Modeling Human Helpfulness with Individual and Contextual Factors for Robot Planning

Amal Nanavati[1], Christoforos Mavrogiannis[1], Kevin Weatherwax[2], Leila Takayama[2], Maya Cakmak[1], Siddhartha Srinivasa[1]

[1] University of Washington, Seattle, USA

[2] University of California, Santa Cruz, USA

amaln@cs.washington.edu

1. Motivation

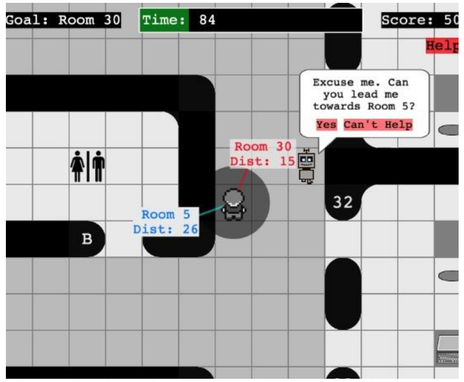

- As robots are deployed in human environments, situations will arise that they are unequipped to handle. In such situations, they can ask for **human help**.
- Yet, if a robot asks for help too frequently, or at the wrong times, **humans can get annoyed** and not help it.
- We present a model of human help, trained on data from a user study, that disaggregates the **individual** and **contextual factors** that influence whether a human helps a robot.
- We then integrate the model into a **BAMDP planning framework**, and evaluate the model and planning framework with another user study.

2. Summary

1. The robot needs help finding a room.

2. The human help model, which is updated based on past interactions with each human, allows it to predict the probability a person will help.

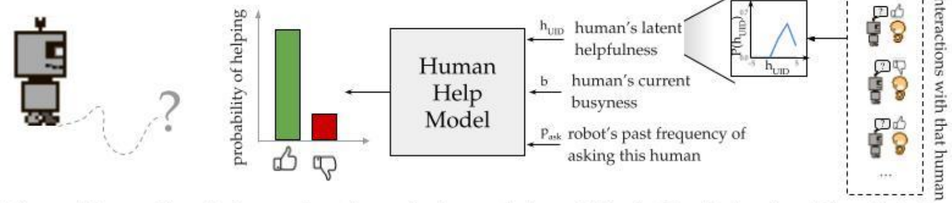

3. For each human the robot encounters, it uses the human help model to decide whether to ask them for help.

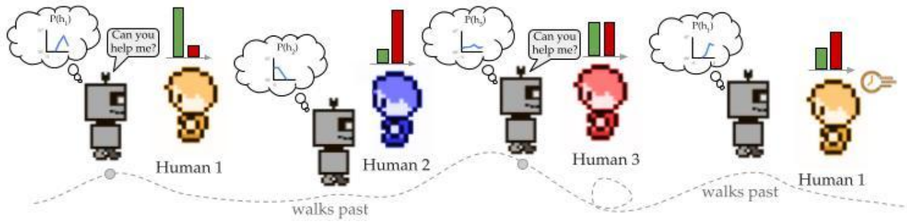

3. The Human Help Model (GLMM)

$$\mathbb{P}(\texttt{human\_helps} \mid h_{\texttt{UID}}, b, \rho_{ask}) = \frac{1}{1+\exp\left(-f(h_{\texttt{UID}}, b, \rho_{ask})\right)}, \quad (4)$$

$$f(h_{\texttt{UID}}, b, \rho_{ask}) = h_{\texttt{UID}} + c_1 + c_2 \cdot b + c_3 \cdot b \cdot \rho_{ask}, \quad (5)$$

$$h_{\texttt{UID}} \sim \mathcal{N}(0, \sigma^2), \quad (6)$$

where c_1, c_2, c_3, and σ^2 are learnt parameters.

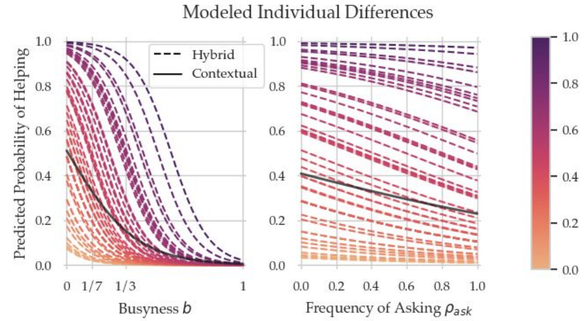

4. Planning With The Human Help Model

- Bayes-Adaptive Markov Decision Process (BAMDP)
- Transition function *uses* the human help model to predict whether the human will help.
- Learns the human's latent helpfulness over repeated interactions.
- Prior belief distribution is learnt by the generalized linear mixed model (GLMM)
- Belief update accounts for unmodeled variability in human behavior.
- Partially Observable Monte-Carlo Planning (POMCP)

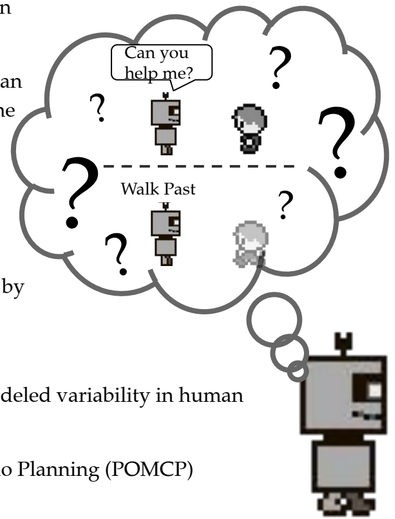

5. Results

- Our model (Hybrid) significantly outperforms baselines by a factor of 1.5x. It asks for help 1.2x fewer times, while still receiving more help on average.
- Users associated "warmth" with Hybrid significantly more than Individual, and "discomfort" with Hybrid significantly *less* than Contextual.

Metric	Hybrid (*n=50*)	Contextual (*n=25*)	Individual† (*n=25*)
Cumulative Reward	**2.74 (3.18)**	1.87 (2.95)*	1.80 (3.04)**
Num Correct Rooms	**4.29 (3.88)**	3.80 (3.70)	3.88 (3.03)
Num Asks	**7.76 (3.76)**	9.64 (4.12)***	10.40 (1.16)**
Num Help Received	**3.83 (3.91)**	3.26 (3.72)	3.48 (3.15)*
Num Help Rejected	**3.93 (1.31)**	6.38 (1.24)***	6.92 (3.36)***

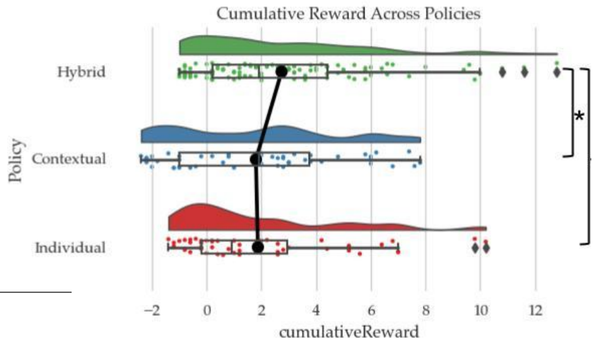

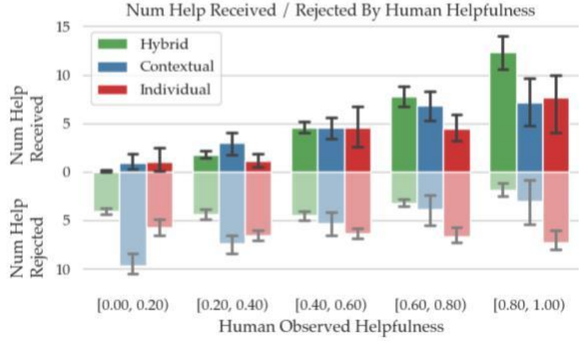

6. Acknowledgements

This work was (partially) funded by the Faithful Steward Endowed Fellowship, National Science Foundation IIS (#2007011), National Science Foundation DMS (#1839371), the Office of Naval Research, US Army Research Laboratory CCDC, Amazon, and Honda Research Institute USA

curious minded machine

PAUL G. ALLEN SCHOOL OF COMPUTER SCIENCE & ENGINEERING