by different sub-
24:    states in human knowledge:
25:     $\{\pi_R(a_R^t|b_{R,0}, g_R^{(l)}) \leftarrow \text{Planner}(b_{R,0}, g_R^{(l)}); \forall s_n^t \in$ 
26:     $K(b_{H,0}^{(l)}(s^t))\}$ 
27:  end for
28:   $m_R \leftarrow (b_{R,0}(s^t), \{b_{H,0}^{(l)}(s^t)\}_{l=1}^L, b(g_H))$ 
29:  All possible human knowledge:  $\hat{K}_H^t = \cup_{l=1}^L K(b_{H,0}^{(l)}(s^t))$ 
30:  Construct the utterance space  $U$  based on the robot knowledge  $K_R^t$ 
31:  and all possible human knowledge  $\hat{K}_H^t$ 
32:  Compute  $R(u, M_R), \forall u \in U$  using the plans generated above based
33:  on Eq. (2-4)
34:  Select robot utterance based on the proxy reward:
35:   $u_R^t = \arg \max_{u \in U} R(u, M_R)$ 
36:  Select robot action based on the average plan:
37:   $a_R^t = \arg \max_{a_R \in \mathcal{A}_R} \sum_{l=1}^L \pi_R(a_R|b_{R,0}, g_R^{(l)})/L$ 
38:  Execute the robot action  $a_R^t$  and send the robot utterance  $u_R^t$ 
39:   $t \leftarrow t + 1$ 
40: until  $t = T_{\max}$  or the true goal has not been reached

```

Formally, let us denote a state partition as $s = \{s_n\}_{n=1}^N$ with N sub-states and $b_{i,0}(s_n)$ as the level-0 belief of the sub-state $s_i n$. For instance, if s_n is object n 's state, then $b_{i,0}(s_n)$ is the belief of the object i ' state. Consequently, we define the knowledge of agent i as

$$K_i = K(b_{i,0}) = \{b_{i,0}(s_n) : \mathcal{H}(b_{i,0}(s_n)) < \mathcal{H}_{\max}, n = 1, \dots, N\}, \quad (1)$$

where \mathcal{H} is the entropy of a belief distribution and \mathcal{H}_{\max} is maximum entropy that is considered to be certain. In the example of object states as sub-states, knowledge consists of objects over which the agent has beliefs with high certainty.

C. Shared Mind Augmented by An Agent's Knowledge

An agent i can imagine a shared mind after acquiring knowledge about a sub-state, $b_{j,0}(s_n) \in K_i$, from another agent j via verbal communication, as both agents would share this knowledge after the communication. We define this as the belief merge operation $b_{i,0}^{+s_n} = \text{Merge}(b_{i,0}, b_{j,0}(s_n))$.

Specifically, this merge operation will set the belief of sub-state n of agent i to that of agent j , i.e., $b_{i,0}^{+s_n}(s_n) = b_{j,0}(s_n)$.

D. Divergence Between Plans as Proxy Reward

It is hard to directly estimate the effect of an utterance on the overall task performance. To directly reason what knowledge is critical for aligning the joint plans between agents, we define a proxy reward for communicating about the knowledge of an agent's knowledge. Since the goal of this work is to generate proactive communication initiated by the robot, we model the proxy reward from the perspective of the robot.

We first define the reward of sharing the robot's knowledge of sub-state s_n with the human user as the Kullback–Leibler (KL) divergence between the human agent's plan after communication and before communication.

$$R(\text{share } s_n, M_R) = \text{KL} \left(\mathbb{E}[\pi_H(a_H|b_{H,0}^{+s_n}, g_H)] || \mathbb{E}[\pi_H(a_H|b_{H,0}, g_H)] \right) - C, \quad (2)$$

where $b_{H,0}^{+s_n} = \text{Merge}(b_{H,0}, b_{R,0}(s_n))$ and C is the cost for communication at a time step.

We then define the reward of requesting possible human knowledge of sub-state s_n to inform the robot's plan:

$$R(\text{request } s_n, M_R) = \text{KL} \left(\pi_R(a_R|b_{R,0}^{+s_n}, g_R) || \pi_R(a_R|b_{R,0}, g_R) \right) - C, \quad (3)$$

where $b_{R,0}^{+s_n} = \text{Merge}(b_{R,0}, b_{H,0}(s_n))$.

The plans used to compute the KL-divergence for the proxy rewards can be generated by running an off-the-shelf planner given the corresponding beliefs and goals (Line 16-23 in Algorithm 1).

We also define the reward for not communicating at a step as follows:

$$R(\text{None}, M_R) = 0. \quad (4)$$

E. Communication Optimization

Given the proxy rewards defined above, we can then choose whether and what to communicate based on the robot's mind at each step (Line 27-30 in Algorithm 1). In particular, the utterance space is $U = \{\text{None}\} \cup \{\text{share } s_n; s_n \in K_R\} \cup \{\text{request } s_n; s \in \hat{K}_H\}$, where \hat{K}_H is the inferred human knowledge estimated from the human belief particles: $\hat{K}_H = \cup_{l=1}^L K(b_{H,0}^{(l)})$. We select the best robot utterance at step t as follows:

$$u_R^t = \arg \max_{u \in U} R(u, M_R). \quad (5)$$

We can further generate a natural language message based on the utterance u_R^t to enable communication with real humans. This can be achieved by using GPT-4 [37] to translate u_R^t to natural language through few-shot prompting.

Recipe Name	Ingredient List
Burger	Cooked(Patty), Cooked(Potato), Chopped(Lettuce), Chopped(Tomato)
Pasta	Cooked(Spaghetti), Cooked(Mushroom), Cooked(Cream), Chopped(Basil)
Ramen	Cooked(Noodle), Cooked(Mushroom), Cooked(Egg), Chopped(Scallion)
Steak & Fries	Cooked(Beef), Cooked(Potato), Chopped(Parsley)

TABLE I: Overcooked recipe specifications.



Fig. 2: Example Overcooked environment. In each environment, there are two rooms. The two agents are always in different rooms. An agent cannot observe the other room and has to rely on verbal communication to infer the states of the objects in the other room.

F. Multimodal Mental Update

At each step, the robot will update its mind based on both its observation o_R^t and the messages it sends and receives. In particular, we extract human knowledge $b_{H,0}(s_n)$ from the human message u_h^t via GPT-4 and use it to update the robot’s level-0 belief $b_{R,0}$ jointly with o_R^t (Line 6 in Algorithm 1) operation. For instance, if the human informs the robot of the location of an object, we can update the robot’s level-0 belief with the knowledge of the object’s location. Additionally, if the robot shares knowledge $b_{R,0}(s_n)$ in its utterance, then the robot can assume that the human’s level-0 belief will also be updated accordingly. Thus, in robot mind M_R , we can update $b(b_{H,0})$ using both the shared robot knowledge and the human observation (Line 15 in Algorithm 1). Note that we can sample possible human observations based on the state inferred by the robot’s level-0 belief (Line 13-14 in Algorithm 1). All beliefs are initialized with a uniform distribution (Line 2 in Algorithm 1).

V. EXPERIMENTS

We evaluate our model in two human-AI domains Overcooked and VirtualHome. These two domains cover two distinct alignment objectives. In both domains, there are two agents – a human user and an embodied AI assistant. In Overcooked, the agent’s goal is to align their plans temporally so that certain joint actions can be performed at similar time steps, whereas in VirtualHome, the agents align their beliefs about the location of the objects they try to collect. We describe each in detail below.

A. Overcooked

Overcooked is a popular multiagent game where agents need to collaborate to prepare and cook ingredients, which is also widely used for evaluating human-AI cooperation (e.g.,

Goals	Goal Specification
Set up table	Put [N forks, N plates, N waterglasses or wineglasses] on [kitchentable, coffeetable]
Put groceries	Put [N apple, N salmon, N pudding, N cupcakes] inside [cabinet, fridge]
Prepare food	Put [N apple, N salmon, N pudding, N cupcakes] on [kitchentable, coffeetable]
Load dishwasher	Put [N forks, N plates, N waterglasses or wineglasses] inside [dishwasher]

TABLE II: VirtualHome goal specifications.

[38], [9]). In the original game, agents have full observability. In this study, we extended the Overcooked simulator from [9] by assuming partial observability where each agent cannot observe the other room as shown in Fig. 2. At each step, the AI assistant may share its progress on the task or ask about the human’s progress.

The goal of the collaborating agents is to complete the dishes in the shortest amount of time. To simulate more realistic cooking scenarios, we augment the existing simulator with dynamics that cooked ingredients will gradually cool down. If cooked ingredients are not at the ideal temperature when the dish is served, the team will receive a penalty. This requires both agents to coordinate to avoid temporal misalignment in their plans. For instance, one agent should not finish making the burger too early if the other agent has not started cooking the French fries. The agents can align their plan by choosing to wait for the other agents (e.g. I will start cooking A as soon as the other agent finishes B). There are four recipes in our experiment (Table I): Burger, Spaghetti, Ramen, and Steak, each in a unique room layout. We simulate a human agent using the planner in [9]. Each recipe is run 10 times with different seeds and we report the aggregate results.

Baselines. We evaluate three baselines: Single-agent, No-Communication (No-Comm), and Heuristic-based Communication (Heur-Comm). In Single-Agent, the human completes all the tasks alone. In the No-Comm baseline, no messages are exchanged. In the Heur-Comm baseline, the AI Assistant follows a simple heuristic that shares updates every time a sub-goal has been completed and periodically asks for the human’s progress. The action planner in all methods including GOMA is the same as the planner in [9].

Metrics. We use two performance metrics: speedup and total plan costs. Speedup is calculated by comparing the plan length in each team condition, where the human is working with one of the four collaborative AI models, to the single agent baseline, i.e. $Speedup = L_{single}/L_{team} - 1$.

Total plan cost is the sum of all action and communication costs with penalties applied for sub-optimal dish states due to time lapse between the completion of a hot sub-task (e.g. cooked noodle) and the end of the trial, i.e. $TotalCost = L + U + \sum_{i \in hot_items} \Delta(L_i, L)$ where U is the total number of utterances in a trial, L is the plan length and L_i refers to the time step where item i is completed.

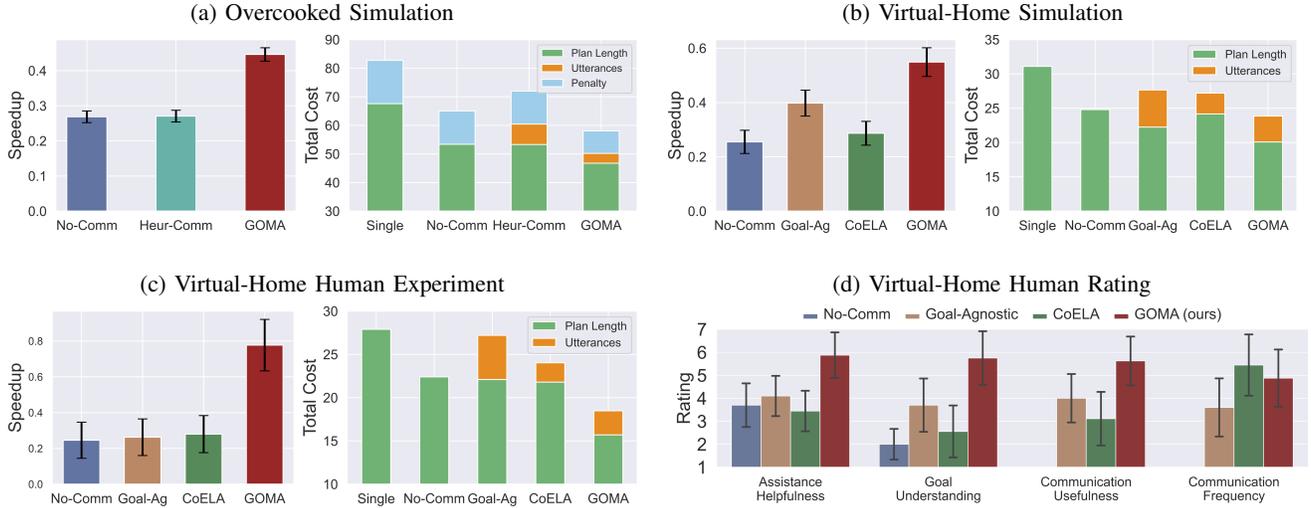


Fig. 3: Experimental results in Overcooked and VirtualHome. The quantitative results from experiments (a, b, c) demonstrate that GOMA led to the greatest speedup (left) and least plan cost (right) compared to other baselines. In human subjective ratings (d), participants find GOMA to be more helpful and communicate more useful information than other models.

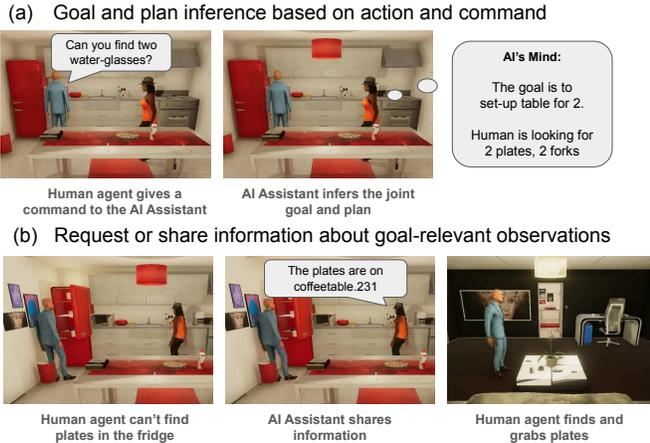


Fig. 4: Example of typical communication enabled by GOMA in VirtualHome. (a) Once the human (in the blue shirt) gives a command to the AI Assistant (in the orange shirt), it infers the human goal and reasons that the human needs 2 plates and 2 forks. (b) As the AI watches the human agent opening the fridge, GOMA informs the human that the plates are on the coffee table. Consequently, the human goes to the coffee table to pick up a plate.

B. VirtualHome

VirtualHome [39] is a multiagent household simulator. In VirtualHome, agents collaborate to complete daily household tasks. In our experiments, we included four common types of household tasks: Set Table, Load Dishwasher, Get Snacks, Stock Fridge. The goal for each task is defined as a set of goal predicates and their counts as defined in Table II. The goal space was defined a priori and accessible to all models. In VirtualHome, each object is associated with a unique object ID, which we used in agents' communication to distinguish the referent from others (e.g. *cabinet.145*).

Simulation Experiment. We simulated 25 collaborative scenarios in VirtualHome across 4 goal types and 5 simulated apartments. Each episode was run 3 times and we reported the averaged results. We simulated the human agent using the MCTS planner from [39]. The simulated human agent requests help by sampling a subset of the goal predicates and replies to the AI assistant's questions. We compared our proposed method against four baselines: Single-agent, No-Communication (No-Comm), Goal-Agnostic (Goal-Ag), and LLM agent (CoELA). The first two were identical to the ones in Overcooked. The Goal-Ag baseline did not infer the joint goals and plan and instead randomly shared information about any objects that the human didn't know. We used CoELA [5] for the LLM agent, which had previously achieved great performance on human-AI collaborative tasks in VirtualHome.

Human Experiment. We developed an online human interface and conducted an experiment with 12 human participants recruited from 3 universities. We adopted a within-subject design where each participant completed 5 trials, each with a different AI assistant model (Single-Agent, No-Comm, Goal-Ag, CoELA, and GOMA) in a randomized order. In total the participants completed 60 trials over 20 unique task scenarios with each scenario receiving 3 data-points. In all collaborative conditions, the interface included a chatbox that allowed the participant and the AI agent to send messages to each other. After completing a trial with an AI assistant, the participants were asked to rate the AI assistant based on four criteria: 1) the assistant is helpful; 2) the assistant understands your goal; 3) the assistant's communication is useful; and 4) the assistant communicates more than necessary. Each criterion is rated on a 7-point Likert scale (1 = Strongly Disagree, 7 = Strongly Agree).

Metrics. In line with previous studies on VirtualHome [39], we evaluated the models' performance by computing

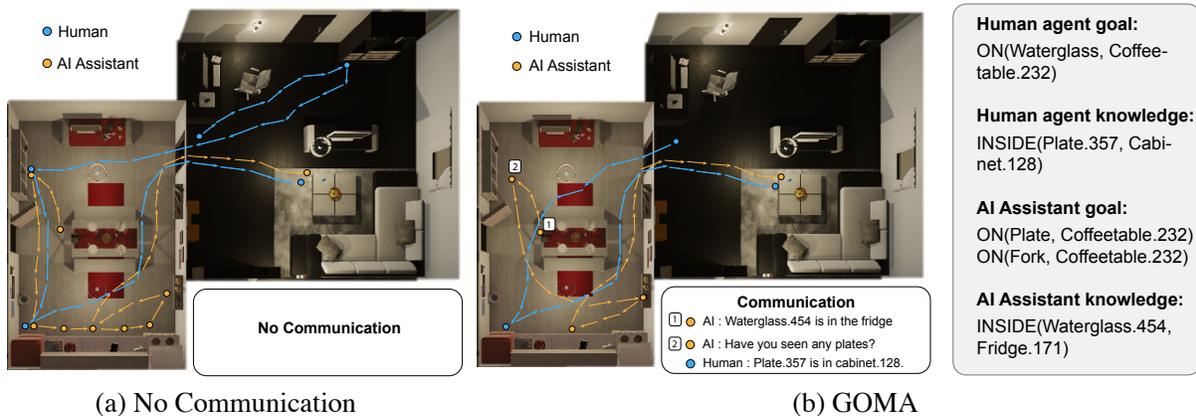


Fig. 5: Agents’ trajectories with No-Comm (left) and with GOMA (right) in a VirtualHome environment. In this example, the team goal is to set up a table for 1 person. The AI Assistant needs to find a plate and a fork while the human is looking for a water glass. Both agents have knowledge about the items that the other agent is looking for but not their own goal objects. In the No-Comm setting, the agents cannot share knowledge and have to open many containers to search for goal items. By inferring the other agent’s goal and communicating goal-relevant knowledge, GOMA drastically reduces the total number of steps taken to complete the task.

1) speedups: counting the number of steps taken to complete the task, and 2) total costs: an overall cost metric that sums up the action and communication cost over the episode.

VI. RESULTS

A. Simulation Experiment

The simulation results are shown in Fig. 3ab. The advantage of collaboration is evident as the Single-agent baseline performed significantly worse than all other collaborative models. Overall, we find that across both Overcooked and Virtual-Home experiments, our model outperformed other baselines in all metrics. The differences between GOMA and other baselines are all statistically significant with $p < 0.01$ across two performance metrics.

In Overcooked, GOMA took on average 46.76 steps to complete the task, achieving a 44.61% speedup. Our model completed the tasks with the lowest costs ($M = 58.06$) compared to the Heuristic-based model ($M = 72.0$) and No-Comm baseline ($M = 65.05$). Additionally, GOMA delivered the dishes in the best condition among all tested models, as signaled by the lowest coldness penalty (7.85).

In VirtualHome, GOMA took on average 20.08 steps to complete the task with a 55.8% speedup. Despite having observed objects relevant to the human agent’s goal, CoELA made few utterances (Mean = 3.03) and focused exclusively on communicating observations of its own goal. For example, when given a command “Please help me find a fork.”, CoELA would respond later “I found fork 323 in cabinet 132.” and did not share any knowledge that may be useful for the human’s subgoal and plan. The Goal-Agnostic model makes frequent (Mean = 5.41) but mostly irrelevant utterances about possible goal objects. However, it did perform slightly better than the No-Comm baseline because, with enough utterances, it occasionally mentions useful information to the human agent.

Unlike baselines, GOMA can communicate and inquire about useful goal-relevant information with the human, leading to improved team performance. We include two qualitative examples of GOMA in VirtualHome simulations in Fig. 4 and 5. In these examples, we show that due to partial observability, the AI Assistant and the human have exclusive knowledge about certain objects relevant to other agent’s subgoals. GOMA allows the AI Assistant to inquire and inform another agent about this goal-relevant information. As a result, the agents can find the goal objects quickly without exhaustively opening and checking all containers.

B. Human Experiment

The human experiment results are shown in Fig. 3cd. Similar to the simulation results, our proposed method had the greatest speedup over a single agent and outperformed all baselines in terms of plan costs. In contrast to the simulation study, the Goal-agnostic model here performed no better than No-Comm and CoELA as participants stopped paying attention to the assistant after it made too many statements irrelevant to the goal. This is shown in the participants’ subjective ratings where participants reported that the Goal-Agnostic baseline communicated more than necessary.

The participants gave a higher subjective rating to our model than other baselines on all 4 items. Interestingly, even though CoELA and GOMA performed goal inference with the same method, the participants thought that only GOMA understood the human’s goal. This is because by communicating goal-relevant information, GOMA implicitly expressed its understanding of the user’s goal, whereas CoELA only communicated the progress of its own subgoal.

VII. CONCLUSION

In this paper, we introduce GOMA, which enables an embodied AI assistant to efficiently and effectively communicate with a human user to achieve optimal cooperation.

GOMA achieves this by reasoning about the other agent’s mental state, assessing the misalignment between mental states, and then proactively initiating necessary communication to exchange goal-relevant information. Our experiments in Overcooked and VirtualHome demonstrate that embodied AI assistants built with GOMA can not only help achieve the human goal faster with lower total plan cost but also receive higher subjective ratings from human participants.

Our study is not without limitations. We have not evaluated GOMA on real-world robot assistants, which we intend to study in the future. We also plan to enhance the flexibility of the communication generation, so that it can communicate about *any* information relevant to the task in an open-ended manner. Finally, we also aim to investigate more general belief representations that go beyond object states.

VIII. ACKNOWLEDGEMENT

This project is funded in part by Lockheed Martin, DARPA Machine Common Science, and Schmidt AI 2050. We thank Hao Liu and Shitian Yang for their contribution.

REFERENCES

- [1] M. Tomasello, *Origins of human communication*. MIT press, 2010.
- [2] L. Ying, T. Zhi-Xuan, V. Mansinghka, and J. B. Tenenbaum, “Inferring the goals of communicating agents from actions and instructions,” in *Proceedings of the AAI Symposium Series*, vol. 2, no. 1, 2023, pp. 26–33.
- [3] D. Hadfield-Menell, S. J. Russell, P. Abbeel, and A. Dragan, “Cooperative inverse reinforcement learning,” in *Advances in Neural Information Processing Systems*, 2016, pp. 3909–3917.
- [4] X. Gao, R. Gong, Y. Zhao, S. Wang, T. Shu, and S.-C. Zhu, “Joint mind modeling for explanation generation in complex human-robot collaborative tasks,” in *2020 29th IEEE international conference on robot and human interactive communication (RO-MAN)*. IEEE, 2020, pp. 1119–1126.
- [5] H. Zhang, W. Du, J. Shan, Q. Zhou, Y. Du, J. B. Tenenbaum, T. Shu, and C. Gan, “Building cooperative embodied agents modularly with large language models,” *arXiv preprint arXiv:2307.02485*, 2023.
- [6] A. Hong, N. Lunscher, T. Hu, Y. Tsuboi, X. Zhang, S. F. dos Reis Alves, G. Nejat, and B. Benhabib, “A multimodal emotional human–robot interaction architecture for social robots engaged in bidirectional communication,” *IEEE transactions on cybernetics*, vol. 51, no. 12, pp. 5954–5968, 2020.
- [7] H. H. Clark, *Using language*. Cambridge university press, 1996.
- [8] M. Kleiman-Weiner, M. K. Ho, J. L. Austerweil, M. L. Littman, and J. B. Tenenbaum, “Coordinate to cooperate or compete: abstract goals and joint intentions in social interaction,” in *COGSCI*, 2016.
- [9] S. A. Wu, R. E. Wang, J. A. Evans, J. B. Tenenbaum, D. C. Parkes, and M. Kleiman-Weiner, “Too many cooks: Bayesian inference for coordinating multi-agent collaboration,” *Topics in Cognitive Science*, vol. 13, no. 2, pp. 414–432, 2021.
- [10] J. MacMillan, E. E. Entin, and D. Serfaty, “Communication overhead: The hidden cost of team cognition.” 2004.
- [11] E. Horvitz and J. Apacible, “Learning and reasoning about interruption,” in *Proceedings of the 5th international conference on Multimodal interfaces*, 2003, pp. 20–27.
- [12] V. V. Unhelkar, S. Li, and J. A. Shah, “Decision-making for bidirectional communication in sequential human-robot collaborative tasks,” in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 2020, pp. 329–341.
- [13] E. C. Williams, N. Gopalan, M. Rhee, and S. Tellex, “Learning to parse natural language to grounded reward functions with weak supervision,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 4430–4436.
- [14] T. Zhi-Xuan, L. Ying, V. Mansinghka, and J. B. Tenenbaum, “Pragmatic instruction following and goal assistance via cooperative language-guided inverse planning,” *arXiv preprint arXiv:2402.17930*, 2024.
- [15] L. Yuan, X. Gao, Z. Zheng, M. Edmonds, Y. N. Wu, F. Rossano, H. Lu, Y. Zhu, and S.-C. Zhu, “In situ bidirectional human-robot value alignment,” *Science robotics*, vol. 7, no. 68, p. eabm4183, 2022.
- [16] C. Zhang, J. Chen, J. Li, Y. Peng, and Z. Mao, “Large language models for human-robot interaction: A review,” *Biomimetic Intelligence and Robotics*, p. 100131, 2023.
- [17] B. Ichter, A. Brohan, Y. Chebotar, C. Finn, K. Hausman, . . . , and C. Kelly, “Do as i can, not as i say: Grounding language in robotic affordances,” in *Proceedings of The 6th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, K. Liu, D. Kulic, and J. Ichnowski, Eds., vol. 205. PMLR, 14–18 Dec 2023, pp. 287–318.
- [18] Z. Mandi, S. Jain, and S. Song, “Roco: Dialectic multi-robot collaboration with large language models,” *arXiv preprint arXiv:2307.04738*, 2023.
- [19] S. Devin and R. Alami, “An implemented theory of mind to improve human-robot shared plans execution,” in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2016, pp. 319–326.
- [20] K. E. Schaefer, E. R. Straub, J. Y. Chen, J. Putney, and A. W. Evans III, “Communicating intent to develop shared situation awareness and engender trust in human-agent teams,” *Cognitive Systems Research*, vol. 46, pp. 26–39, 2017.
- [21] T. Ullman, “Large language models fail on trivial alterations to theory-of-mind tasks,” *arXiv preprint arXiv:2302.08399*, 2023.
- [22] C. L. Baker, J. Jara-Ettinger, R. Saxe, and J. B. Tenenbaum, “Rational quantitative attribution of beliefs, desires and percepts in human mentalizing,” *Nature Human Behaviour*, vol. 1, no. 4, pp. 1–10, 2017.
- [23] T. Zhi-Xuan, J. Mann, T. Silver, J. Tenenbaum, and V. Mansinghka, “Online bayesian goal inference for boundedly rational planning agents,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [24] T. Shu, A. Bhandwaldar, C. Gan, K. Smith, S. Liu, D. Gutfreund, E. Spelke, J. Tenenbaum, and T. Ullman, “Agent: A benchmark for core psychological reasoning,” in *International conference on machine learning*. PMLR, 2021, pp. 9614–9625.
- [25] C. Jin, Y. Wu, J. Cao, J. Xiang, Y.-L. Kuo, Z. Hu, T. Ullman, A. Torralba, J. B. Tenenbaum, and T. Shu, “Mmtom-qa: Multimodal theory of mind question answering,” *arXiv preprint arXiv:2401.08743*, 2024.
- [26] L. Ying, K. M. Collins, M. Wei, C. E. Zhang, T. Zhi-Xuan, A. Weller, J. B. Tenenbaum, and L. Wong, “The neuro-symbolic inverse planning engine (nipe): Modeling probabilistic social inferences from linguistic inputs,” *arXiv preprint arXiv:2306.14325*, 2023.
- [27] L. Ying, T. Zhi-Xuan, L. Wong, V. Mansinghka, and J. Tenenbaum, “Grounding language about belief in a bayesian theory-of-mind,” *arXiv preprint arXiv:2402.10416*, 2024.
- [28] A. Dragan and S. Srinivasa, “Generating legible motion,” 2013.
- [29] F. Stulp, J. Grizou, B. Busch, and M. Lopes, “Facilitating intention prediction for humans by optimizing robot motions,” in *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2015, pp. 1249–1255.
- [30] M. Kwon, S. H. Huang, and A. D. Dragan, “Expressing robot incapability,” in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018, pp. 87–95.
- [31] Y. Zhang, S. Sreedharan, A. Kulkarni, T. Chakraborti, H. H. Zhuo, and S. Kambhampati, “Plan explicability and predictability for robot task planning,” in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 1313–1320.
- [32] X. Gao, L. Yuan, T. Shu, H. Lu, and S.-C. Zhu, “Show me what you can do: Capability calibration on reachable workspace for human-robot collaboration,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2644–2651, 2022.
- [33] P. J. Gmytrasiewicz and P. Doshi, “A framework for sequential planning in multi-agent settings,” *Journal of Artificial Intelligence Research*, vol. 24, pp. 49–79, 2005.
- [34] P. Doshi and P. J. Gmytrasiewicz, “Monte Carlo sampling methods for approximating interactive POMDPs,” *Journal of Artificial Intelligence Research*, vol. 34, pp. 297–337, 2009.
- [35] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, “Planning and acting in partially observable stochastic domains,” *Artificial intelligence*, vol. 101, no. 1-2, pp. 99–134, 1998.
- [36] A. Bosch-Domenech, J. G. Montalvo, R. Nagel, and A. Satorra, “One, two,(three), infinity,. . . : Newspaper and lab beauty-contest experi-

- ments,” *American Economic Review*, vol. 92, no. 5, pp. 1687–1701, 2002.
- [37] OpenAI, “Gpt-4 technical report,” 2023.
- [38] M. Carroll, R. Shah, M. K. Ho, T. Griffiths, S. Seshia, P. Abbeel, and A. Dragan, “On the utility of learning about humans for human-ai coordination,” *Advances in neural information processing systems*, vol. 32, 2019.
- [39] X. Puig, T. Shu, S. Li, Z. Wang, Y.-H. Liao, J. B. Tenenbaum, S. Fidler, and A. Torralba, “Watch-and-help: A challenge for social perception and human-ai collaboration,” *arXiv preprint arXiv:2010.09890*, 2020.