

Adaptation and Communication in Human-Robot Teaming to Handle Discrepancies in Agents' Beliefs about Plans

Yuening Zhang, Brian Williams

Massachusetts Institute of Technology
zhangyn@mit.edu, williams@mit.edu

Abstract

When agents collaborate on a task, it is important that they have some shared mental model of the task routines – the set of feasible plans towards achieving the goals. However, in reality, situations often arise that such a shared mental model cannot be guaranteed, such as in ad-hoc teams where agents may follow different conventions or when contingent constraints arise that only some agents are aware of. Previous work on human-robot teaming has assumed that the team has a set of shared routines, which breaks down in these situations. In this work, we leverage epistemic logic to enable agents to understand the discrepancy in each other's beliefs about feasible plans and dynamically plan their actions to adapt or communicate to resolve the discrepancy. We propose a formalism that extends conditional doxastic logic to describe knowledge bases in order to explicitly represent agents' nested beliefs on the feasible plans and state of execution. We provide an online execution algorithm based on Monte Carlo Tree Search for the agent to plan its action, including communication actions to explain the feasibility of plans, announce intent, and ask questions. Finally, we evaluate the success rate and scalability of the algorithm and show that our agent is better equipped to work in teams without the guarantee of a shared mental model.

Introduction

When agents collaborate on a task, it is important that they have some shared mental model of the task routines – the set of feasible plans towards achieving the goals. However, in reality, situations often arise that such a shared mental model cannot be guaranteed. For example, in online multi-player games or search-and-rescue missions, people trained separately could form an ad-hoc team where they may follow different conventions. Even if the team has a set of shared routines, novel situations may still occur in which some contingent constraint that forbids certain plans to be taken becomes known only by some agents. In these situations, experienced teammates keep in mind what others know and what actions they may take, and communicate when necessary to make sure the team converges on a feasible plan of action.

Previous work on human-robot teaming, *Pike* (Levine and Williams 2018), assumed that agents share common knowl-

edge of the feasible plans for the task encoded in a knowledge base. Under an equal partner setting, each agent observes the actions taken by others and adapts their actions accordingly, only taking what is still feasible. This approach allows fluid human-robot interaction but breaks down when the common knowledge assumption no longer holds.

In this work, we generalize the approach to handle situations where there may be discrepancies in agents' beliefs about plans by incorporating epistemic logic (Van Ditmarsch et al. 2015), as it provides an explicit representation of agents' nested beliefs towards each other and a mechanism to model communication between agents.

The contribution of this paper is threefold: (1) We propose the formalism of conditional doxastic logic (Baltag and Smets 2008) extended to knowledge bases in order to represent agents' nested beliefs on the set of feasible plans for the task and the state of execution. (2) We model both execution and a rich set of communication actions within the framework, including explanation, intent announcement, and question-asking actions, that allows agents to explicitly talk about the feasibility of plans and exchange their intent. (3) We provide an online execution algorithm based on Monte Carlo Tree Search (MCTS) for the agent to dynamically plan its action to adapt to others or communicate to resolve the discrepancy. Finally, we evaluate the success rate and performance of the algorithm through experiments.

Motivating Example

Consider a pedagogical example where a robot (our agent) and a human collaborate to prepare a drink. The robot has a manipulator arm that can fetch a mug or a glass as the container, and the human can brew some coffee or take some orange juice from the fridge for the drink. For the task to succeed, it must satisfy that: (C1) the mug has to go with the coffee and the glass has to go with the orange juice. Under an equal partner setting, from the robot's perspective:

Case1 If the human doesn't believe constraint C1 holds and thinks that any container can go with any drink, then the robot can adapt to the human by waiting for the human to take the drink first, then fetch the corresponding container. The robot can also explain to the human about constraint C1, especially if the task requires the robot to fetch the container first. The robot can also announce the intent for the human

to choose coffee, in which case it can just fetch the mug.

Case2 If the human has determined a choice of coffee or juice, but the robot doesn't know which one, the robot may wait for the human to pick first so that it can distinguish their intent, or it can ask the human about their intent.

Case3 If the human picked up the juice but doesn't know that the robot couldn't reach the glass, the robot may explain the constraint and that the task has failed.

Background

In order to represent the agents' nested beliefs, our representation builds on top of conditional doxastic logic (Baltag and Smets 2008), which is one variant in the broader epistemic logic literature. Compared to epistemic logic, it allows the modeling of false beliefs and belief revision by pre-encoding the conditional belief of the agents within the model. Given a set of agents Ag and a set of atomic propositions At , *conditional doxastic logic* $\mathcal{L}(At, Ag)$ is defined by the following Backus-Naur Form (BNF):

$$\varphi := p \mid \neg\varphi \mid (\varphi \wedge \psi) \mid B_a^\psi\varphi,$$

where $p \in At$, $a \in Ag$. $B_a^\psi\varphi$ reads as "agent a believes φ given ψ ". Denoting \top as tautology, $B_a\varphi := B_a^\top\varphi$ means that agent a believes φ . Its semantic model is a *plausibility model*, which is a tuple $M = \langle W, \{\leq_a\}_{a \in Ag}, L \rangle$, where

- W : a non-empty set of possible worlds,
- $\leq_a \subseteq W \times W$: binary relation on W imposing a relative plausibility order between any two worlds for agent a ,
- $L : W \rightarrow 2^{At}$: valuation function mapping each world to the set of atomic propositions that hold in the world.

$w \leq_a v$ means that agent a considers w to be at least as plausible as v . $<_a := \leq_a \cap \not\geq_a$ denotes a strict plausibility order. $\simeq_a := \leq_a \cap \geq_a$ denotes an equi-plausibility order. $\sim_a := \leq_a \cup \geq_a$ denotes epistemic indistinguishability, and $cc_a(w) := \{v \in W \mid w \sim_a^* v\}$ is the set of worlds that agent a finds (possibly more or less) plausible given world w , where \sim_a^* is the transitive closure of \sim_a . Note that the plausibility relation is reflexive, transitive, locally connected, that is, $v \in cc_a(w)$ implies $v \leq_a w$ or $w \leq_a v$, and well-founded, that is, $min_a(S) := \{w \in S \mid \forall v \in S : v \not<_a w\}$ is well-defined, which is the subset of worlds in S that agent a finds most plausible. A pair (M, w) is a *pointed plausibility model*, which describes a conditional doxastic state with a pointed view at world $w \in W$, i.e. taking w as the true world. The truth of a formula $\varphi \in \mathcal{L}(At, Ag)$ on (M, w) , i.e. $(M, w) \models \varphi$, can be defined inductively as follows:

- $(M, w) \models p$ iff $p \in L(w)$
- $(M, w) \models \neg\varphi$ iff $(M, w) \not\models \varphi$
- $(M, w) \models \varphi \wedge \psi$ iff $(M, w) \models \varphi$ and $(M, w) \models \psi$
- $(M, w) \models B_a^\psi\varphi$ iff $min_a([\psi]_M \cap cc_a(w)) \subseteq [\varphi]_M$, where $[\varphi]_M := \{w \in W \mid (M, w) \models \varphi\}$ is the set of worlds in M in which φ holds.

Figure 1 shows an example state represented by a pointed plausibility model (M, w_1) with agents a and b . The two worlds w_1 and w_2 are labeled with the atomic propositions

that hold in the respective worlds. The pointed world w_1 highlighted in bold represents the true world in which p holds. The single arrow pointing from w_1 to w_2 labeled with b indicates that agent b considers w_2 to be strictly more plausible than w_1 . We say $(M, w_1) \models B_b\neg p$ since $min_b(cc_b(w_1)) = \{w_2\} \subseteq [\neg p]_M$. If it is instead a double-headed arrow, then it means that agent b considers w_1 and w_2 to be equally plausible. The lack of any arrow between w_1 and w_2 for agent a indicates that $w_1 \not\sim_a w_2$, that is, when in w_1 or w_2 , agent a does not consider the other world plausible at all. Since the plausibility relation is reflexive, the self-loops indicate that whichever world it is, the agents find the world plausible. $(M, w_1) \models B_a p$ since $min_a(cc_a(w_1)) = \{w_1\} \subseteq [p]_M$, and $(M, w_1) \models B_a B_b \neg p$ since $min_a(cc_a(w_1)) = \{w_1\} \subseteq [B_b \neg p]_M = \{w_1, w_2\}$.




Figure 1: Example pointed plausibility model with legend

An action is defined by a *plausibility action model* $A = \langle \Sigma, \{\leq_a\}_{a \in Ag}, pre, post \rangle$, which has a similar structure except instead of a set of worlds W , it has a set of events Σ representing possible events that may occur in the action. pre and $post$ are functions that assign to each event $\sigma \in \Sigma$ a precondition and a postcondition in $\mathcal{L}(At, Ag)$ respectively, where the postcondition of an event is restricted to a conjunction of literals over At or \top . A *pointed plausibility action model* is a pair (A, σ) , $\sigma \in \Sigma$, which describes an action where σ is the true event.

In general, a pointed plausibility model for state or action can point at multiple worlds, such as (M, W_d) or (A, Σ_d) . W_d and Σ_d are called the *designated* worlds or events. For example, given a state (M, w) , (M, W_d) with $W_d = min_a(cc_a(w))$ represents agent a 's local perspective of the state, where W_d includes all the worlds that agent a finds the most plausible. (M, W_d) is a *global* state if $|W_d| = 1$. Additionally, $(M, W_d) \models \varphi$ iff $(M, w) \models \varphi$ for all $w \in W_d$. An action act updates a state s through *action-priority update* $s \otimes act$, which we refer the readers to the details in (Baltag and Smets 2008; Bolander and Andersen 2011).

Approach Overview

Our solution requires answering three questions: (1) what representation to use to capture the agents' nested beliefs of the set of feasible plans and the state of execution, (2) how to model execution and communication actions and how they update the state, and (3) how to strategically choose the next action. Our key insight is to extend conditional doxastic logic to describe knowledge bases, and use the knowledge bases to encode the feasible plans and state of execution, so that we can describe agents' beliefs on the plan space instead of their beliefs on state. As a result of this new logic,

execution and communication actions can be defined which operate by adding or removing constraints from the knowledge bases. With the state and action models defined, we use an MCTS-based algorithm to simulate forward in the next k -step horizon to decide what is the best action to take.




Figure 2: Plausibility model for nested beliefs on plans

For example, Figure 2 captures the agents' nested beliefs on plans from Case1. Each world in the plausibility model is now a knowledge base that contains the constraints of the task. Since H (human) finds w_2 more plausible, H believes that constraint C1 does not need to hold. An example action where the robot announces the intent of coffee is shown in Figure 3 (left). The action has a single event whose precondition is that R (robot) must believe that adding the constraint of coffee is satisfiable given its belief of the current feasible plans. As a result of the action, all worlds now have the constraint of coffee added, including w_2 that H believes in.




Figure 3: intent announcement (left) & resulting state (right)

Representing Team's Nested Beliefs on Plans

We describe our representation in two parts: (1) conditional doxastic logic on knowledge bases and its semantics, (2) our task representation and its encoding in the knowledge base.

Conditional Doxastic Logic on knowledge Bases

Given a finite set of atomic propositions At , and a finite set of agents Ag , *conditional doxastic logic on knowledge bases* $\mathcal{L}_{KB}(At, Ag)$ is defined by the BNF:

$$\varphi := in(c) | entailed(c) | \neg\varphi | (\varphi \wedge \varphi) | B_a^c \varphi,$$

in which $a \in Ag$, $c \in \mathcal{C}(At)$, where $\mathcal{C}(At)$ is the classical propositional logic $c := p | \neg c | (c \wedge c)$, $p \in At$. Note that the formulation naturally extends to constraint systems with finite-domain variables, which is what we use. We hence refer to c as a constraint. $in(c)$ means constraint c is an explicit member of the constraints in the knowledge base, $entailed(c)$ means constraint c is entailed by the knowledge base, and we define $sat(c) := \neg entailed(\neg c)$, which means constraint c is satisfiable by the knowledge base.

The plausibility model for $\mathcal{L}_{KB}(At, Ag)$ is a tuple $M = \langle W, \{\leq_a\}_{a \in Ag}, KB \rangle$, where W and \leq_a are the same as before and $KB : W \rightarrow \mathbf{KB}_{\mathcal{C}(At)}$ is a function that maps each

world to an associated knowledge base in $\mathcal{C}(At)$. When determining the truth of a formula $\varphi \in \mathcal{L}_{KB}(At, Ag)$ on a pointed plausibility model (M, w) , we replace the first rule on $(M, w) \models p$ in the inductive rules with the following:

- $(M, w) \models in(c)$ iff $c \in KB(w)$
- $(M, w) \models entailed(c)$ iff $KB(w) \models c$

We can say the following about the state in Figure 2:

- $B_R in((mug \wedge coffee) \vee (glass \wedge juice))$
- $B_R B_H \neg in((mug \wedge coffee) \vee (glass \wedge juice))$
- $\neg B_R entailed(mug \wedge coffee)$
- $B_R \neg sat(mug \wedge juice) \wedge B_R B_H sat(mug \wedge juice)$

Task Representation & Encoding

The set of feasible plans towards achieving the goals of the task forms a *plan library* for the task. Additionally, actions may be ordered in the plan, such as requiring the container to be picked up first before the drink. Therefore, our task representation is a *temporal plan library* $\langle V, E, O, C \rangle$, where:

- V is a set of decision variables with $domain(v), v \in V$.
- E is a set of time points with guard condition $guard(e)$ for each $e \in E$, a conjunction of decision variable assignments. e should be executed iff $guard(e)$ is satisfied.
- O is a set of ordering constraints $o = \langle e_i, e_j, guard(o) \rangle$, $o \in O$, requiring time point e_i to precede time point e_j in execution order if its guard condition $guard(o)$ is satisfied. We assume $guard(o) \models guard(e_i) \wedge guard(e_j)$.
- C is a set of constraints scoped on V .

The time points represent the actual events of taking the actions. In multi-agent case, a *multi-agent temporal plan library* $\langle V, E, O, C, Ag, f \rangle$ additionally has a set of agents Ag and a function $f : E \rightarrow Ag$ that maps each time point to an agent that it belongs to. In our formulation, the decision variables do not have ownership. This reflects our equal partner setting in which decisions do not belong to any agent and an announced intent can affect multiple agents' actions.

The plan library represents a set of *candidate subplans* G , where a subplan $g \in G$ is a full assignment to all the decision variables V . We use E_g and O_g to denote the set of time points and ordering constraints activated by g , i.e. those whose guard conditions are satisfied. A subplan induces a set of total orderings on E_g that satisfies O_g , which we denote by T_g . A subplan g is *feasible* iff all the constraints C are satisfied, i.e. $\forall c \in C, g \models c$, and there exists a total ordering of E_g that satisfies O_g , i.e. $T_g \neq \emptyset$.

Execution As execution progresses, decision variables are gradually grounded either implicitly from the execution of time points or explicitly from announcement of intent. The *execution state* is a tuple $\langle t, C_I \rangle$, where t is an *execution history*, which is a total ordering of time points (e_i, e_j, \dots, e_k) that have been executed, and C_I is the set of intents that have been announced during execution. An intent, in its most general form, can be an arbitrary constraint scoped on V , but is commonly an assignment to a specific decision variable. The subplans that are *feasible with respect to* $\langle t, C_I \rangle$ include any feasible subplan g such that there exists $t_g \in T_g$, where t

is the prefix of t_g , and g satisfies C_I . Execution *fails* when there exists no feasible subplan with respect to $\langle t, C_I \rangle$. Execution *succeeds* when there exists a feasible subplan g with respect to $\langle t, C_I \rangle$ such that $t \in T_g$. Note that execution can succeed without ever converging to a unique subplan, and it is possible for further time points to be executed and move away from the success state.

Encoding in Knowledge Base We encode the plan library and the execution state in the knowledge base, so that at any point during execution, the knowledge base contains all the feasible subplans with respect to $\langle t, C_I \rangle$. We ensure that the knowledge base is consistent iff execution has not failed.

The variables of the encoding include (1) a discrete variable for each decision variable $v_i \in V$ with the same domain $domain(v_i)$, and (2) a boolean variable for each time point $e_i \in E$ with domain $\{T, F\}$, representing if the time point is executed. We add the following constraints to the knowledge base prior to execution:

- The constraints C as defined in the plan library.
- For each time point e_i , $((e_i = T) \rightarrow guard(e_i))$, i.e. if time point e_i is executed, its guard condition must hold.
- Negation of *nogoods* (Katsirelos and Bacchus 2005) that represent any combination of choices of V that would result in an inconsistent ordering of time points. This can be computed from the ordering constraints O .

During execution, we may additionally add to the KB:

- Announced intents C_I .
- For each execution of time point e_j , a conjunction of (1) assignment of variable e_j to T, and (2) the negation of the guard condition $\neg guard(o)$ for any ordering constraint $o = \langle e_i, e_j, guard(o) \rangle$ in which the predecessor e_i has not been executed by the time e_j is executed.

The last rule ensures that for any ordering constraint $o = \langle e_i, e_j, guard(o) \rangle$, if the guard condition holds, then if e_j is executed, e_i must also have been executed, hence satisfying the ordering constraint. Note that we only encode the set of time points that have been executed, instead of their actual order of execution. With the above encoding, given a knowledge base KB, execution fails iff $KB \models \perp$. Execution succeeds iff there exists subplan g , i.e. a full assignment of V , such that $KB \wedge g \not\models \perp$ and $\forall e_i \in E_g, KB \models (e_i = T)$. We denote the success condition by $suc_{(V,E)}$, and say that execution succeeds iff $KB \models suc_{(V,E)}$. Additionally, $(M, w) \models suc_{(V,E)}$ iff $KB(w) \models suc_{(V,E)}$.

Take Case 1 as an example, the knowledge base contains discrete variables *container* with domain $\{mug, glass\}$ and *drink* with domain $\{coffee, juice\}$, and boolean variables $e_{mug}, e_{glass}, e_{coffee}, e_{juice}$ representing the events of picking up each item. Using *mug* as a shorthand for $(container = mug)$ and similarly for others for the purpose of decluttering, the constraints include:

1. $(mug \wedge coffee) \vee (glass \wedge juice)$
2. $(e_{mug} = T) \rightarrow mug$, similarly for other time points

Note that we use the same shorthands throughout the rest of the paper. In this example, when the robot picks up the




Figure 4: Action representations

mug, constraint $(e_{mug} = T)$ is added to the knowledge base. From 2 above, we now have $KB \models mug$, and consequently from 1, we have $KB \models coffee$, which limits the human's choice of drink to coffee. Picking up juice is no longer feasible since $KB \wedge (e_{juice} = T) \models \perp$. Consider another case where the robot's action must precede the human's action, i.e. there are ordering constraints $\langle e_{mug}, e_{coffee}, mug \wedge coffee \rangle, \langle e_{glass}, e_{coffee}, glass \wedge coffee \rangle$, etc. If the human picks up the coffee before the robot takes any action, then $(e_{coffee} = T) \wedge \neg(mug \wedge coffee) \wedge \neg(glass \wedge coffee)$ is added, resulting in an inconsistent knowledge base.

Dynamic Model of Evolution

In this section, we describe how the model evolves as a result of execution or communication actions. We first introduce the plausibility action model for our extended logic, then describe how to model each type of action. In this work, we assume that agents observe all actions that are taken, that is, all actions are public.

Plausibility Action Model for Knowledge Bases

A *plausibility action model* A for $\mathcal{L}_{KB}(At, Ag)$ is a tuple $\langle \Sigma, \{\leq_a\}_{a \in Ag}, pre, post \rangle$, where Σ and \leq_a are the same, and pre and $post$ are functions that map each event to a formula in $\mathcal{L}_{KB}(At, Ag)$. The postcondition is restricted to a conjunction of $in(c)$, which adds constraint c to the knowledge base, and $\neg in(c)$, which removes constraint c from the knowledge base if it exists, as well as \top , i.e. nothing changes. For this paper, we further restrict the postcondition to be either $in(c)$ or \top , i.e. adding at most one constraint to the knowledge base. The *action-priority update* updates the knowledge bases as described accordingly.

Execution Action An execution action is the action of an agent executing a time point, such as robot picking up *mug*. Recall that in our setting, each time point is assigned to an agent who can execute it. Given that the time point being executed is $e_i \in E$, and the agent who executes it is $a = f(e_i)$, the simplest case of execution of time point e_i that has no potential predecessors is shown in Figure 4 (left). We assume agents are rational and for agent a to execute time point e_i , it needs to believe that executing e_i is feasible, i.e. $B_a sat(e_i = T)$. All agents observing the action also observe the truth of agent a having such belief. For the postcondition, as e_i is executed, the constraint $(e_i = T)$ is added to the knowledge base.

When there are potential predecessors for e_i , we need to make sure the corresponding ordering constraints are satisfied. The postcondition of the event should always add




Figure 5: Example execution action of e_i with ordering constraints $\langle e_j, e_i, \text{guard}(o_j) \rangle$ and $\langle e_k, e_i, \text{guard}(o_k) \rangle$

($e_i = T$) to the knowledge base, and for any ordering constraint $o = \langle e_j, e_i, \text{guard}(o) \rangle$, add ($\neg \text{guard}(o)$) on condition that $\neg \text{entailed}(e_j = T)$, that is, if e_j has not been executed. While a more succinct action model specification is possible (Van Ditmarsch and Kooi 2008), we use the standard form defined above by taking the cross product of all the predecessors, and creating equi-plausible events for them as shown in Figure 5. Even though the size of the action model is exponential to the number of potential predecessors each time point has, because the preconditions of these events are mutually exclusive, the model size for the updated state will not increase as a result of the action update.

Intent Announcement The model for an intent announcement action is shown in Figure 4 (middle). For agent a to announce the intent, it must believe that it is satisfiable, hence the precondition $B_a \text{sat}(c)$. The intent is added as a postcondition. Figure 3 shows an example intent announcement.

Explanation The model for an explanation action where agent a explains its belief of $\varphi \in \mathcal{L}_{KB}(At, Ag)$ is shown in Figure 4 (right). The precondition says that agent a has to believe φ , i.e. agents cannot lie about their belief. To the other agents, the explanation is essentially a public announcement that agent a believes φ . This means that whether a particular agent adopts the explainer’s belief depends on the conditional belief pre-encoded in the initial pointed plausibility model, which specifies how an agent’s belief gets revised when a new piece of evidence is received.




Figure 6: Explanation action (left) and resulting state (right)

An example explanation action where the robot explains constraint C1 is shown in Figure 6. Based on the pointed plausibility model in Figure 2, upon the announcement that the robot in fact believes C1 holds, w_2 is eliminated as it does not satisfy the precondition, and the human is left with w_1 in which C1 holds. Depending on the initial conditional

belief, it is also possible to have situations where the human does not trust the robot and does not adopt its belief.

In this paper, we restrict the explained formula to be of the BNF form $\varphi := \neg\varphi | B_a\varphi | \text{in}(c)$, where $a \in Ag$, $c \in \mathcal{C}(At)$. This simplifies the explanations in that (1) the explained formula cannot be arbitrarily complex such as $B_a \text{in}(c) \rightarrow B_b \text{in}(c)$, (2) the explanation must be about whether the knowledge base contains a constraint or not, instead of the satisfiability or entailment of an arbitrary constraint. This is similar in spirit to the idea of abductive explanations, where we want to give an explanation c such that together with the existing theory T , it explains an explanandum O , i.e. $T \cup \{c\} \models O$. In this case, what is satisfiable or entailed is often the explanandum, and what constraints should or should not be in the theory is what we explain.

Question-Asking An agent can ask another agent about something that it is uncertain of. Since we assume public actions, the answer is observed by all agents. Given that agent a is asked about its belief on formula $\varphi \in \mathcal{L}_{KB}(At, Ag)$, the pointed plausibility action model is shown in Figure 7. We place the same restriction on φ as in the explanation actions.




Figure 7: Question-asking action




Figure 8: Example pointed plausibility model where the robot is uncertain about the human’s choice

Using Case 2 as an example, the robot does not know which choice of drink the human has determined on, which is represented by the pointed plausibility model in Figure 8. The robot can ask a question about the human’s belief on $\text{in}(\text{coffee})$, i.e. whether its intent is to take coffee. The resulting state would have the double-headed arrow in the middle labeled with R removed, i.e. the robot will be able to distinguish human’s intent.

Online Execution Problem

We assume that execution is asynchronous, all actions are public, and communication has a cost. We also assume the discrepancies in beliefs come only from the agents’ initial beliefs on constraints C , i.e. they share the belief on the rest of the plan library such as the ordering constraints and the guard for the time points. In this paper, we assume that a task involves two agents (e.g. robot and human), though there is no theoretical barrier to applying it to more agents. The online execution problem from a single agent’s perspective,

say agent a , involves taking the input of the following prior to execution:

- A multi-agent temporal plan library $\langle V, E, O, \{\}, Ag, f \rangle$.
- A pointed plausibility model $s_0 = (M, W_d)$ capturing agents' initial nested beliefs on constraints C from agent a 's perspective, such that $W_d = cc_a(w), \forall w \in W_d$.

Note that constraints C is empty in the plan library as it is captured by the input of s_0 . W_d includes any world that agent a finds plausible (not necessarily most plausible), and we assume that (M, W_d) captures the ground truth state (M, w^*) as one of its possibilities, so that the agent's belief can also be revised if needed. During execution, the agent receives the input of a stream of actions that are taken by itself or others in real-time, including execution and communication actions. Each action triggers a callback and the agent outputs an action to be taken or none.

Overall Algorithm Upon receiving an action, our agent determines how it should act next – either take an execution action, communicate, or wait for others to act. It simulates forward to predict the utility of each possible action, e.g. if others may follow up with incorrect actions, or if many communication actions will be needed. The algorithm draws insight from epistemic planning for implicit coordination (Engesser et al. 2017) and relies on the agent's ability to take others' perspectives to predict their actions.

The overall algorithm is illustrated in Algorithm 1, which we name it *Epistemic Pike (EPike)* after *Pike* (Levine and Williams 2018). Prior to execution, the agent compiles the initial state s from s_0 and the plan library as described in the knowledge base encoding section (line 1). Upon receiving an action act' , the updated state (line 3) is checked for several conditions. If the agent believes that execution has failed (line 4), then it explains the failure when some agent might not know (line 5 - 6). If the agent is unsure about whether execution has failed (line 7), then it sees if it can ask someone to distinguish it (line 8). If both conditions do not apply, then execution has not failed, and the agent checks if execution has succeeded (line 9). If so, then the agent explains it when some agent might not know (line 10 - 11). If execution has not succeeded, then the agent searches for the next action to take, if any, to progress toward completing the task (line 13). Note that when a question-asking action is taken, we wait until the answer is announced before encoding the answer as an explanation action that gets observed.

Each online subroutine in Algorithm 1 calls an MCTS algorithm with a different configuration. The MCTS algorithm simulates the team's possible execution in the next k -step horizon, and based on the result of the simulations, the agent decides if it should take an action now and which action to take. Our MCTS algorithm can be configured on (1) the termination conditions, including the horizon k , (2) which types of actions to consider for both the ego agent and the other agents, and (3) the penalties for communication actions, since we assume communication has a cost. We add a fifth type of action, *noop*, to represent agent taking no action and waiting for others to act.

For SEARCHACTION, a node terminates if its state s satisfies either $s \models entailed(\perp)$, which gives a utility of 0 (ex-

Algorithm 1: Online Execution for Ego Agent a

Input : $V, E, O, Ag, f, s_0 = \langle M, W_d \rangle$, agent a ,
 Online: act' , an observed action
Output: Online: act , an action or *None*
 1 **Offline:** $s \leftarrow \text{COMPILEINITIALSTATE}(s_0, V, E)$
 2 **Online upon observing act' :**
 3 $s \leftarrow s \otimes act'$
 4 **if** $s \models B_a entailed(\perp)$ **then**
 5 **if** $s \not\models B_a (\wedge_{i \in Ag} B_i entailed(\perp))$ **then**
 6 **return** EXPLAINFAILURE(s)
 7 **else if** $s \not\models B_a \neg entailed(\perp)$ **then**
 8 **return** ASKIFFFAILURE(s)
 9 **else if** $s \models B_a suc_{V,E}$ **then**
 10 **if** $s \not\models B_a (\wedge_{i \in Ag} B_i suc_{V,E})$ **then**
 11 **return** EXPLAINSUCCESS(s)
 12 **else**
 13 **return** SEARCHACTION(s)
 14 **return** *None*;

ecution fails), or $s \models suc_{(V,E)}$, which gives a utility of 1 (execution succeeds), or if simulation reaches a horizon of k , which gives a utility of 1 (execution has not failed). Note that only execution actions increment the horizon, since we care about the outcome after the next three physical actions. During search, we consider all five types of actions (including noop) from all agents, except for the intent announcement and question-asking actions from the other agents. They can be reasonably omitted to reduce the search space, since they may be unpredictable and ignoring it does not prevent the simulated execution to reach success state.

For the rest of the subroutines, a node terminates if its state s satisfies $s \models \wedge_{i \in Ag} B_i entailed(\perp)$ for EXPLAINFAILURE, $s \models B_a entailed(\perp) \vee B_a \neg entailed(\perp)$ for ASKIFFFAILURE, and $s \models \wedge_{i \in Ag} B_i suc_{(V,E)}$ for EXPLAINSUCCESS, all giving a utility of 1. For EXPLAINFAILURE and EXPLAINSUCCESS, only explanation and question-asking actions of the ego agent are considered. Asking a question may still be useful if the agent is uncertain about what others currently believe. For ASKIFFFAILURE, only question-asking actions for the ego agent are considered. In these cases, since the ego agent is just looking to inform others or ask a question, it is reasonable to ignore what other agents may do. To penalize communication, we set a penalty factor of 0.9 for explanation actions and question-asking actions, and 0.85 for intent announcement actions, though the values may change depending on applications. Penalty is a multiplication factor to the utility of the node. Execution actions and noop action are not penalized.

Search Tree We describe the expansion of the search tree using SEARCHACTION as an example, before discussing the details of MCTS. A partially expanded search tree is shown in Figure 9. There are four types of nodes in the tree: root decision node (bold circle), split nodes (diamonds), predict nodes (squares), and decision nodes (circles). Each node has




Figure 10: Success rate and failure rate of EPike and Pike on a set of hand-crafted test cases




Figure 11: Success rate and failure rate of EPike and Pike on randomly generated test cases

means that if every agent prefers a noop action, then the predict node has a score of 0, i.e. execution is stuck. Note that in reality, agents may decide to act if nobody else does instead of waiting forever. We do not take into account such interactive behavior, but assume this is a reasonable way to approximate the utility of the predict node.

Implicit Belief Revision We consider explanations to be an *explicit* way of revising others’ beliefs. We consider it an *implicit* belief revision when an unexpected execution action or intent announcement action causes a less plausible world of an agent to be promoted to be a most plausible world. We assume agents do not wish to surprise others and penalize an action that causes implicit belief revision with a score of 0. This makes sure that our agent always explains its action before taking it if it is not expected by others. However, during execution, implicit belief revision may still occur, such as when others take an action unexpected by our agent which revises our agent’s belief.

Performance Optimization The performance of our algorithm largely depends on the speed of solving constraint satisfaction problems (CSPs) from the knowledge bases. To optimize the performance, we implement incremental checking and caching for CSPs, since the CSPs are largely similar throughout the process. At the decision node, we lazily expand the actions in the order of execution actions, noop action, then communication actions. For example, communication actions do not need to be expanded if higher-priority actions have a better score than the maximum possible score for a communication action due to penalty.

Experiment Results

We describe our experiment results on the success rate and the scalability of our algorithm EPike compared to Pike (Levine and Williams 2018). We implemented our own Pike as a naive version of EPike that assumes what it believes is believed by all, that is, it may falsely assume common knowledge when there is not. Note that the original Pike

supports some additional functionalities not accounted for by this paper, such as scheduling. We use z3 as our CSP solver (Moura and Bjørner 2008). We use an exploration parameter of 4 for SEARCHACTION, and $\sqrt{2}$ for the rest. We use a horizon of $k = 3$ for SEARCHACTION. We limit our focus to a 2-agent team. The experiments are run by instantiating two EPike or Pike agents that execute together in a task. We measure the runtime in seconds for one agent being the ego agent, who we assume gets to be the first agent to act if it decides to after each action is taken.

Success Rate Since MCTS is an anytime algorithm, we evaluate the success rate and failure rate of EPike and Pike, under different timeout in seconds for MCTS (or if it reaches 1000 iterations of simulations, whichever comes first). The experiments are run for the domains of (1) **Breakfast**, which includes variations of our motivating example, (2) **Word Puzzles**, (3) **Search-and-Rescue (SAR)**, (4) Randomly generated sequential tasks. We run each hand-crafted test case for 20 times for both Pike and EPike with no timeout, with results shown in Figure 10. We generate random test cases that vary in the size of the task (number of variables V) and the number of constraints that agents differ on, ranging from [0, 3], for 10 cases per condition, and report the result after running each case for 2 times for both Pike and EPike under different timeout, as shown in Figure 11. Note that it is possible for execution to neither succeed nor fail, in which case execution hangs as no agents plan to act. This could be because (1) MCTS algorithm is stopped by the timeout before it finds a feasible next step, (2) EPike believes execution is bound to fail no matter its action, such as when the other agent would not trust its explanation, or (3) EPike falsely believes that the other agent will act. In practice, we can adopt mitigations such as allowing EPike to take the next best action after having waited for a long time.

From Figure 11, we see that as timeout increases, EPike’s success rate increases, especially for larger-sized tasks, and is higher than Pike’s success rate given enough time. Meanwhile, its failure rate is consistently low and always lower than Pike. This shows that EPike is conservative, and when it does not succeed, it is mainly because it has not found a good action to take within the timeout, but it does not take an incorrect action rashly as Pike tends to do. This is consistent with the result of the hand-crafted test cases in Figure 10.

Scalability To see how EPike scales, we run the MCTS algorithm for a fixed number of 500 iterations to see how long it takes to reach a certain level of certainty under different model parameters, such as the size of the initial plausibility model (embodied by the number of constraints that agents differ on, *Diff* shown by hue of the plot), the size of




Figure 12: Runtime of EPike and Pike on random tests

the task (the number of variables V , *Num Variables*), concurrency level of the actions (the number of ordering constraints O , *Num Orders*), and the number of constraints C in the task (*Num Constraints*). We measure the average runtime for each callback for Pike and EPike. As shown in Figure 12, runtime for EPike is heavily affected by how much agents’ beliefs differ, and it also increases as the task size increases. Pike takes less time than EPike, as expected, but when under common knowledge, EPike’s runtime is closer to Pike and can also finish relatively quickly.

Related Work

Human-Robot Teaming Our work is related to human-robot teaming as it considers the collaborative process of humans and robots working together to achieve tasks. Work in this field focuses primarily on recognizing and adapting to humans’ intent, and in some cases, communicating about intent, which we inherit in our work. Pike inspired us to take a constraint-based approach for concurrent intent recognition and adaption, in which a library of precompiled plans are encoded in a knowledge base (Levine and Williams 2018). Pike is later extended to a probabilistic setting, called *Riker* (Levine 2019), where the robot can further ask the human about their intent (Broida 2021). Other than constraint-based approaches, work has been proposed using techniques from classical planning and MDP, such as PReTCIL (Freedman and Zilberstein 2017) that uses probabilistic plan recognition and NOPA (Puig et al. 2023) that leverages inverse planning for goal recognition. In (Unhelkar, Li, and Shah 2020b,a), a human behavior model is learned through semi-supervised learning and incorporated into the robot’s POMDP model that supports bi-directional communication on intent. However, most work assumes common knowledge of the task, as opposed to implementing an explicit Theory of Mind.

Epistemic Planning In the field of epistemic planning, there are two main categories of approaches – the semantic approach based on Dynamic Epistemic Logic (DEL) (Bolander and Andersen 2011; Le et al. 2018; Fabiano et al. 2020) and the symbolic approach (Muise et al. 2015). Our work leverages the DEL approach and carries over their insight on how to model announcement and question-asking actions. In particular, we take an implicit coordination approach, following from the work of Engesser et al., where the agent takes into account the spontaneous cooperation of other agents in achieving the goal, which requires recursive perspective-taking in order to predict their actions (Engesser et al. 2017). In (Bolander et al. 2018), the authors further dis-

cussed the impact of eager and lazy agents in the framework, and in (Reifsteck et al. 2019), an MCTS algorithm is developed that shares similar insights as our work. Compared to the work by Engesser et al., we differ in that our framework based on conditional doxastic logic allows the modeling of false beliefs and the revision of false beliefs, and our explanations refer directly to the plan space instead of states as a result of extending the logic to knowledge bases.

XAIP Our work is related to Explainable AI in Planning (XAIP), especially to the work on plan explanations taking into account the differences in agents’ mental models. In (Chakraborti et al. 2017), model reconciliation is proposed that allows robots to explain the model differences upon misalignment between the human’s mental model of the robot and the robot’s actual model. In (Vasileiou et al. 2022), a logic-based approach to model reconciliation is proposed, where the planning problem is encoded as a SAT problem using SatPlan, and the model differences are computed with respect to the human and the robot’s knowledge bases. Since these approaches consider the entire PDDL planning model, plan explanations go beyond explaining about the differences in the initial states but can also be about agents’ discrepancies in goal states and action models. In (Chakraborti, Sreedharan, and Kambhampati 2019), model reconciliation is balanced with explicable planning, which allows robots to find (potentially sub-optimal) plans that are expected by the human based on the human’s understanding of the robot (Zhang et al. 2017), and in (Sreedharan et al. 2020), the two are unified in an expectation-aware planning framework with additional explanatory actions. This inspired us in thinking about how the robot can balance its adaptation and communication with the human. However, most of their work considers humans as observers without much human-robot cooperation. In (Zahedi, Sreedharan, and Kambhampati 2022), the authors pointed out the importance of a richer mental modeling framework that allows human-robot collaboration, which we provide a viable way of filling the gap.

Another line of work from Shvo et al. (Shvo, Klassen, and McIlraith 2020) provides explanations by considering agents’ Theory of Mind represented using epistemic logic. In particular, to resolve the human’s misconceptions about plans, a symbolic epistemic planner RP-MEP (Muise et al. 2015) is used for the robot to either take actions to align the true state to the human’s belief or explain the true state to the human (Shvo, Klassen, and McIlraith 2022). However, their explanations are also on states rather than plans.

Conclusion

In this work, we combine insights from epistemic logic and knowledge-base encoding of plans to allow agents to understand discrepancies in their beliefs of feasible plans. We develop an online execution algorithm Epistemic Pike for the agent to dynamically plan its actions to adapt to others and communicate to resolve any discrepancy. We show that our agent is effective in working in teams where a shared mental model of plans cannot be guaranteed. A natural next step is to consider cases where actions are partially observable.

Acknowledgements

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001120C0035. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA).

References

- Baltag, A.; and Smets, S. 2008. A qualitative theory of dynamic interactive belief revision. *Logic and the foundations of game and decision theory (LOFT 7)*, 3: 9–58.
- Bolander, T.; and Andersen, M. B. 2011. Epistemic planning for single-and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1): 9–34.
- Bolander, T.; Engesser, T.; Mattmüller, R.; and Nebel, B. 2018. Better eager than lazy? How agent types impact the successfulness of implicit coordination. In *Sixteenth International Conference on Principles of Knowledge Representation and Reasoning*.
- Broida, J. 2021. *Active Policy Querying for Dynamic Human-Robot Collaboration Tasks*. Master’s thesis, Massachusetts Institute of Technology.
- Chakraborti, T.; Sreedharan, S.; and Kambhampati, S. 2019. Balancing Explicability and Explanations in Human-Aware Planning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*.
- Chakraborti, T.; Sreedharan, S.; Zhang, Y.; and Kambhampati, S. 2017. Plan Explanations as Model Reconciliation: Moving Beyond Explanation as Soliloquy. In *International Joint Conference on Artificial Intelligence*.
- Engesser, T.; Bolander, T.; Mattmüller, R.; and Nebel, B. 2017. Cooperative Epistemic Multi-Agent Planning for Implicit Coordination. In *M4M@ICLA*.
- Fabiano, F.; Burigana, A.; Dovier, A.; and Pontelli, E. 2020. EFP 2.0: A Multi-Agent Epistemic Solver with Multiple E-State Representations. In *International Conference on Automated Planning and Scheduling*.
- Freedman, R. G.; and Zilberstein, S. 2017. Integration of Planning with Recognition for Responsive Interaction Using Classical Planners. In *AAAI Conference on Artificial Intelligence*.
- Katsirelos, G.; and Bacchus, F. 2005. Generalized nogoods in CSPs. In *AAAI*, volume 5, 390–396.
- Le, T.; Fabiano, F.; Son, T. C.; and Pontelli, E. 2018. EFP and PG-EFP: Epistemic Forward Search Planners in Multi-Agent Domains. In *International Conference on Automated Planning and Scheduling*.
- Levine, S. J. 2019. *Risk-bounded coordination of human-robot teams through concurrent intent recognition and adaptation*. Ph.D. thesis, Massachusetts Institute of Technology.
- Levine, S. J.; and Williams, B. C. 2018. Watching and acting together: Concurrent plan recognition and adaptation for human-robot teams. *Journal of Artificial Intelligence Research*, 63: 281–359.
- Moura, L. d.; and Bjørner, N. 2008. Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*, 337–340. Springer.
- Muise, C.; Belle, V.; Felli, P.; McIlraith, S. A.; Miller, T.; Pearce, A. R.; and Sonenberg, L. 2015. Planning Over Multi-Agent Epistemic States: A Classical Planning Approach. In *AAAI Conference on Artificial Intelligence*.
- Puig, X.; Shu, T.; Tenenbaum, J. B.; and Torralba, A. 2023. NOPA: Neurally-guided Online Probabilistic Assistance for Building Socially Intelligent Home Assistants. In *International Conference on Robotics and Automation*.
- Reifsteck, D.; Engesser, T.; Mattmüller, R.; and Nebel, B. 2019. Epistemic multi-agent planning using monte-carlo tree search. In *KI 2019: Advances in Artificial Intelligence: 42nd German Conference on AI*.
- Shvo, M.; Klassen, T. Q.; and McIlraith, S. A. 2020. Towards the role of theory of mind in explanation. In *International Workshop on Explainable, Transparent Autonomous Agents and Multi-Agent Systems*, 75–93. Springer.
- Shvo, M.; Klassen, T. Q.; and McIlraith, S. A. 2022. Resolving misconceptions about the plans of agents via Theory of Mind. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 32, 719–729.
- Sreedharan, S.; Chakraborti, T.; Muise, C.; and Kambhampati, S. 2020. Expectation-aware planning: A unifying framework for synthesizing and executing self-explaining plans for human-aware planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Unhelkar, V. V.; Li, S.; and Shah, J. A. 2020a. Decision-making for bidirectional communication in sequential human-robot collaborative tasks. In *2020 15th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 329–341. IEEE.
- Unhelkar, V. V.; Li, S.; and Shah, J. A. 2020b. Semi-supervised learning of decision-making models for human-robot collaboration. In *Conference on Robot Learning*, 192–203. PMLR.
- Van Ditmarsch, H.; and Kooi, B. 2008. Semantic results for ontic and epistemic change. *Logic and the foundations of game and decision theory (LOFT 7)*, 3: 87–117.
- Van Ditmarsch, H.; van der Hoek, W.; Halpern, J. Y.; and Kooi, B. 2015. *Handbook of epistemic logic*. College Publications.
- Vasileiou, S. L.; Yeoh, W.; Son, T. C.; Kumar, A.; Cashmore, M.; and Magazzeni, D. 2022. A Logic-Based Explanation Generation Framework for Classical and Hybrid Planning Problems. *Journal of Artificial Intelligence Research*, 73: 1473–1534.
- Zahedi, Z.; Sreedharan, S.; and Kambhampati, S. 2022. A Mental-Model Centric Landscape of Human-AI Symbiosis. *CoRR*, abs/2202.09447.
- Zhang, Y.; Sreedharan, S.; Kulkarni, A.; Chakraborti, T.; Zhuo, H. H.; and Kambhampati, S. 2017. Plan explicability and predictability for robot task planning. In *2017 IEEE international conference on robotics and automation (ICRA)*, 1313–1320. IEEE.