Toward Natural Explanations for a Robot's Navigation Plans

Raj Korpan
The Graduate Center, CUNY
rkorpan@gradcenter.cuny.edu

Susan L. Epstein
The Graduate Center and Hunter College, CUNY
susan.epstein@hunter.cuny.edu

ABSTRACT

People more readily accept and trust a robot that explains its behavior in natural language. This paper introduces Why-Plan, a method that compares the perspectives of an autonomous robot and a person when they plan a path for navigation. An implementation of Why-Plan demonstrates its ability to produce meaningful, human-friendly explanations quickly in natural language.

KEYWORDS

Human-robot interaction, robot navigation, explanations

1 INTRODUCTION

To build trust and understanding in its human collaborator, a robot should produce *natural explanations*, human-friendly descriptions of its behavior in natural language [6]. Such transparent, intelligible communication enables the robot to gain social acceptance and reduce confusion about its abilities. Our thesis is that a cognitive basis for a *robot controller* (autonomous decision system) facilitates the production of natural explanations. In particular, a controller that uses human-like rationales to make decisions can readily produce natural explanations. This paper focuses on navigation, an increasingly common task for a robot among people. The principal contribution reported here is Why-Plan, a method that addresses the question "Why does your plan go this way?"

Why-Plan exploits differences between planning objectives to produce clear, concise, natural explanations quickly. Consider the running example in Figure 1, where darker shading indicates greater anticipated traffic. The robot (denoted by a black box) faces in the direction of the arrow and is about to travel to its goal (the star). Despite a shorter available route (the dotted line) to its goal, the robot has planned a longer one (the solid line) to avoid the more crowded route that could hinder its travel and cause discomfort to people. Why-Plan's response contrasts two objectives: "take the shortest path" and "avoid crowds."

An autonomous robot can either *deliberate* (plan ahead to reach its goals) or *react* (choose one action at a time). Deliberation capitalizes on the robot's knowledge, while reactivity senses and responds to the environment (*world*). In a dynamic world, the robot must correct or abandon its plan when people move or unanticipated obstacles appear. The modern approach is a *hybrid* robot controller that integrates the flexibility and robustness of reactivity with the foresight of deliberation. As a result, the robot can plan a route to its goal, travel along that route, and manage the unexpected.

HRI'18, March 2018, Chicago, Illinois USA

Automated planners rely on representations that are unfamiliar and opaque to most people. Previous approaches that explained automatically generated plans have represented and described them in first-order logic [7], interpreted classical planning formulations to produce explanations [4], or generated classical plans that sacrificed optimality to be more human-like [8]. None of these, however, implements explanations for navigation plans in natural language. The work described here addresses that gap within the context of the robot's most recent action and its long-range perspective, but does not provide a description of the complete plan.

2 PLANNING AND REASONING

A *plan* here is a finite ordered sequence of *waypoints*, locations through which the robot intends to pass during travel on its way to a goal location G. A^* is a traditional search algorithm to find a plan that performs well with respect to some objective O, such as "minimize distance traveled" or "avoid crowds." The two routes in our running example arose separately from those two objectives.

A* relies on a graph whose nodes represent possible locations in a static world. An edge between two nodes in the graph indicates that the robot can move directly from one location to the other. Each edge has a weight that represents the anticipated cost of that action from O's perspective. Because there are many possible objectives, A* can be seen as a family of heuristic planners that may produce different plans.

SemaFORR is a cognitively-based hybrid robot controller that learns a human-like spatial model [3]. Given a goal and a set of possible actions, SemaFORR makes a decision with two sets of reactive *Advisors*, procedures that rely on good reasons to take an action. First, the rule-based Advisors mandate clearly correct actions or veto unacceptable ones (e.g., "go directly to the goal in front of you" or "do not move into that wall"). Then, if no action has been mandated and multiple choices remain, Advisors that are based on commonsense (e.g., "get closer to the goal") and on SemaFORR's spatial model (e.g., "enter that area through that door") vote to select an action. Our earlier work, Why, produced natural explanations for SemaFORR's reactive decisions, based on this reasoning framework and the Advisors' rationales [5].

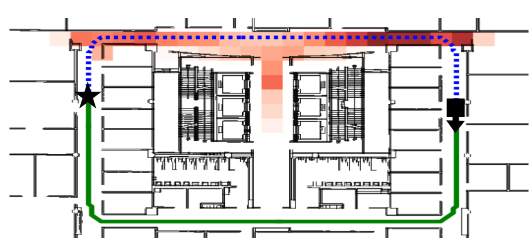

Figure 1: Why-Plan explains, "Although there may be a somewhat shorter way, I think mine is a lot less crowded."

SemaFORR is also deliberative, however, with multiple planners available to it. One of them is *CSA**, which applies A* to a graph whose edge weights incorporate both distance and the crowd density the robot observes in discretized spaces as it travels through the world [1]. Given a plan P, SemaFORR's Advisors determine how to travel from one waypoint to the next. WHY-PLAN uses P to explain, in natural language, the robot's long-range perspective.

3 EXPLANATIONS BASED ON OBJECTIVES

The premise of WHY-PLAN is that a human plans from the perspective of objective $\mathcal{H}$, while the robot plans from the perspective of objective $\mathcal{R}$. As a running example, we let $\mathcal{H}$ be "take the shortest path" and $\mathcal{R}$ be *crowd sensitivity*, the objective of CSA*. A question about the robot's plan could arise anywhere along its intended path. At time t, the robot is in *state* s_t described by its location, the direction it faces, and the obstructions its sensors currently detect. WHY-PLAN takes as input the goal's location G, the robot's current state s_t, its planning objective $\mathcal{R}$, and an alternative planning objective $\mathcal{H}$ attributed to the human questioner. We assume that the robot and the person share knowledge relevant to $\mathcal{R}$ and $\mathcal{H}$ (e.g., distances and crowding). This allows the robot to construct two A* plans: its own plan P_R based on $\mathcal{R}$, and P_H, the robot's approximation of the human's implicit plan, based on $\mathcal{H}$.

WHY-PLAN compares the robot's plan and the human's plan from both perspectives. The core of its explanation is how the planner's objectives differ. As in Table 1, given the human's plan P_H and the robot's plan P_R, $H(P_H)$ and $H(P_R)$ measure how well those two plans would perform under the human's objective (e.g., their length if $\mathcal{H}$ is shortest path). Given our shared knowledge assumption and planners that seek to minimize their respective objectives, under the human's objective the human plan should perform at least as well as the robot's ($D_H = H(P_R) - H(P_H) \geq 0$). Similarly, $R(P_H)$ and $R(P_R)$ measure how well the two plans would perform under the robot's objective (e.g., crowd sensitivity). Under the robot's objective, its own plan should perform at least as well as the human's ($D_R = R(P_R) - R(P_H) \leq 0$).

In reply to "Why does your plan go this way?" WHY-PLAN uses the differences D_H and D_R to produce an explanation in natural language. Each objective has an adjective (h for $\mathcal{H}$ and r for $\mathcal{R}$, e.g., h = "short" and r = "crowded") and a set of user-specified numeric intervals used to transform the numeric differences into natural language. For example, in Table 1, two plans for which $D_R = -799$ would map the robot's plan to "somewhat less crowded."

Although $D_H \geq 0$ and $D_R \leq 0$ produce four possible cases, each with its own language template, we restrict discussion here to the three that occur under CSA*. If both D_H and D_R are 0, then the plans equally address the two objectives, and WHY-PLAN explains "I decided to go this way because I think it is just as h and equally r." If only D_H = 0, WHY-PLAN explains the robot's objective: "I think mine is [D_R phrase] r'." Otherwise, D_R is negative, D_H is positive, and the robot prefers its plan because "Although there may be a [D_H phrase] h' way, I think mine is [D_R phrase] r'" where h' and r' compare h and r, respectively (e.g., "shorter" and "less crowded").

SemaFORR and WHY-PLAN are implemented as separate packages in *ROS*, the state-of-the-art Robot Operating System. The work reported here used a ROS-integrated simulator [2] for a real-world robot (Fetch Robotics' Freight) as it navigated to 40 destinations in a complex, 60m×90m office world. WHY-PLAN produced explanations in real time (μ = 4.9 ms, σ = 2.5 ms, n = 3327) while the robot traveled. WHY-PLAN's speed allows it to compute an explanation for each state, which it provides only when asked. In Figure 1, P_R (solid line) plans to avoid the crowd, while P_H (dotted line) takes a shorter path through it. Because $D_R = -6729$ and $D_H = 9.6$, the robot explains, "Although there may be a somewhat shorter way, I think mine is a lot less crowded." Further inspection of the simulated results indicates that WHY-PLAN's other natural explanations for SemaFORR's plans are similarly intelligible and transparent.

WHY-PLAN presumes that questions arise from a difference between the human's and the robot's objectives, but they could also stem from a violation of our shared knowledge assumption. A broader system for human-robot collaboration would seek the cause of such a mismatch, use plan explanations to resolve it, and then adjust the robot's responses based on feedback from its human partner. WHY-PLAN could also use known planning objectives to detect a more complex human objective. For example, given a plan P from a person or an unknown heuristic planner, WHY-PLAN could use the individual objectives in its repertoire to tease apart and then characterize how P weighted its objectives (e.g., "You appear to consider distance twice as important as travel time.").

Current work on WHY-PLAN addresses a variety of issues, including integration of WHY-PLAN with WHY, more thorough evaluation in different worlds, and violations of assumptions made during planning (e.g., now-blocked passageways or empty areas that are typically crowded). WHY-PLAN will also be expanded to answer other important questions, such as "Why doesn't your plan go that way?", "What makes your plan better than mine?", "What's another

	Performance					Phrasing		
	Planning objective	Plan	Human evaluation	Robot evaluation	Difference	Numeric intervals	Comparators	Descriptor
Human $\mathcal{H}$	short path	P_H	$H(P_H)$	$R(P_H)$	$D_H = H(P_R) - H(P_H)$	$(0, 1]$ $(1, 10]$ $(10, +\infty)$	a bit somewhat a lot	h = short h' = shorter
Robot $\mathcal{R}$	uncrowded path	P_R	$H(P_R)$	$R(P_R)$	$D_R = R(P_R) - R(P_H)$	$(-\infty, -1500]$ $(-1500, -500]$ $(-500, 0)$	a lot somewhat a bit	r = crowded r' = less crowded

Table 1: WHY-PLAN maps from planning objectives to language

way we could go?", and "How sure are you about your plan?" To measure improvements in understanding and trust, we also intend to evaluate WHY-PLAN's impact on human subjects. Meanwhile, WHY-PLAN produces nuanced but direct natural explanations for the plans of a robot navigator.

REFERENCES

[1] Anoop Aroor and Susan L Epstein. 2017. Toward Crowd-Sensitive Path Planning. In *AAAI Fall Symposium on Human-Agent Groups: Studies, Algorithms and Challenges*. 238–245.

[2] Anoop Aroor, Susan L Epstein, and Raj Korpan. 2017. MengeROS: A Crowd Simulation Tool for Autonomous Robot Navigation. In *AAAI Fall Symposium on Artificial Intelligence for Human-Robot Interaction*. 123–125.

[3] Susan L Epstein, Anoop Aroor, Matthew Evanusa, Elizabeth I Sklar, and Simon Parsons. 2015. Learning Spatial Models for Navigation. In *12th International Conference on Spatial Information Theory*. Springer, 403–425.

[4] Maria Fox, Derek Long, and Daniele Magazzeni. 2017. Explainable Planning. In *IJCAI-17 Workshop on Explainable AI (XAI)*. 24–30.

[5] Raj Korpan, Susan L Epstein, Anoop Aroor, and Gil Dekel. 2017. WHY: Natural Explanations from a Robot Navigator. In *AAAI Fall Symposium on Natural Communication for Human-Robot Collaboration*.

[6] Todd Kulesza, Simone Stumpf, Margaret Burnett, Sherry Yang, Irwin Kwan, and Weng-Keen Wong. 2013. Too Much, Too Little, or Just Right? Ways Explanations Impact End Users' Mental Models. In *2013 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 3–10.

[7] Bastian Seegebarth, Felix Müller, Bernd Schattenberg, and Susanne Biundo. 2012. Making Hybrid Plans More Clear to Human Users- A Formal Approach for Generating Sound Explanations. In *Twenty-Second International Conference on Automated Planning and Scheduling*. 225–233.

[8] Sarath Sreedharan, Tathagata Chakraborti, and Subbarao Kambhampati. 2017. Balancing Explicability and Explanation in Human-Aware Planning. In *AAAI Fall Symposium on Artificial Intelligence for Human-Robot Interaction*. 61–68.