

# Informative Communication of Robot Plans

Michele Persiani<sup>1</sup>[0000-0001-5993-3292] and Thomas  
Hellström<sup>1</sup>[0000-0001-7242-2200]

Umeå University, Umeå, Sweden  
michelep@cs.umu.se, thomas.hellstrom@umu.se

**Abstract.** When a robot is asked to verbalize its plan it can do it in many ways. For example, a seemingly natural strategy is incremental, where the robot verbalizes its planned actions in plan order. However, an important aspect of this type of strategy is that it misses considerations on what is effectively informative to communicate, because not considering what the user knows prior to explanations. In this paper we propose a verbalization strategy to communicate robot plans informatively, by measuring the information gain that verbalizations have against a second-order theory of mind of the user capturing his prior knowledge on the robot. As shown in our experiments, this strategy allows to understand the robot's goal much quicker than by using strategies such as increasing or decreasing plan order. In addition, following our formulation we hint to what is informative and why when a robot communicates its plan.

**Keywords:** Plan Verbalization; Human-Robot Interaction; Bayesian Network; Mirror Agent Model

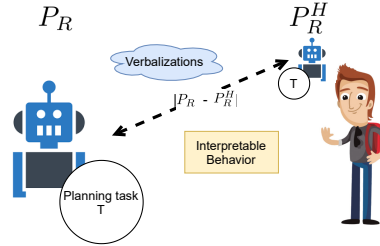
## 1 Introduction

With its ever-growing advancements, Artificial Intelligence (AI) is proving to be a promising partner in our lives. In envisioning this relationship, AI agents should support us by being able to perform their assigned tasks both efficiently and accurately, but, on top of that, a crucial aspect for this relationship to be successful is the degree by which agents are able to make themselves understood by their human users, either through explanations, or by behavior that is interpretable from the human perspective [6].

An interpretable behavior is a behavior that is expected by the user, because fulfilling his expectations about the agent [3, 2]. Once implemented in a model, these expectations describes how the user is modeling the agent, therefore informing the agent about how it is being perceived and explained. It follows that behaviors fitting this expectations model are intrinsically explainable and therefore should not require additional explanations; however, other types behaviors could require explanations. This is because, while fulfilling other relevant properties such as being optimal, do not fit the user's expectations, and thus result

being uninterpretable. It becomes therefore crucial to complement such behavior with explanations targeting its alignment with the user’s expectation model, such that it becomes interpretable in it [3].

A relevant measure of the explanation process is its degree of *informativeness*, that is the amount of information that explanations transfer to the explainee [9]. To be informative, explanations can leverage the user’s expectation model on the agent when computing what to communicate and how: given that many candidate verbalizations are possible to communicate, the agent has to determine which one is most informative, in the sense that it increases by a largest amount the user’s understanding of the agent. We see the process of explanation as having a complementary role with respect to interpretable behavior, and explanations and interpretable behavior have the same objective of keeping the agent model  $P_R$ , and the expectations that the user has about the agent  $P_R^H$ , as similar as possible. This is because, when aligned, this alignment signals that the agent is being understood because fitting the expectations.



**Fig. 1.**  $P_R$ : model of the robot working on task  $T$ .  $P_R^H$ : second-order theory of mind of the robot. The goal of interpretable behavior and verbalizations is to reduce the distance  $|P_R - P_R^H|$ . Here we consider the minimum amount of actions, and their order, that the robot should communicate such that  $|P_R - P_R^H| < \theta$ .

Informativeness in explanations is still sparsely considered inside the literature, and previous work on explanations focuses more on the consistency of explanations from the agent’s perspective, rather than generating explanations that are informative to the user [4, 12]. With the goal of addressing this gap, in this paper we propose an algorithm to communicate a robot’s plan using a criteria based on the informativeness of communications. We achieve this by firstly grounding the understanding that the user has about the agent as a measure of difference between agent model and user expectation model. Later, we select which actions of the plan to communicate, and their order, based on how their communication affect this distance measure. In addition, by the same concepts we define what is informative in the case of explanations of robot plans.

The rest of the paper is structured as follows. In Section 2 we provide brief, relevant background in theory of mind for robotics and verbalizations of robot plans. In Section 3 we present our method for informative communication of the robot plan. Section 5 describes experiments and measures performed on different tests. Finally, Section 7 and Section 8 contains discussion on the results and conclusions.

## 2 Background

Previous research in understandable robots [6] propose to consider the concept of *understanding* as a distance measure between a robot’s intentional model  $P_R$  and the user’s model of  $P_R$ ,  $P_R^H$ . The more similar these models are, the higher is the degree of understanding between user and robot. The limit case where the distance is 0 signals the robot that it is being fully understood in the sense that the user’s expectations of it are the same as its actual intentional model, while a measure  $> 0$  signals the robot that some parts of its intentional model are mismatching with the user’s expectation, and therefore require additional verbalizations. Here we see verbalizations as sequences of communicative actions<sup>1</sup> aimed at communicating parts of the intentional model. While having the same function of interpretable behavior, that is to reduce  $|P_R - P_R^H|$ , they are purely communicative in nature, such as through speech, text, images, etc.

The model  $P_R^H$  constitutes a second-order theory of mind [6], and shown in recent literature, can take many forms. For example, in [18] and in [5] it is a context-dependent label describing whether the user understands the robot. After a training phase through human annotations, predicting the label’s value allows to augment the cost function of the planning procedure with explainability awareness. Alternatively, more structured forms are for example presented in [3], where the theory of mind is a complete planning instance.

The work on understandable robots has connections also to research in interpretability and explainability. Given that a model of the user’s expectations  $P_R^H$  is provided to the robot (Fig. 1), interpretable behaviors keep the distance measure  $|P_R - P_R^H|$  low by modifying its behavior, while explanations through sequences of communicative actions (that we refer to as verbalizations) [7, 16].

In this paper we focus on explanations, and in particular on the verbalizations of the robot’s task experience, such as its recent course of action, its goal, or plan. For this type of question, the strategies to verbalize explanations that are proposed in the literature are often incremental i.e. orderly from the first to the last actions. For example, [15, 11] propose strategies to verbalize robot plans. In addition, the authors keep into consideration possible categories of users by predefining, for a given plan, multiple types of explanations, one for every type of user. The resulting explanations span over standardized dimensions of verbalization such as abstraction, locality, and specificity. However, while these strategies can consider different types of users and what can be informative to them, they don’t consider the information that users (in any given category) possess prior to the explanations, and therefore verbalizations can only be incremental, or with hand-coded strategies. [19] addresses this problem by a manual approach: we can ask the users which elements of the plan are of interest before its verbalization, in this way, the successive incremental verbalization can be filtered to contain only elements that are relevant to the specific user.

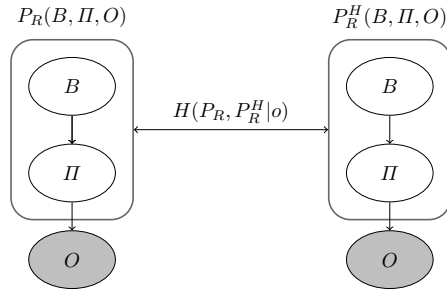
---

<sup>1</sup> A communicative action is an action, either verbal or physical, intended to decrease the models distance  $|P_R - P_R^H|$  [6]

As we show in this paper the order by which actions are verbalized can be controlled to maximize the amount of information that verbalizations produce inside the second-order theory of mind. A similar idea is explored in [17], where the authors propose how the explanation process can be re-ordered by keeping into consideration the cognitive load of the user. While we also confront the question of ordering verbalizations, in this paper we are interested in making the user understand the robot’s internal state with a minimum amount of utterances.

### 3 Method

We consider the case in which the robot’s model  $P_R$  and user’s estimated model of the robot model,  $P_R^H$ , are probabilistic Belief-Desire-Intention (BDI) models, described by two Equivalent Bayesian Networks and forming a second-order theory of mind setting (Figure 2). The networks use the same random variables, however, these variables can be differently distributed in  $P_R$  and  $P_R^H$ , thus possibly reflecting a mismatch between the user’s beliefs over the robot model and the true robot model. We refer to this setting as the *Mirror Agent Model* because defining the second-order theory of mind of the agent as a “mirror” of the agent’s true model. This model is easy to deploy and can be obtained *as a service* [1] on a variety of agents, which is an important property for its usability.



**Fig. 2.** Robot model ( $P_R$ ) and second-order theory of mind of the user ( $P_R^H$ ) as equivalent Bayesian Networks. The cross-entropy  $H$  measures the difference between robot’s and user’s estimated state of mind a posteriori of a verbalization  $o$ .

A setting which mirrors  $P_R$  and  $P_R^H$  has clear advantages when the robot decides what to communicate in order to increase the models similarity: if we are able to directly compare the robot and its theory of mind, we could immediately know what part of the model is being understood, and which part instead requires a verbalizations. However, the assumption this setting makes is that the user has internalized a robot model of the same form of the original. For this paper we take this assumption to hold, and model the user as knowing the exact robot model<sup>2</sup>. The robot’s probabilistic BDI model is defined  $P_R(O, II, B) = P_R(O|II)P_R(II|B)P_R(B)$  where every  $b \in B$  include both belief and goal of

the robot,  $\pi \in II$  a possible intention (we consider intentions as plans) and  $o \in O$

<sup>2</sup> This assumption doesn’t make the proposed method to lose generality, because from two different  $P_R$  and  $P_R^H$  we can always create a super-agent which comprehend both models, and that is used in the mirror model. Then, we would set inside  $P_R$  and  $P_R^H$  the corresponding probabilities to 0.

the observations that plans yield to the user. Since we here consider the case where the robot verbalizes its plans, observations are verbalizations of actions, that in our case have the form of textual descriptions of sequences of actions. The only observable random variable is  $O$ , while  $B$  and  $\Pi$  are variables internal to the robot's *state of mind* [6], and can be interacted with only through observations.  $P_R^H$  is similarly defined.

As a simplified case, we don't consider how the robot forms and updates its beliefs through its sensors, and we assume a fixed probabilistic belief  $P_R(B)$ . The robot can form plans using the planning model  $P_R(\Pi|B)$ , where the probability of a plan  $P_R(\pi|b)$  is  $> 0$  only if the plan is consistent with both belief and goal condition of  $b$ . While computing plans we model the robot to be a rational agent ie. associating higher probabilities to cheap plans, and accordingly set  $P_R(\pi|b) \propto R(\pi, b)$  where  $R(\cdot)$  is a function measuring the rationality of  $\pi$  when executed in  $b$  [14]. The model  $P_R(O|\Pi)$  instead describes how executing plans create observations.  $P_R(o|\pi) > 0$  means that  $o$  verbalizes a set of actions  $o = \{a_1, \dots, a_n\}$  that are a subset of the plan  $\pi$ , with  $P_R(o|\pi) = \prod_i P_R(a_i|\pi)$ ,  $P_R(a_i|\pi) > 0$  iff  $a_i \in \pi$ . The model  $P_R^H$  is similarly defined. The cross-entropy measure  $H(P_R, P_R^H)$  determines the similarity between  $P_R$ 's and  $P_R^H$ 's states of mind (ie. their corresponding intentions and beliefs). For simplicity, we consider the case of a deterministic  $P_R$ , with  $P_R(B = b_R) = 1$  and  $P_R(\Pi = \pi_R) = 1$ , and probabilistic  $P_R^H$ . In this setting, the cross-entropy between  $P_R$  and  $P_R^H$  is:

$$H(P_R, P_R^H) = \sum P_R(\pi, b) \log P_R^H(\pi, b) = -\log P_R^H(\pi_R|b_R) - \log P_R^H(b_R) \quad (1)$$

Therefore, to increase the similarity between  $P_R$  and  $P_R^H$  we must increase the probability of  $\pi_R$  and  $b_R$  in the theory of mind; the limit case where also the theory of mind deterministically produces  $\pi_R$  and  $b_R$  is the case in which the user is estimated to have an exact understanding of the robot. In this case the cross-entropy is 0.

### 3.1 Informative Communication

As described previously, the main objective of a verbalization  $o \in O$  is to make  $P_R^H$  as similar as possible to  $P_R$ . This is achieved by decreasing the cross-entropy measure between the models. The associated Information Gain (IG) of the verbalization is  $IG(P_R, P_R^H, o) = H(P_R, P_R^H) - H(P_R, P_R^H|o)$ . Using the information gain, the most informative verbalization can be selected as the one with maximal information gain:

$$\begin{aligned} \hat{o} &= \arg \max_{o \in O} H(P_R, P_R^H) - H(P_R, P_R^H|o) \\ &= \arg \max_{o \in O} \log P_R^H(\pi_R, b_R|o) \\ &= \arg \max_{o \in O} \log P_R^H(o|\pi_R) - \log \mathbb{E}[P_R^H(o|\pi)] \end{aligned} \quad (2)$$

Two important properties of informative communication emerge from Eq. 2. The first is that only verbalizations that correctly reflect the model of the robot are informative. Communicating an action not belonging to the robot’s plan, with  $P_R^H(a_i|\pi_R) = 0$ , induce an informativeness of  $-\infty$  of the full verbalization. The second is that actions that are informative to communicate are those that tend to appear exclusively in  $\pi_R$ .

A simple way to find which is the best verbalization of  $N$  actions to communicate is to enumerate all the possible combinations of actions belonging to  $\pi_R$ , to then select the combination with highest IG. Algorithm 1 implements this procedure. The output of the algorithm is the sequence of the  $N$  most informative actions ordered in plan order.

---

**Algorithm 1** Find the most informative communication of size  $N$ , by enumerating and sorting the combinations of actions of size  $N$

---

```

1: procedure VERBALIZE-PLAN( $P_R, P_R^H, N$ )
2:    $\pi_R \leftarrow \text{PLAN}(P_R)$ 
3:    $o \leftarrow \text{FIND-MOST-INFORMATIVE}(\pi_R, P_R^H, N)$ 
4:    $o_{sorted} \leftarrow \text{PLAN-SORT}(o)$ 
5:    $\text{VERBALIZE}(o_{sorted})$ 
6: end procedure
1: procedure FIND-MOST-INFORMATIVE( $\pi_R, P_R^H, N$ )
2:    $C \leftarrow \text{COMBINATIONS}(\pi_R, N)$   $\triangleright |C| = \binom{|\pi_R|}{N}$ 
3:    $Q \leftarrow \emptyset$ 
4:   for  $c \in C$  do
5:      $h_c = \text{INFORMATION-GAIN}(P_R^H | o = c)$   $\triangleright \text{Eq. 2}$ 
6:      $\text{APPEND}(Q, \langle h_c, c \rangle)$ 
7:   end for
8:    $Q \leftarrow \text{SORT}(Q)$ 
9:    $h_{best}, c_{best} \leftarrow \text{POP}(Q)$ 
10:  yield  $c_{best}$ 
11: end procedure

```

---

## 4 Implementation in PDDL

We implemented the probabilistic BDI models  $P_R(O, \Pi, B)$  and  $P_R^H(O, \Pi, B)$  by specifying planning instances using the Planning Domain Description Language (PDDL). PDDL [10] is a standard language to specify planning domains for what is usually referred to as classical planning. A planning instance for the robot is obtained by specifying the tuple  $\langle \mathcal{P}_R, \mathcal{A}_R, I_R, \mathcal{G}_R, \mathcal{O}_R \rangle$ . Where  $I_R$  and  $\mathcal{G}_R$  are set of ground predicates and correspond to the initial and goal state respectively,  $\mathcal{O}_R$  is the set of objects available to ground the predicates  $\mathcal{P}_R$ , while  $\mathcal{A}_R$  is the set of available actions to transition between states. Similarly, the second order theory of mind model has components  $\langle \mathcal{P}_R^H, \mathcal{A}_R^H, I_R^H, \mathcal{G}_R^H, \mathcal{O}_R^H \rangle$ .

We set the descriptive components of the planning instances of robot and theory of mind to be equivalent. ie.  $\mathcal{P}_R = \mathcal{P}_R^H, \mathcal{A}_R = \mathcal{A}_R^H$  and  $\mathcal{O}_R = \mathcal{O}_R^H$ ,

with the only probabilistic parts being  $I_R$ ,  $\mathcal{G}_R$ ,  $I_R^H$  and  $\mathcal{G}_R^H$ . In this setting, the corresponding probability distribution over the possible PDDL instances describing the robot state is obtained by a combination of Bernoulli distribution for the beliefs  $I_R$ , and a categorical distribution for the possible goals  $\mathcal{G}_R$  (the same for  $I_R^H$  and  $\mathcal{G}_R^H$  respectively).

$$\begin{aligned} P_R(B) &= P_R(I)P_R(G) & P_R(G; \theta_R) &= P(G|\{g_0, \dots, g_m\}) \\ P_R(I; \theta_R) &= \prod_i P(p_i \in I_R; \theta_{p_i}) & P_R(G = \langle \mathcal{G}_j \rangle | \{g_0, \dots, g_m\}) &= \theta_j \\ P(p_i \in I_R) &= \theta_i & \sum_j \theta_j &= 1 \end{aligned}$$

where  $\theta$  are the distribution parameters. Sampling a belief from  $P_R(B)$  yields an initial state and a goal state for the PDDL planner. The planning model  $P_R(\Pi|B)$  is instead implemented by a planner of choice compatible with the underlying PDDL requirements. The probability of a plan  $P_R(\Pi = \pi|B = b)$  is defined as a function of rationality [14]. Sampling from the planning model can for example be done through Diverse Planning techniques [8].  $P_R^H(B)$  and  $P_R^H(\Pi|B)$  are similarly defined.

## 5 Experiments

We test our algorithm for informative communication over two sets of tests. First, we run a set of automated benchmarks on the PUCRS dataset, which is a curated dataset of planning domains [13], discussing then some generally valid metrics and results. After that we show the results of a user study simulating the case where an operator meets a robot working in a warehouse, and not knowing what it is doing, asks for its plan.

### 5.1 Tests on the PUCRS dataset

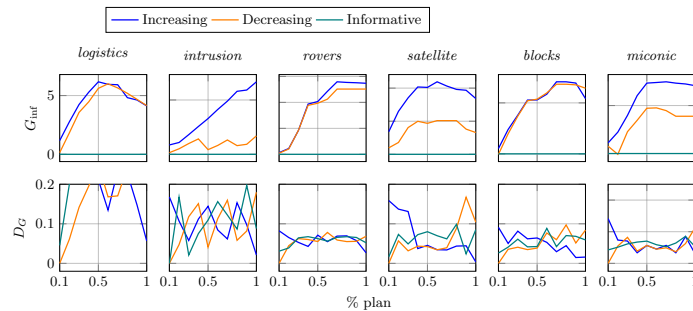
We benchmarked Algorithm 1 on the PUCRS dataset [13], that is a dataset of PDDL planning domains and problems, with the goal to search for properties of informative communication that generalizes across domains. We selected 6 standard PDDL domains: *logistics*, *intrusion-detection*, *rovers*, *satellite*, *blocks-world*, *satellite*. Table 1 shows relevant average measures of the domains, such as number of predicates and operators, size of the initial state, number of tested goals and length of optimal plans.

For every domain we perform the tests on 10 planning instances, averaging the results. In each test  $P_R$  is initialized by randomly selecting a goal from the pool of available goals and using the original instance’s initial state.  $P_R^H$  is initialized by randomly selecting 3 goals from the goal pool plus the goal of  $P_R$ , for a total of 4 possible goals. These goals are used with equal probabilities in  $P_R^H(G)$ . In addition, 4 random predicates belonging to the robot plan’s preconditions are made as probabilistic ie.  $P_R^H(I) = \prod_{i \in 1..4} P(p_i)$ , with  $P(p_i = 1) = 0.5$ .

Two baseline communication strategies are tested alongside informative communication: increasing order, which communicates actions in plan order, and decreasing order, which communicates the actions in reversed plan order. Figure 3 shows the average measures of the benchmark for the tested strategies. On the first row is shown the gain on entropy obtained by using the informative communication strategy rather than the others, measured as  $G_{\text{inf}} = H(P_R, P_H^R) - H_{\text{inf}}(P_R, P_H^R)$ . The second row instead shows the distance measure  $D_G(a_i)$  between the predicates affected by a

communicated action  $a_i$  and the goal predicates of the instances.  $D_G(a_i) = \frac{k}{|\pi_R|}$  means that, along the plan and starting counting from  $a_i$ , the first action directly affecting the predicates of the instance’s goal is the  $k$ -th, eg.  $D_G(a_i) = 0$  if  $a_i$  has in its effects a predicate contained in the instance’s goal.

The plots highlight an additional relevant property of the informative strategy, that is a consistent low distance-to-goal measure for the actions communicated earliest. This means that the informative strategy tends to communicate earlier the actions that directly affect the predicates of the goal. If we think about it, these actions are clearly informative, because communicating the goal predicates (namely, communicating the goal) of the plan clearly mostly disambiguates the goal from the others. This order of communication is similar to the decreasing strategy in its early actions, and as an average, after communicating the first 10% of actions (that roughly correspond to one or two depending on the domain) the informative and the decreasing strategies are equally informative ( $G_{\text{inf}} = 0$ ). Notice also that because of this reason the decreasing strategy is consistently more informative than the increasing strategy.



**Fig. 3.** Entropy gains and distance-to-goal measures for the automated benchmarks.

Domain	$ \mathcal{P} $	$ \mathcal{A} $	$ I $	$ G $	$ \bar{\pi} $
<i>miconic</i>	8	4	517	6	35.57
<i>logistics</i>	3	6	22.16	10.39	24.41
<i>intrusion</i>	11	9	1	16.67	13.07
<i>blocks-world</i>	5	4	14.58	20.28	14.50
<i>rovers</i>	32	9	172.71	6	24.93
<i>satellite</i>	12	5	58	6.43	16.89

**Table 1.** Average instance measures over the tested planning domains. The columns, from left to right are: number of operators, number of predicates, size of the initial state, number of tested goals, length of optimal plans.

## 6 User Study

To further test our method, we performed a user study that simulates the case in which an operator encounters a robot in a corridor of a warehouse, and not knowing what it’s doing asks it to verbalize its plan. The study’s participants were given 12 scenarios of this type, illustrated on a Graphical User Interface (Figure 3). In each of the scenarios there are two rooms connected by a corridor and containing 1 to 3 random objects of the same color. Objects can be circles, squares or triangles, and can be either red or blue. There are 3 possible exit doors at the corners. A recharge station is in the central corridor. The robot has 3 possible actions: moving from corridors and rooms, grabbing objects and recharging at the recharge station. In every scenario the robot had the goal of grabbing two random objects from those available, and exit the floor from a random door. It was also set to need to recharge at the recharge station in half of the scenarios. Both participants and robot had full observability on the initial state of the scenario, with robot position and available objects. However, the participants weren’t informed on the belief of the robot, its goal or its plan.

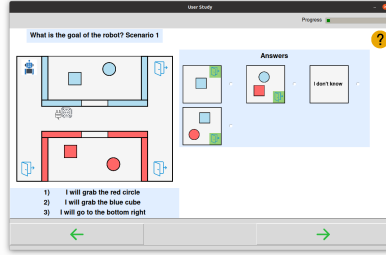
Increasing order	Decreasing order	Informative order
I will navigate to the central corridor.	I will exit from the bottom door.	I will grab a red circle.
I will reach the warehouses doors.	I will navigate to the right corridor.	I will grab a blue cube.
I will enter the blue warehouse.	I will reach the warehouses doors.	I will exit from the bottom door.
I will grab a blue cube.	I will grab a red circle.	I will enter the blue warehouse.
I will reach the warehouses doors.	I will enter the red warehouse.	I will enter the red warehouse.
I will enter the red warehouse.	I will reach the warehouses doors.	I will navigate to the central corridor.
I will grab a red circle.	I will grab a blue cube.	I will navigate to the right corridor.
I will reach the warehouses doors.	I will enter the blue warehouse.	I will reach the warehouses doors.
I will navigate to the right corridor.	I will reach the warehouses doors.	I will reach the warehouses doors.
I will exit from the bottom door.	I will navigate to the central corridor.	I will reach the warehouses doors.

**Table 2.** Verbalizations of a plan from the user study, sorted by the increasing, decreasing and informative strategies.

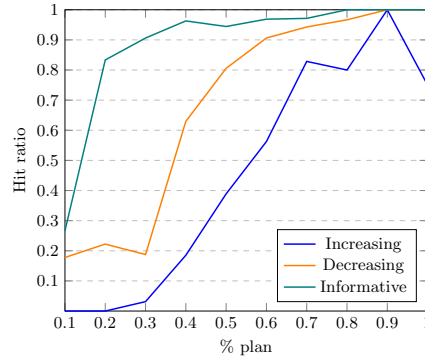
At every step of a scenario participants were given a picture of the scenario and a verbalization of the robot’s plan, in the form of its plan in textual form (see Table 2). The verbalization was of increasing size at each of the scenario’s steps (at the first step of size 1, second step of size 2, etc.), but, randomly for every scenarios, the verbalization could be given in one of the strategies *informative*, *increasing* and *decreasing*. Using picture and verbalization at every step the participant were asked to predict the robot’s goal among a set of 3 possible goals, or to state they didn’t know. Participants utilized the GUI to see the current scenario and verbalization, and to provide their prediction—for example, Figure 3 shows the participants view of a scenario after the third verbalization of the robot.

We conducted the experiments with 10 participants. The participants were given a thorough introduction to the domain and the task they were asked to perform. At any time they could ask any additional clarification. In the attempt of better simulating an online interaction, the participants were told that the experiment was going to be evaluated also in term of the time required to answer the questions, and that they should have been both accurate and fast. Figure 4

shows the average hit ratio of the participants as a function of the percentage of communicated plan. The hit ratio was measured as the number of times the correct goal was predicted divided by the number of predictions, averaged among all the participants. The average length of plans was 9.14 actions.



**Table 3.** Graphical user interface utilized during the user study. The interface shows the positions of robot, objects and doors (left). For every step of a scenario the participants could select a prediction about the robot’s goal (right) from a verbalization as a written sequence of actions (bottom).



**Table 4.** Hit ratio of the participants as a function of percentage of communicated actions.

The results highlight the effectiveness of the informative communication strategy among the participants. After communicating 20% of the plan, corresponding to 2 actions as an average, the hit ratio for the informative strategy was already measured as 0.85. In comparison, the decreasing strategy takes half of the plan to be as effective, while the increasing strategy almost all of the plan. The statistical significance of the informative strategy is tested against the null hypotheses  $H_{inc}$ : no difference with the increasing strategy ( $\mu_{inf} = \mu_{inc}$ ), and  $H_{dec}$ : no difference with the decreasing strategy ( $\mu_{inf} = \mu_{dec}$ ).  $\mu$  is the average earliest step at which the participants answered correctly, depending on the strategy used. The corresponding found p-values are:  $p(n \leq \mu_{inf} | H_{inc}) = 0.0014$  and  $p(n \leq \mu_{inf} | H_{dec}) = 0.04$ , showing statistical significance of the results.

## 7 Discussion

Our observations are that the most informative actions to communicate are the ones that, for a planning instance, directly affect the goal predicates, therefore likely at the end of the communicated plan. Inside a planning instance communicating the goal is what transfers the highest information because a large part of the plan can be inferred as a product of it. This makes sense to us because the last actions of a plan are the most expensive to commit to. Or, from a rationality perspective inside the observer’s expectation model (of which inference

of plans inside the models is based on), it would be irrational to see achieving a goal and then commit to another, excluding the case where the first goal is a sub-goal of the second. After the actions affecting the goal predicates have been communicated, the actions that mostly discriminate the beliefs become the most informative. This is reflected also in tests of the user study, where the informative strategy systematically selected as first actions to communicate the robot's exit point and the gathered objects.

Overall, the proposed informative communication works by communicating what most discriminates the robot's true state of mind among the set of other candidates, inside the estimated user's second-order theory of mind. The user study demonstrated that this strategy is effectively informative also for human participants. The proposed order of communicating plans based on informativeness, and formalized by the delivered algorithm, is mostly different from the seemingly natural way of communicating incrementally, that is also what is commonly used in the literature. Surprisingly, from our tests it results that the incremental strategy is the most inefficient among those being tested. To deepen the investigation, future work could probe whether human-human communication in task-driven scenarios follows more an informative rather than incremental strategy.

## 8 Conclusion

In this paper we proposed a strategy for the communication of robot plans that is based on the informativeness of verbalizations when contextualized inside a theory of mind. We compared the informative communication strategy with two baselines that followed increasing and decreasing plan order, finding that the informative strategy was effectively better at making users understand the robot's state using the least amount of verbalizations.

Our results further indicated that the communication strategy most commonly proposed in the literature, namely incremental in plan order, was the most inefficient at transferring information compared to the others. This is because the informative strategy (and partly the decreasing strategy) leverage the fact that the plan is grounded in a planning instance, and communicating as earliest key parts of it such as its goal allows the explainee to infer most of the remaining parts. This was also confirmed by a user study showing an effective correlation between informativeness and the capacity of participants to predict the robot's state.

There are several interesting directions for future work based on informative communication. A first direction is about a deepened investigation on how informative communication relates to human communication in task-driven scenarios. e.g. is it more natural to use an incremental or an informative strategy? Additionally, tests in a real human-robot interaction setting could be performed. And finally, additional tests using the proposed theory of mind model on more complex scenarios could also be performed.

## References

1. Cashmore, M., Collins, A., Krarup, B., Krivic, S., Magazzeni, D., Smith, D.: Towards explainable ai planning as a service. arXiv preprint arXiv:1908.05059 (2019)
2. Chakraborti, T., Kulkarni, A., Sreedharan, S., Smith, D.E., Kambhampati, S.: Explicability? legibility? predictability? transparency? privacy? security? the emerging landscape of interpretable agent behavior. In: Proceedings of the international conference on automated planning and scheduling. vol. 29, pp. 86–96 (2019)
3. Chakraborti, T., Sreedharan, S., Zhang, Y., Kambhampati, S.: Plan explanations as model reconciliation: Moving beyond explanation as soliloquy. arXiv preprint arXiv:1701.08317 (2017)
4. Gavriilidis, K., Carreno, Y., Munafo, A., Pang, W., Petrick, R.P., Hastie, H.: Plan verbalisation for robots acting in dynamic environments (2021)
5. Gong, Z., Zhang, Y.: Robot signaling its intentions in human-robot teaming. In: HRI Workshop on Explainable Robotic Systems (2018)
6. Hellström, T., Bensch, S.: Understandable robots-what, why, and how. *Paladyn, Journal of Behavioral Robotics* **9**(1), 110–123 (2018)
7. Kambhampati, S.: Synthesizing explainable behavior for human-ai collaboration. In: Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems. pp. 1–2 (2019)
8. Katz, M., Sohrabi, S.: Reshaping diverse planning. In: AAAI. pp. 9892–9899 (2020)
9. Lawless, W.F., Mittu, R., Sofge, D., Hiatt, L.: Artificial intelligence, autonomy, and human-machine teams—interdependence, context, and explainable ai. *Ai Magazine* **40**(3), 5–13 (2019)
10. McDermott, D.: Pddl-the planning domain definition language (1998)
11. Meza, I., Flores, J.G., Gangemi, A., Pineda, L.A.: Towards narrative generation of spatial experiences in service robots. In: IJCAI 2016 Workshop on Autonomous Mobile Service Robots (2016)
12. Moon, J., Magazzeni, D., Cashmore, M., Buksz, D., Lee, B.H., Moon, Y.S., Roh, S.H.: Towards explanations of plan execution for human-robot teaming. *CEUR Workshop Proceedings* (2019)
13. Pereira, R., Oren, N., Meneguzzi, F.: Landmark-based heuristics for goal recognition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 31 (2017)
14. Persiani, M., Hellström, T.: Inference of the intentions of unknown agents in a theory of mind setting. In: International Conference on Practical Applications of Agents and Multi-Agent Systems. pp. 188–200. Springer (2021)
15. Rosenthal, S., Selvaraj, S.P., Veloso, M.M.: Verbalization: Narration of autonomous robot experience. In: IJCAI. vol. 16, pp. 862–868 (2016)
16. Sreedharan, S., Kambhampati, S., et al.: Balancing explicability and explanation in human-aware planning. In: 2017 AAAI Fall Symposium Series (2017)
17. Zakershaharak, M., Marpally, S.R., Sharma, A., Gong, Z., Zhang, Y.: Order matters: Generating progressive explanations for planning tasks in human-robot teaming. arXiv preprint arXiv:2004.07822 (2020)
18. Zhang, Y., Sreedharan, S., Kulkarni, A., Chakraborti, T., Zhuo, H.H., Kambhampati, S.: Plan explicability and predictability for robot task planning. In: 2017 IEEE international conference on robotics and automation (ICRA). pp. 1313–1320. IEEE (2017)
19. Zhu, Q., Perera, V., Wächter, M., Asfour, T., Veloso, M.: Autonomous narration of humanoid robot kitchen task experience. In: 2017 IEEE-RAS 17th international conference on humanoid robotics (humanoids). pp. 390–397. IEEE (2017)