Robot Signaling its Intentions in Human-Robot Teaming

Ze Gong
CIDSE ASU <zgong11@asu.edu>

Yu Zhang
CIDSE ASU <yzhan442@asu.edu>

ABSTRACT

Facilitating a shared team understanding is an important task in human-robot teaming. In order to achieve efficient collaboration between the human and robot, it requires not only the robot to understand what the human is doing, but also the robot's behavior to be understood by the human. To address the latter problem, we propose an approach to signaling robot's intentions by generating natural language sentences without changing the plan itself. This is in contrast to recent approaches to generating explicable and legible plans. We formulate human interpreting robot's actions as a labeling process. Skip-chain Conditional Random Fields (CRFs) are used to capture the long dependencies of action labeling. In order to find out *when* and *what* to signal, possible timing and content of signaling are explored during the inference phase of skip-chain CRFs to find the ones that maximize the human understanding of the robot plan. In our preliminary result, we tested how signaling may help achieve better teaming by reducing the criticism on robot actions that may appear undesirable but is required. Evaluation with Amazon MTurk showed that is indeed the case.

KEYWORDS

Human-Robot Teaming, Intention Signaling and Projection

1 INTRODUCTION

A mutual understanding between the human and robot can largely increase trust as well as improve effectiveness and efficiency for teaming. There has already been lots of research work focusing on making robot actions comprehensible to humans. Legible motions [3], unambiguous natural language sentences [4, 6] and visual cues [1] can be generated to convey the objectives of the robot to human. For robot task planning, Zhang *et al.* [7] proposed two metrics which can be used to measure the explicability and predictability of a plan and in turn guide the robot to generate more explicable plans to reduce the cognitive load of its human teammate. Chakraborti *et al.* [2] formulated the explanation generation problem (for a robot to explain its actions) as a "model reconciliation problem". The goal of an explanation is to make the robot's plan optimal according to a changed human model.

Those methods, however, have to make modifications to the human model or the (optimal) plan generated by the robot. It may introduce an explicable plan that is too costly or an explanation that is too difficult to understand, and hence may not always be possible or efficient. Alternatively, a robot may always choose to signal every plan step to the human using natural language sentences or visual projections [1, 6]. These methods, however, do not answer *when* and *what* to signal and thus can significantly increase the human cognitive load when not done properly. In this paper, we propose an approach that enables a robot to convey its own intentions by signaling only when applicable and with what is necessary.

Consider an indoor domain, a human asks a robot to make coffee for him. His expectation is that the robot will navigate to the kitchen to make coffee and then get back and hand it over. The first action of the robot, however, is to navigate to a locker in the living room since the coffee maker was placed there after it was cleaned last time. The human may be confused at first if the robot doesn't explain its actions. However, note that if the robot returns later with the coffee maker from the living room, the human will understand the robot's actions even though they were confusing in the first place. If the robot in this scenario can explain that "*I am gonna move to the living room to fetch the coffee maker*" before it starts the execution, its actions will be better understood by the human. Such explicit signaling of the robot's intentions can largely improve the efficiency of human-robot interaction.

Inspired by this scenario, we formulate human interpreting robot's actions as a labeling process. Since the human interpretation of a robot action may be influenced not only by the previous actions but also future actions, different from our prior work [7], we use skip-chain conditional random fields (CRFs) [5] to learn this process. The two key questions, *when* and *what* to signal during the plan execution, can then be formulated as an inference problem with skip-chain CRFs while maximizing the human understanding of the robot plan. We evaluate our approach using a synthetic domain on Amazon Mechanical Turk (MTurk).

2 OUR APPROACH

In a human-robot teaming task, given the initial state $\mathcal{I}$ and goal state $\mathcal{G}$, the robot will generate a plan $\pi_{\mathcal{M}_R}$ using its own model $\mathcal{M}_R$. It is assumed that humans interpret each robot action in its plan by associating it with a task label using $\mathcal{M}_R^h$, which can be considered as the human's understanding of $\pi_{\mathcal{M}_R}$. In other words, we can only understand a robot action if we can associate it with a label. When we cannot assign any label to an action a, the label set of a will be $\emptyset$ and we say that a is inexplicable in the plan. An example of labeling a plan with 6 actions is shown in Figure 1. In order to make a plan explicable, plan signaling enables the robot to signal its intentions by providing context information *from the future* to maximize the probability that these inexplicable actions can now be labeled.

Learning: The label for each action is treated as a hidden variable. For each training sample, we assume that the plan is executed by a robot and a human subject observes the robot and assigns each action to a label. The sequential labeling process can easily be modeled by a CRFs. To capture long-term dependencies, we use skip-chain CRFs.

HRI'18, March 2018, Chicago, IL USA

ACM ISBN 123-4567-24-567/08/06.
https://doi.org/10.475/123_4

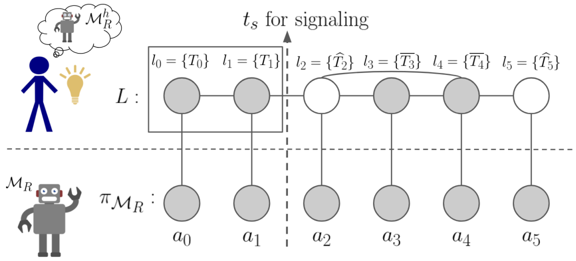

Figure 1: Robot signaling its intentions at t_s. l_2 was the inexplicable action. Robot chooses a timing t_s ($t_s = 1$ in this example) for signaling its intentions (i.e., "*I perform a_2 in order to achieve T_3 and T_4*"). All the labels before t_s which are in the bounding box will remain the same for computing Eq. (1). $\overline{L} = \{l_3, l_4\}$ is the set of task variables the robot chooses to provide more context information for signaling. The values of $\overline{L}$ is explored over all possible task labels in order to make labels for the remaining task variables (i.e., $\widehat{L} = \{l_2, l_5\}$) nonempty, resulting in $\pi_{\mathcal{M}_R}$ being fully explicable.

Inference: When robot tries to find the optimal timing and information to signal, it searches over all time slices before the inexplicable action, denoted by a_i (the i*th* action in the plan), and all the possible task label assignments (i.e., task labels after i) that can make the inexplicable action to be explicable. An example is shown in Figure 1. At each iteration, the robot picks $t_s \in \{0, \cdots, i-1\}$ to be the timing it performs signaling and $\overline{L}$ to be the set of task variables it will assign labels to (i.e., these labels determine the content of signaling). The optimal time and content can be determined by the choices that maximize the probability that the plan becomes fully explicable. We formulate the inference problem as a function s:

$$\underset{\{\overline{L}, T_{\overline{L}}, t_s\}}{\text{argmax}} \; s(\overline{L}, T_{\overline{L}}, t_s, T, \pi) \tag{1}$$

$s(\overline{L}, T_{\overline{L}}, t_s, T, \pi)$ can be computed as follows:

$$s(\overline{L}, T_{\overline{L}}, t_s, T, \pi) = P(\widehat{L} \neq \{\emptyset\} | L_{1:t_s} = T_{1:t_s}, \overline{L} = T_{\overline{L}}, \pi) \tag{2}$$

$$= \frac{\sum_{\widehat{L} \neq \{\emptyset\}} P(\widehat{L}, L_{1:t_s} = T_{1:t_s}, \overline{L} = T_{\overline{L}} | \pi)}{\sum_{\widehat{L'}} P(\widehat{L}', L_{1:t_s} = T_{1:t_s}, \overline{L} = T_{\overline{L}} | \pi)} \tag{3}$$

where $L_{1:t_s}$ is the set of task variables from start to t_s, $T_{1:t_s}$ is the set assigned to $L_{1:t_s}$ from T, and $\widehat{L}$ is the set of remaining variables that must be non-empty. It measures the entropy over the distribution of $\widehat{L} \neq \{\emptyset\}$ given π by modeling the human's interpretation of π.

After the robot getting what information should be signaled, it will generate a natural language sentence to convey its intentions. We create a set of templates that can translate the task labels and actions into natural language sentences, such as, "*I will [ACTION] soon*", "*I will [ACTION] soon in order to achieve [TASK]*".

3 EVALUATION

We evaluation our approach on Amazon Mechanical Turk (MTurk). We construct a synthetic smart house environment where there is a living room, a kitchen and three bedrooms, shown in Figure 2. There is also a vacuum cleaner, a coffee maker, coffee beans, and three keys to each bedroom in the environment (locations not shown). The robot will be randomly assigned tasks such as making coffee

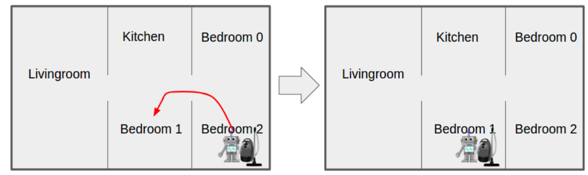

Figure 2: Synthetic domain for evaluation. (e.g., the robot fetches the vacuum in bedroom 2 and navigate to bedroom 1).

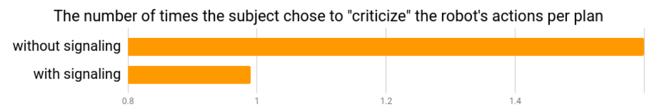

Figure 3: Results for robot signaling evaluation.

and cleaning an assigned room. The robot can generate its own plan. Given this synthetic domain, workers on MTurk are presented the robot's actions one by one. We provide two choices for them based on their evaluation of the actions. Whenever the workers think the robot's behavior is not explainable, they may choose "QUESTIONABLE". Otherwise, they can select "CONTINUE". We post two evaluation settings with the same set of scenarios for with and without signaling. We evaluate the performance of our approach by comparing the average number of "QUESTIONABLE" actions per plan. The result is shown in Figure 3. For the plans without signaling, there are on average 1.60 "QUESTIONABLE" actions the number is 0.99 with signaling.

4 CONCLUSIONS AND FUTURE WORK

In order to make the robot's plan explicable to its human teammate, we proposed an approach that enables the robot to explain its behavior by signaling its intentions. In this research, we formulate human interpreting robot's actions as a labeling process, similar to [7]. Given the plan, the robot first determines whether it is explicable to the human by checking if every action in the plan can be assigned at least one task label. If not, it will search for an optimal timing and content to signal its intentions in order to make the human better understand its plan. A skip-chain CRFs model is used for capturing long-term dependencies. To further evaluate our approach, we plan to also compare the performance of this signaling method with previous work [1, 2, 7].

REFERENCES

[1] R. S. Andersen, O. Madsen, T. B. Moeslund, and H. B. Amor. 2016. Projecting robot intentions into human environments. In *RO-MAN*. 294–301.
[2] T. Chakraborti, S. Sreedharan, Y. Zhang, and S. Kambhampati. 2017. Plan explanations as model reconciliation: Moving beyond explanation as soliloquy. In *IJCAI*. 156–163.
[3] A. Dragan and S. Srinivasa. 2013. Generating Legible Motion. In *Robotics: Science and Systems*.
[4] Z. Gong and Y. Zhang. 2018. Temporal Spatial Inverse Semantics for Robots Communicating with Humans. In *ICRA*.
[5] C. Sutton, A. McCallum, et al. 2012. An introduction to conditional random fields. *Foundations and Trends® in Machine Learning* 4, 4 (2012), 267–373.
[6] S. Tellex, R. Knepper, A. Li, D. Rus, and N. Roy. 2014. Asking for Help Using Inverse Semantics. In *Robotics: Science and Systems*.
[7] Y. Zhang, S. Sreedharan, A. Kulkarni, T. Chakraborti, H. H. Zhuo, and S. Kambhampati. 2017. Plan explicability and predictability for robot task planning. In *ICRA*. 1313–1320.