WHY: Natural Explanations from a Robot Navigator

Raj Korpan[1], Susan L. Epstein[1,2], Anoop Aroor[1], and Gil Dekel[2]
[1]The Graduate Center and [2]Hunter College, City University of New York
rkorpan@gradcenter.cuny.edu, susan.epstein@hunter.cuny.edu, aaroor@gradcenter.cuny.edu, gil.dekel47@myhunter.cuny.edu

Abstract

Effective collaboration between a robot and a person requires natural communication. When a robot travels with a human companion, the robot should be able to explain its navigation behavior in natural language. This paper explains how a cognitively-based, autonomous robot navigation system produces informative, intuitive explanations for its decisions. Language generation here is based upon the robot's commonsense, its qualitative reasoning, and its learned spatial model. This approach produces natural explanations in real time for a robot as it navigates in a large, complex indoor environment.

Introduction

Successful human-robot collaboration requires *natural explanations*, human-friendly descriptions of the robot's reasoning in natural language. In *collaborative navigation*, a person and an autonomous robot travel together to some destination. The thesis of this paper is that natural explanations for collaborative navigation emerge when a *robot controller* (autonomous navigation system) is cognitively based. This paper introduces WHY, an approach that accesses and conveys the robot's reasoning to provide its human companion with insight into its behavior. The principal results presented here are natural explanations from an indoor robot navigator.

Even in unfamiliar, complex spatial environments (*worlds*), people travel without a map to reach their goals successfully (Conlin 2009). Efficient human navigators reason over a mental model that incorporates commonsense, spatial knowledge, and multiple heuristics (Golledge 1999). They then use the same model to explain their chosen path and their reasons for decisions along the way. Our research goal is an autonomous robot navigator that communicates with its human companions much the way people do.

WHY explains a navigation decision in natural language. It anticipates three likely questions from a human companion: "Why did you decide to do that?" "Why not do something else?" and "How sure are you that this is the right decision?" WHY generates its answers with *SemaFORR*, a robot controller that learns a spatial model from sensor data as it travels through a partially-observable world without a map (Epstein et al. 2015). SemaFORR's cognitively-based reasoning and spatial model facilitate natural explanations.

WHY is an interpreter; it uses SemaFORR's cognitive foundation to bridge the perceptual and representational gap between human and robot navigators. WHY and SemaFORR could accompany any robot controller to provide natural explanations. More broadly, WHY can be readily adapted to explain decisions for other applications of FORR, SemaFORR's underlying cognitive architecture.

The next section of this paper reviews related work. Subsequent sections describe SemaFORR and formalize WHY. Finally, we evaluate WHY-generated explanations and give examples of them as our mobile robot navigates through a large, complex, indoor world.

Related Work

When a robot represents and reasons about space similarly to the way people do, it facilitates human-robot collaboration (Kennedy et al. 2007). Communication with a robot allows people to build a mental model of how it perceives and reasons, and thereby helps to establish trust (Kulesza et al. 2013; Bussone, Stumpf, and O'Sullivan 2015). A recent approach grounded perceived objects between the robot and a person to build a mutual mental model, and then generated natural language descriptions from it (Chai et al. 2016). Although that supported natural dialogue, it did not explain the reasoning that produced the robot's behavior.

Despite much work on how a robot might understand natural language from a human navigator (Boularias et al. 2016; Duvallet et al. 2016; Thomason et al. 2015), natural explanations from a robot navigator to a person remain an important open problem. Such work has thus far required detailed logs of the robot's experience, which only trained researchers could understand (Landsiedel et al. 2017; Scalise, Rosenthal, and Srinivasa 2017). It is unreasonable, however, to expect people to decipher robot logs.

Natural language descriptions of a robot's travelled path have addressed abstraction, specificity, and locality (Rosenthal, Selvaraj, and Veloso 2016; Perera et al. 2016). A similar approach generated path descriptions to improve sentence correctness, completeness, and conciseness (Barrett et al. 2017). Those approaches, however, used a labeled map to generate descriptions and did not explain the robot's reasoning. Other work visually interpreted natural-language navigation commands with a semantic map that showed the robot's resulting action (Oh et al. 2016). Although a person might eventually unpack the robot's reasoning process this way, no natural language explanation was provided.

Researchers have generated navigation instructions in natural language from metric, topological, and semantic information about the world (Daniele, Bansal, and Walter 2016) or rules extracted from human-generated instructions (Dale, Geldof, and Prost 2005). Other work has focused on human spatial cognition (Look 2008), or on simplicity and understandability (Richter and Duckham 2008). None of these approaches, however, can explain how the instructions were generated, nor can they justify a particular instruction.

More generally, researchers have sought human-friendly explanations for systems that learn. Trust in and understanding of a learning system improved when people received an explanation of why a system behaved one way and not another (Lim, Dey, and Avrahami 2009). Several approaches to sequential tasks have explained Markov decision processes, but the resultant language was not human-friendly and was not based on human reasoning (Ramakrishnan and Shah 2016; Dodson et al. 2013; Khan et al. 2011). In summary, although intelligent systems should be able to provide natural explanations during collaborative navigation, to the best of our knowledge no work has focused on explanations for the robot's decisions. WHY addresses that gap.

SemaFORR

SemaFORR is a robot controller implemented in *ROS*, the state-of-the-art Robot Operating System. SemaFORR selects one action at a time to move the robot to its target location. Instead of a world map, SemaFORR uses local sensor data, learned knowledge, and reactive, heuristic reasoning to contend with any obstacles and reach its target. The resultant behavior is satisficing and human-like rather than optimal.

A *decision state* records the robot's current sensor data and its *pose* $\langle x, y, \theta \rangle$, where $\langle x, y \rangle$ is its location and θ is its orientation with respect to an allocentric, two-dimensional coordinate system. As the robot travels, its *path* to a target is recorded as a finite sequence of decision states. SemaFORR makes decisions based on a hierarchical reasoning framework and a spatial model that it learns while it navigates. WHY uses them both to generate its explanations.

Spatial Model

SemaFORR learns its compact, approximate spatial model from experience. The model captures many of the features of a cognitive map, the representation that people construct as they navigate (Foo et al. 2005). Instead of a metric map,

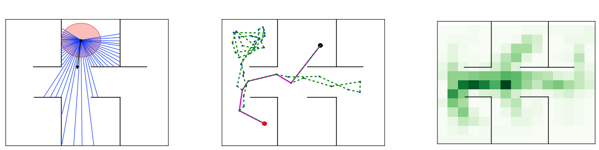

(a) (b) (c)

Figure 1: Examples of affordances in a simple world (a) a region (b) a path and its trail (c) conveyors

SemaFORR's model is a set of *spatial affordances*, abstract representations that preserve salient details and facilitate movement. As the robot travels or once it reaches its target, it learns spatial affordances from local sensor readings and stores them as episodic memory. Figure 1 gives examples.

A *region* is an unobstructed area where the robot can move freely, represented as a circle. A region's center is the robot's location in a decision state; its radius is the minimum distance sensed from the center to any obstacle. An *exit* is a point that affords access to and from a region, learned as a point where the robot crossed the region's circumference .

A *trail* refines a path the robot has taken. It is an ordered list of *trail markers*, decision states selected from the robot's path. The first and last trail markers are the initial and final decision states on the path. Trail learning works backward from the end of the path; it creates a new trail marker for the earliest decision state that could have sensed the current trail marker. The resultant trail is usually shorter than the original path and provides a more direct route to the target.

A *conveyor* is a small area that facilitates travel. It is represented in a grid superimposed on the world, where each cell tallies the frequency with which trails pass through it. High-count cells in the grid are conveyors.

The spatial model combines affordances to produce more powerful representations. For example, a *door* generalizes over the exits of a region. It is represented as an arc along the region's circumference. The door-learning algorithm introduces a door when the length of the arc between two exits is within some small ϵ. Once generated, a door incorporates additional exits if they are within ϵ of it. Another example is the *skeleton*, a graph that captures global connectivity with a node for each region. An edge in the skeleton joins two nodes if a path has ever moved between their corresponding regions. Along with commonsense qualitative reasoning, affordances are used to select the robot's next action.

Reasoning Framework

SemaFORR is an application of *FORR*, a cognitive architecture for learning and problem solving (Epstein 1994). FORR is both reactive and deliberative. Reactivity supports flexibility and robustness, and is similar to how people experience and move through space (Spiers and Maguire 2008). Deliberation makes plans that capitalize on the robot's experience; it is the focus of current work (Aroor and Epstein, in press).

The crux of any FORR-based system is that good decisions in complex domains are best made reactively, by a mixture of good reasons. FORR represents each good reason by a procedure called an *Advisor*. Given a decision state and a discrete set of possible actions, an Advisor expresses its opinions on possible actions as *comments*. In a *decision cycle*, SemaFORR uses those comments to select an action. Possible actions are alternately a set of forward moves of various lengths or a set of turns in place of various rotations. A move with distance 0 is equivalent to a pause. Thus, in any given decision state, SemaFORR chooses only the intensity level of its next move or turn. The resultant action sequence is expected to move the robot to its target.

SemaFORR's Advisors are organized into a three-tier hierarchy, with rules in tier 1 and commonsense, qualitative heuristics in tier 3. Tier 1 invokes its Advisors in a predetermined order; each of them can either mandate or veto an action. If no action is mandated, the remaining, unvetoed actions are forwarded to tier 3. (Natural explanations for tier 2, SemaFORR's deliberative layer, are a focus of current work.) Table 1 lists the Advisors' rationales by tier.

Each tier-3 Advisor constructs its comments on the remaining possible actions with its own commonsense rationale. Comments assign a *strength* in [0,10] to each available action. Strengths near 10 indicate actions that are in close agreement with the Advisor's rationale; strengths near 0 indicate direct opposition to it. For n Advisors, m actions, and comment strength c_{ij} of Advisor i on action j, SemaFORR selects the action with the highest total comment strength:

$$argmax_{j \in m} \sum_{i=1}^{n} c_{ij}.$$

Because ties are broken at random, tier 3 introduces uncertainty into action selection. For further details on SemaFORR, see (Epstein et al. 2015).

Approach

This section describes how WHY exploits SemaFORR to generate natural explanations. Each of the three questions below focuses on a different aspect of a robot controller. The result is a rich, varied set of natural explanations.

Why did you do that?

The first question asks why the robot chose a particular action. WHY constructs its answer from the rationales and comments of the Advisors responsible for that choice, with templates to translate actions, comments, and decisions into natural language. Given the robot's current pose, WHY maps each possible action onto a descriptive phrase for use in any [action] field. Examples include "wait" for a forward move of 0.0 m, "inch forward" for a forward move of 0.2 m, and "shift right a bit" for a turn in place of 0.25 rad.

Tier 1, in order	
VICTORY	Go toward an unobstructed target
AVOIDWALLS	Do not go within ϵ of an obstacle
NOTOPPOSITE	Do not return to the last orientation
Tier 3	
Based on commonsense reasoning	
BIGSTEP	Take a long step
ELBOWROOM	Get far away from obstacles
EXPLORER	Go to unfamiliar locations
GOAROUND	Turn away from nearby obstacles
GREEDY	Get close to the target
Based on the spatial model	
ACCESS	Go to a region with many doors
CONVEY	Go to frequent, distant conveyors
ENTER	Go into the target's region
EXIT	Leave a region without the target
TRAILER	Use a trail segment to approach the target
UNLIKELY	Avoid dead-end regions

Table 1: SemaFORR's Advisors and their rationales. Tier 2 is outside the scope of this paper.

Algorithm 1 is pseudocode for WHY's responses. WHY takes as input the current decision state, target location, and spatial model, and then calculates its response based on the comments from SemaFORR's Advisors. There are three possibilities: tier 1 chose the action, tier 1 left only one unvetoed action, or tier 3 chose the action. SemaFORR only makes a decision in tier 1 if VICTORY mandates it or AVOIDWALLS has vetoed all actions but the pause. The applicable templates in those cases are "I could see our target and [action] would get us closer to it" and "I decided to wait because there's not enough room to move forward."

The inherent uncertainty and complexity of a tier-3 decision, however, requires a more nuanced explanation. For a set of m actions, assume tier-3 Advisor D_i outputs comment with strengths $c_{i1}, \ldots, c_{im} \in [0, 10]$. D_i's *t*-support for action a_k is the *t*-statistic $t_{ik} = (c_{ik} - \bar{c_i})/\sigma_i$ where $\bar{c_i}$ is the mean strength of D_i's comments in the current decision state and σ_i is their standard deviation. (This is not a z-score because sampled values replace the unavailable true population mean and standard deviation.) WHY can compare different Advisors' *t*-supports because they have common mean 0 and standard deviation 1. If $|t_{ik}|$ is large, Advisor D_i has a strong opinion about action a_k relative to the other actions: supportive for $t_{ik} > 0$ and opposed for $t_{ik} < 0$.

Table 2 provides a running example. It shows the original comment strengths from four Advisors on four actions, and the total strength C_k for each action a_k. Tier 3 chooses action a_4 because it has maximum support. While D_1 and D_2 support a_4 with equal strength, the *t*-support values tell a different story: D_1 prefers a_4 much more ($t_{14} = 1.49$) than D_2 does ($t_{24} = 0.71$). Moreover, D_3 and D_4 actually oppose a_4 (-0.34 and -0.78, respectively).

For each measure, we partitioned the real numbers into three intervals and assigned a descriptive natural language phrase to each one, as shown in Table 3. This partitioning allows WHY to hedge in its responses, much the way people explain their reasoning when they are uncertain (Markkanen and Schröder 1997). WHY maps the *t*-support values into

Algorithm 1: WHY's Explanation Procedure

Input: *decision state, target location, spatial model*
Output: *explanation*
switch *mode(decision)* **do**
  case *tier 1 decides action*
    explanation ← sentence based on VICTORY
  case *only 1 unvetoed action remains after tier 1*
    explanation ← sentence based on vetoes
  otherwise
    Compute *t*-statistics for tier-3 Advisors' strengths
    Categorize the support level for the chosen action
    Complete template for each Advisor with its support level and rationale
    explanation ← combined completed templates
  endsw
endsw
return *explanation*

these intervals. For a_4, D_1's *t*-support of 1.49 is translated as "want" and D_4's -0.78 is translated as "don't want". WHY then completes the clause template "I [phrase] to [rationale]" for each Advisor based on Table 1 and less model-specific language from Table 3. For example, if D_1 were GREEDY, then the completed clause template for a_4 would be "I want to get close to the target."

Finally, WHY combines completed clause templates into the final tier-3 explanation, but omits language from Advisors with *t*-support values in (-0.75, 0.75] because they contribute relatively little to the decision. WHY concatenates the remaining language with appropriate punctuation and conjunctions to produce its tier-3 explanation: "(Although [language from opposed Advisors],) I decided to [action] because [language from supporting Advisors]". The portion in parentheses is omitted if no opposition qualifies. If the Advisors in the running example were GREEDY, ELBOWROOM, CONVEY, and EXPLORER, in that order, and a_4 were move forward 1.6 m, then the natural explanation is "Although I don't want to go somewhere I've been, I decided to move forward a lot because I want to get close to our target." (Note that D_2's support fails the -0.75 filter and so is excluded.)

This approach can also respond to "What action would you take if you were in another context?" Given the decision state and the target location, WHY would reuse its current spatial model, generate hypothetical comments, and process them in the same way. The sentence template would substitute "I would [action]" for "I decided to [action]."

How sure are you that this is the right decision?

The second question from a human collaborator is about the robot's confidence in its decision, that is, how much it trusts that its decision will help reach the target. Again, WHY responds based on the tier that selected the action. Tier 1's rule-based choices are by definition highly confident. If VICTORY chose the action then the response is "Highly confident, since our target is in sensor range and this would get us closer to it." If AVOIDWALLS vetoed all forward moves except the pause, then the explanation is "Highly confident, since there is not enough room to move forward."

Again, tier-3's uncertainty and complexity require more nuanced language, this time with two measures: level of agreement and overall support. The extent to which the tier-3 Advisors agree indicates how strongly the robot would like to take the action. WHY measures the level of that agreement with Gini impurity, where values near 0 indicate a high level of agreement and values near 0.5 indicate disagreement (Hastie, Tibshirani, and Friedman 2009). For n tier-3 Advisors and maximum comment strength 10, the *level of agreement* $G_k \in$ [0,0.5] on action a_k is defined as

$$G_k = 2 \cdot \left[\frac{\sum_{i=1}^{n} c_{ik}}{10n} \right] \cdot \left[1 - \frac{\sum_{i=1}^{n} c_{ik}}{10n} \right].$$

In the example of Table 2, the level of agreement on a_4 is $G_4 = 2 \cdot \left[\frac{22}{40}\right] \cdot \left[1 - \frac{22}{40}\right] \approx 0.50$. This indicates considerable disagreement among the Advisors in Table 2.

The second confidence measure is SemaFORR's overall support for its chosen action compared to other possibilities, defined as a t-statistic across all tier-3 comments. Let μ_C be the mean total strength of all actions *C* under consideration by tier 3, and σ_C be their standard deviation. We define the *overall support* for action a_k as $T_k = (C_k - \mu_C)/\sigma_C$. T_k indicates how much more the Advisors as a group would like to perform a_k than the other actions. In Table 2, the overall support T_4 for a_4 is 0.66, which indicates only some support for a_4 over the other actions.

WHY weights level of agreement and overall support equally to gauge the robot's confidence in a tier-3 decision with *confidence level* $L_k = (0.5 - G_k) \cdot T_k$ for a_k. It then maps each of L_k, G_k, and T_k to one of three intervals and then to natural language, as in Table 3, with implicit labels *low* < *medium* < *high* in order for each statistic. Two statistics *agree* if they have the same label; one statistic is lower than the other if its label precedes the other's in the ordering.

All responses to this question use a template that begins "I'm [L_k adverb] sure because...." If G_k and T_k both agree with L_k, the template continues "[G_k phrase]. [T_k phrase]." For example, "I'm really sure about my decision because I've got many reasons for it. I really want to do this the most." If only one agrees with L_k, the template continues "[phrase for whichever of G_k or T_k agrees]." For example, "I'm not sure about my decision because my reasons conflict." Finally, if neither agrees with L_k, it concludes "even though [phrase for whichever of G_k or T_k is lower than L_k], [G_k phrase or T_k phrase that is higher than L_k]." For example, "I am only somewhat sure about my decision because,

	c_{ik}				t_{ik}			
	a_1	a_2	a_3	a_4	a_1	a_2	a_3	a_4
D_1	0	1	1	10	-0.64	-0.43	-0.43	1.49
D_2	0	8	9	10	-1.48	0.27	0.49	0.71
D_3	2	0	10	2	-0.34	-0.79	1.47	-0.34
D_4	3	10	1	0	-0.11	1.44	-0.55	-0.78
C_k	5	19	21	22				

Table 2: Example of comments from tier-3 Advisors D_i on actions a_k, where c_{ik} is strength and t_{ik} is *t*-support

t-support	$(-\infty, -1.5]$	really don't want
$t_{ik} \leq 0$:	$(-1.5, -0.75]$	don't want
opposed	$(-0.75, 0]$	somewhat don't want
t-support	$(0, 0.75]$	somewhat want
$t_{ik} > 0$:	$(0.75, 1.5]$	want
supportive	$(1.5, +\infty)$	really want
Level of agreement G_k	$(0.45, 0.5]$	My reasons conflict
	$(0.25, 0.45]$	I've only got a few reasons
	$[0, 0.25]$	I've got many reasons
Overall support T_k	$(-\infty, 0.75]$	don't really want
	$(0.75, 1.5]$	somewhat want
	$(1.5, +\infty)$	really want
Confidence level L_k	$(-\infty, 0.0375]$	not
	$(0.0375, 0.375]$	only somewhat
	$(0.375, +\infty)$	really
Difference in	$(0, 0.75]$	slightly more
overall support	$(0.75, 1.5]$	more
$T_k - T_j$	$(1.5, +\infty)$	much more

Table 3: Phrase mappings from value intervals to language

even though I've got many reasons, I don't really want to do this the most." For a_4 in Table 2, L_4 is near 0, G_4 = 0.50, and T_4 = 0.66. This produces the natural explanation "I'm not sure about my decision because my reasons conflict. I don't really want to do this more than anything else."

Why not do something else?

A human collaborator makes decisions with her own mental model of the world. When her decision conflicts with another team member's, she tries to understand why they made a different decision. WHY's approach explains SemaFORR's preference for action a_k over an alternative a_j. If tier 1 chose a_k, the explanation uses VICTORY's rationale: "I decided not to [action$_j$] because I sense our goal and another action would get us closer to it." If AVOIDWALLS or NOTOPPOSITE vetoed a_j, then the natural explanation is "I decided not to [action] because [rationale from Advisor that vetoed it]."

The other possibility is that a_j had lower total strength in tier 3 than a_k did. In this case, WHY generates a natural explanation with the tier-3 Advisors that, by their comment strengths, discriminated most between the two actions. WHY calculates $t_{ik} - t_{ij}$ for each Advisor D_i. If the result lies in [-1, 1] then D_i's support is similar for a_k and a_j; otherwise D_i displays a *clear preference*. The natural explanation includes only those Advisors with clear preferences.

The explanation template is "I thought about [action$_j$] (because it would let us [rationales from Advisors that prefer action$_j$]), but I felt [phrase] strongly about [action$_k$] since it lets us [rationales from Advisors that prefer action$_k$]." The [phrase] is the extent to which SemaFORR prefers a_k to a_j. It is selected based on $T_k - T_j$, the difference in the actions' overall support, and mapped into intervals as in Table 3. The portion in parentheses is only included if any Advisors showed a clear preference for action$_j$.

For "Why didn't you take action a_2?" on our running example, WHY calculates the difference in overall support between a_4 and a_2 at 0.38, which maps to "slightly more." The differences in t-support between a_4 and a_2 are 1.92, 0.44, 0.45, and -2.22. Thus, if D_1 is GREEDY and prefers a_4, while D_4 is EXPLORER and prefers a_2, the natural explanation is "I thought about a_2 because it would let us go somewhere new, but I felt slightly more strongly about a_4 since it lets us get closer to our target."

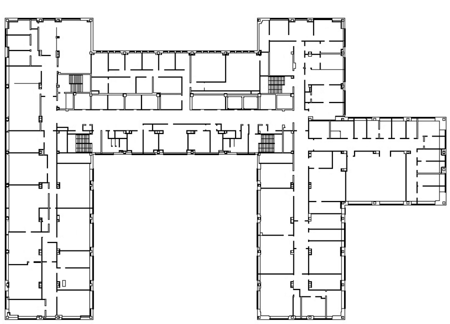

Figure 2: Tenth floor of Hunter College's North Building

Results

Implemented as a ROS package, WHY explains SemaFORR's decisions in real time. We evaluated WHY in simulation for a real-world robot (Fetch Robotics' Freight). When the robot navigated to 230 destinations in the complex 60m×90m office world of Figure 2, WHY averaged less than 3 msec per explanation.

WHY's many distinct natural explanations simulate people's ability to vary their explanations based on their context (Malle 1999). Table 4 provides further details. The Coleman-Liau index measures text readability; it gauged WHY's explanations over all three questions at approximately a 6th-grade level (Coleman and Liau 1975), which should make them readily understandable to a layperson.

For action a_k chosen in tier 3 and every possible alternative a_j, Table 5 shows how often the values of G_k, T_k, L_k, $t_{ik} - t_{ij}$, and $T_k - T_j$ fell in their respective Table 3 intervals. The Advisors disagreed ($G_k > 0.45$) on 67.15% of decisions. Strong overall support ($T_k > 1.5$) made SemaFORR strongly confident in 2.44% of its decisions ($L_k > 0.375$) and somewhat confident in 42.64% of them. When asked about an alternative, individual Advisors clearly preferred ($t_k - t_j > 1$) the original decision 39.50% of the time; SemaFORR itself declared a strong preference ($T_k - T_j > 1.5$) between the two actions 61.13% of the time.

Table 6 illustrates WHY's robust ability to provide nuanced explanations for tier-3 decisions. The target appears as an asterisk and the black box and arrow show the robot's pose. Decision 1 was made when the robot had not yet learned any spatial affordances; decision 2 was made later, when the spatial model was more mature. In decision 3, the Advisors strongly disagreed, while in decision 4 the spatial model-based Advisors disagreed with a commonsense-based Advisor.

Tier where made	1	3	All
Number of decisions	22,982	84,920	107,902
Avg. computation time (ms)	0.45	3.08	2.52
Unique phrasings			
Why?	14	31,896	31,910
Confidence?	2	11	13
Something else?	19	124,086	124,105
Total	35	155,993	156,028
Average readability			
Why?	8.18	5.02	5.70
Confidence?	10.39	7.63	8.22
Something else?	3.91	6.44	5.96
Overall	5.36	6.41	6.21

Table 4: Empirical explanations

	Low	Medium	High
G_k	67.15%	30.41%	2.44%
T_k	2.34%	60.09%	37.57%
L_k	54.92%	42.64%	2.44%
$t_{ik} - t_{ij}$	16.09%	44.41%	39.50%
$T_k - T_j$	18.48%	20.40%	61.13%

Table 5: Metric distributions by interval in tier-3 decisions

Decision State			
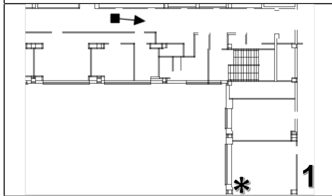	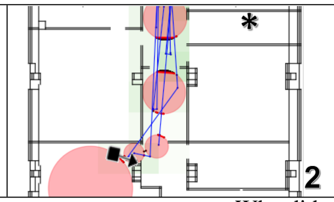	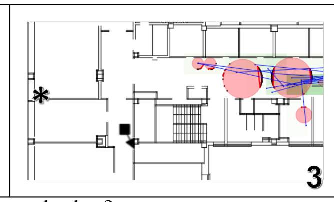	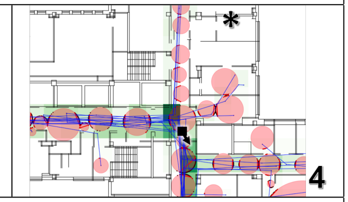
Why did you do that?			
Although I don't want to go close to that wall, I decided to bear right because I really want to take a big step.	Although I don't want to turn towards this wall, I decided to turn right because I want to go somewhere familiar, I want to get close to our target, and I want to follow a familiar route that gets me closer to our target.	Although I really don't want to go close to that wall and I really don't want to get farther from our target, I decided to move forward a lot because I really want to go to an area I've been to a lot, I really want to take a big step, and I really want to go somewhere new.	Although I don't want to get farther from our target, I decided to bear left because I really want to go somewhere familiar and I want to leave since our target isn't here.
How sure are you?			
I'm not sure because my reasons conflict.	I'm only somewhat sure because, even though my reasons conflict, I really want to do this most.	I'm not sure because my reasons conflict.	I'm only somewhat sure in my decision because I've only got a few reasons. I somewhat want to do this most.
Why not do something else?			
I thought about turning left because it would let us stay away from that wall and get close to our target, but I felt more strongly about bearing right since it lets us take a big step and get around this wall.	I thought about shifting left a bit because it would let us get around this wall, but I felt much more strongly about turning right since it lets us go somewhere familiar and get close to our target.	I decided not to move far forward because the wall was in the way.	I thought about turning hard right because it would let us get close to our target, but I felt much more strongly about bearing left since it lets us go somewhere familiar, leave since our target isn't here, go somewhere new, and get around this wall.

Table 6: Explanations for decision states and any current spatial model, enlarged from Figure 2

Discussion

WHY is applicable more broadly than we have indicated thus far. Any robot controller could have SemaFORR learn the spatial model in parallel, and use it with WHY to produce transparent, cognitively-plausible explanations. If the alternative controller were to select action a_j when SemaFORR selected a_k, WHY could still explain a_j with any Advisors that supported it, and offer an explanation for a_k as well. Furthermore, once equipped with Advisor phrases and possibly with new mappings, any FORR-based system could use WHY to produce explanations. For example, Hoyle is a FORR-based system that learns to play many two-person finite-board games expertly (Epstein 2001). For Hoyle, WHY could explain "Although I don't want to make a move that once led to a loss, I decided to do it because I really want to get closer to winning and I want to do something I've seen an expert do."

Because SemaFORR's spatial model is approximate and its Advisors are heuristic, precise natural language interpretations for numeric values are ad hoc. For Table 3, we inspected thousands of decisions, and then partitioned the computed values as appeared appropriate. We intend to fine-tune both intervals and phrasing with empirical assessment by human subjects. Because natural explanations have improved people's trust and understanding of other automated systems, we will then evaluate WHY with human subjects.

SemaFORR and WHY are both ongoing work. As heuristic planners for tier 2 are developed, we will extend WHY to incorporate plans in its explanations. We also anticipate revisions in WHY's phrasing to reflect changes in SemaFORR's possible action set. Finally, WHY could be incorporated into a more general dialogue system that would facilitate part of a broader conversation between a human collaborator and a robot. A FORR-based system for human-computer dialogue, could prove helpful there (Epstein et al. 2011).

In summary, WHY produces natural explanations for a robot's navigation decisions as it travels through a complex world. These explanations are essential for collaborative navigation and are made possible by the robot controller's cognitively-based reasoning. The approach presented here generates explanations that gauge the robot's confidence and give reasons to take an action or to prefer one action over another. As a result, a human companion receives informative, user-friendly explanations from a robot as they travel together through a large, complex world in real time.

Acknowledgements. This work was supported in part by NSF 1625843. The authors thank the reviewers for their insightful comments.

References

Aroor, A., and Epstein, S. L. In Press. Toward Crowd-Sensitive Path Planning. In *2017 AAAI Fall Symposium on Human-Agent Groups: Studies, Algorithms and Challenges*.

Barrett, D. P.; Bronikowski, S. A.; Yu, H.; and Siskind, J. M. 2017. Driving Under the Influence (of Language). *IEEE Transactions on Neural Networks and Learning Systems*.

Boularias, A.; Duvallet, F.; Oh, J.; and Stentz, A. 2016. Learning Qualitative Spatial Relations for Robotic Navigation. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, 4130–4134. AAAI Press.

Bussone, A.; Stumpf, S.; and O'Sullivan, D. 2015. The Role of Explanations on Trust and Reliance in Clinical Decision Support Systems. In *2015 International Conference on Healthcare Informatics (ICHI)*, 160–169. IEEE.

Chai, J. Y.; Fang, R.; Liu, C.; and She, L. 2016. Collaborative Language Grounding Toward Situated Human-Robot Dialogue. *AI Magazine* 37(4).

Coleman, M., and Liau, T. L. 1975. A Computer Readability Formula Designed for Machine Scoring. *Journal of Applied Psychology* 60(2):283.

Conlin, J. A. 2009. Getting Around: Making Fast and Frugal Navigation Decisions. *Progress in Brain Research* 174:109–117.

Dale, R.; Geldof, S.; and Prost, J.-P. 2005. Using Natural Language Generation in Automatic Route Description. *Journal of Research and Practice in Information Technology* 37(1):89.

Daniele, A. F.; Bansal, M.; and Walter, M. R. 2016. Natural Language Generation in the Context of Providing Indoor Route Instructions. In *Proceedings Robotics: Science and Systems Workshop on Model Learning for Human-Robot Communication*.

Dodson, T.; Mattei, N.; Guerin, J. T.; and Goldsmith, J. 2013. An English-Language Argumentation Interface for Explanation Generation with Markov Decision Processes in the Domain of Academic Advising. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 3(3):18.

Duvallet, F.; Walter, M. R.; Howard, T.; Hemachandra, S.; Oh, J.; Teller, S.; Roy, N.; and Stentz, A. 2016. Inferring Maps and Behaviors from Natural Language Instructions. In *Experimental Robotics*, 373–388. Springer.

Epstein, S. L.; Passonneau, R.; Gordon, J.; and Ligorio, T. 2011. The Role of Knowledge and Certainty in Understanding for Dialogue. In *2011 AAAI Fall Symposium Series*.

Epstein, S. L.; Aroor, A.; Evanusa, M.; Sklar, E. I.; and Parsons, S. 2015. Learning Spatial Models for Navigation. In *12th International Conference on Spatial Information Theory*, 403–425. Springer.

Epstein, S. L. 1994. For the Right Reasons: the FORR Architecture for Learning in a Skill Domain. *Cognitive Science* 18(3):479–511.

Epstein, S. L. 2001. Learning to Play Expertly: A Tutorial on Hoyle. In *Machines That Learn to Play Games*, 153–178. Nova Science Publishers, Inc.

Foo, P.; Warren, W. H.; Duchon, A.; and Tarr, M. J. 2005. Do Humans Integrate Routes into a Cognitive Map? Map-Versus Landmark-Based Navigation of Novel Shortcuts. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31(2):195–215.

Golledge, R. G. 1999. Human Wayfinding and Cognitive Maps. *Wayfinding Behavior: Cognitive Mapping and Other Spatial Processes* 5–45.

Hastie, T.; Tibshirani, R.; and Friedman, J. 2009. *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer, 2 edition.

Kennedy, W. G.; Bugajska, M. D.; Marge, M.; Adams, W.; Fransen, B. R.; Perzanowski, D.; Schultz, A. C.; and Trafton, J. G. 2007. Spatial Representation and Reasoning for Human-Robot Collaboration. In *Proceedings of the Twenty Second Conference on Artificial Intelligence*, volume 7, 1554–1559.

Khan, O.; Poupart, P.; Black, J.; Sucar, L.; Morales, E.; and Hoey, J. 2011. Automatically Generated Explanations for Markov Decision Processes. *Decision Theory Models for Applications in Artificial Intelligence: Concepts and Solutions* 144–163.

Kulesza, T.; Stumpf, S.; Burnett, M.; Yang, S.; Kwan, I.; and Wong, W.-K. 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)*, 3–10. IEEE.

Landsiedel, C.; Rieser, V.; Walter, M.; and Wollherr, D. 2017. A Review of Spatial Reasoning and Interaction for Real-World Robotics. *Advanced Robotics* 31(5):222–242.

Lim, B. Y.; Dey, A. K.; and Avrahami, D. 2009. Why and Why Not Explanations Improve the Intelligibility of Context-Aware Intelligent Systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2119–2128. ACM.

Look, G. W. K. 2008. *Cognitively-Inspired Direction Giving*. Ph.D. Dissertation, Massachusetts Institute of Technology.

Malle, B. F. 1999. How People Explain Behavior: a New Theoretical Framework. *Personality and Social Psychology Review* 3(1):23–48.

Markkanen, R., and Schröder, H. 1997. *Hedging and Discourse: Approaches to the Analysis of a Pragmatic Phenomenon in Academic Texts*, volume 24. Walter de Gruyter.

Oh, J.; Howard, T. M.; Walter, M. R.; Barber, D.; Zhu, M.; Park, S.; Suppe, A.; Navarro-Serment, L.; Duvallet, F.; Boularias, A.; et al. 2016. Integrated Intelligence for Human-Robot Teams. In *International Symposium on Experimental Robotics*, 309–322. Springer.

Perera, V.; Selveraj, S. P.; Rosenthal, S.; and Veloso, M. 2016. Dynamic Generation and Refinement of Robot Verbalization. In *25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 212–218. IEEE.

Ramakrishnan, R., and Shah, J. 2016. Towards Interpretable Explanations for Transfer Learning in Sequential Tasks. In *2016 AAAI Spring Symposium Series*.

Richter, K.-F., and Duckham, M. 2008. Simplest Instructions: Finding Easy-To-Describe Routes for Navigation. *Geographic Information Science* 274–289.

Rosenthal, S.; Selvaraj, S. P.; and Veloso, M. 2016. Verbalization: Narration of Autonomous Mobile Robot Experience. In *Proceedings of IJCAI*, volume 16.

Scalise, R.; Rosenthal, S.; and Srinivasa, S. 2017. Natural Language Explanations in Human-Collaborative Systems. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, 377–378. ACM.

Spiers, H. J., and Maguire, E. A. 2008. The Dynamic Nature of Cognition During Wayfinding. *Journal of Environmental Psychology* 28(3):232–249.

Thomason, J.; Zhang, S.; Mooney, R.; and Stone, P. 2015. Learning to Interpret Natural Language Commands Through Human-Robot Dialog. In *Proceedings of the 24th International Conference on Artificial Intelligence*, 1923–1929. AAAI Press.