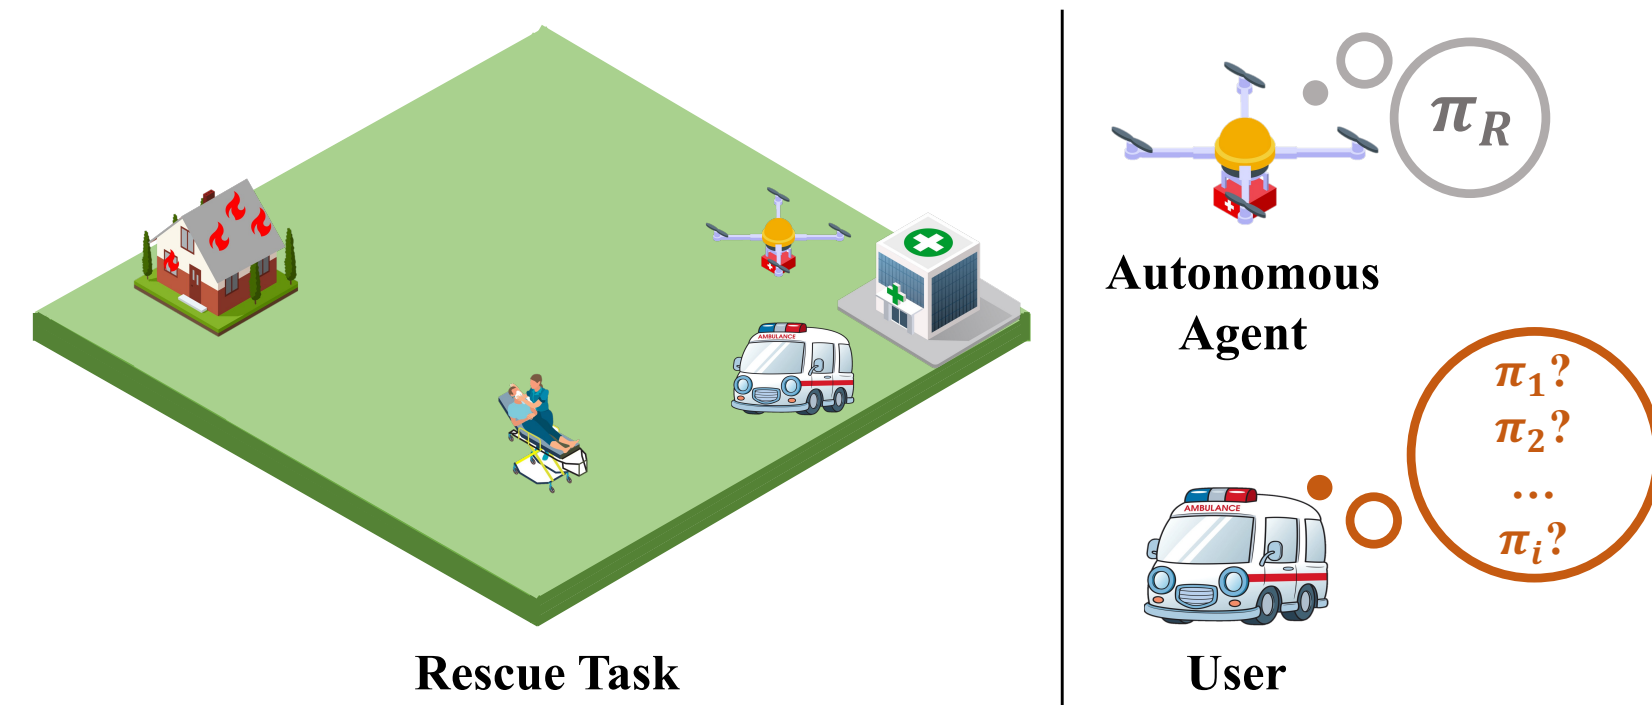


PPS: Personalized Policy Summarization for Explaining Sequential Behavior of Autonomous Agents

Peizhu “Pam” Qian,* Harrison Huang,* and Vaibhav Unhelkar
 {pqian, hhuang, unhelkar}@rice.edu

Pam is on the market for faculty + postdoc positions and Harrison is applying for PhD programs! Chat with us!

Motivating Example




A first responder wants to deploy the autonomous rescue robot to assist in disaster response. However, she is unsure how the robot behaves. To effectively deploy the robot, the first responder must be able to predict robot behavior in familiar and new situations.

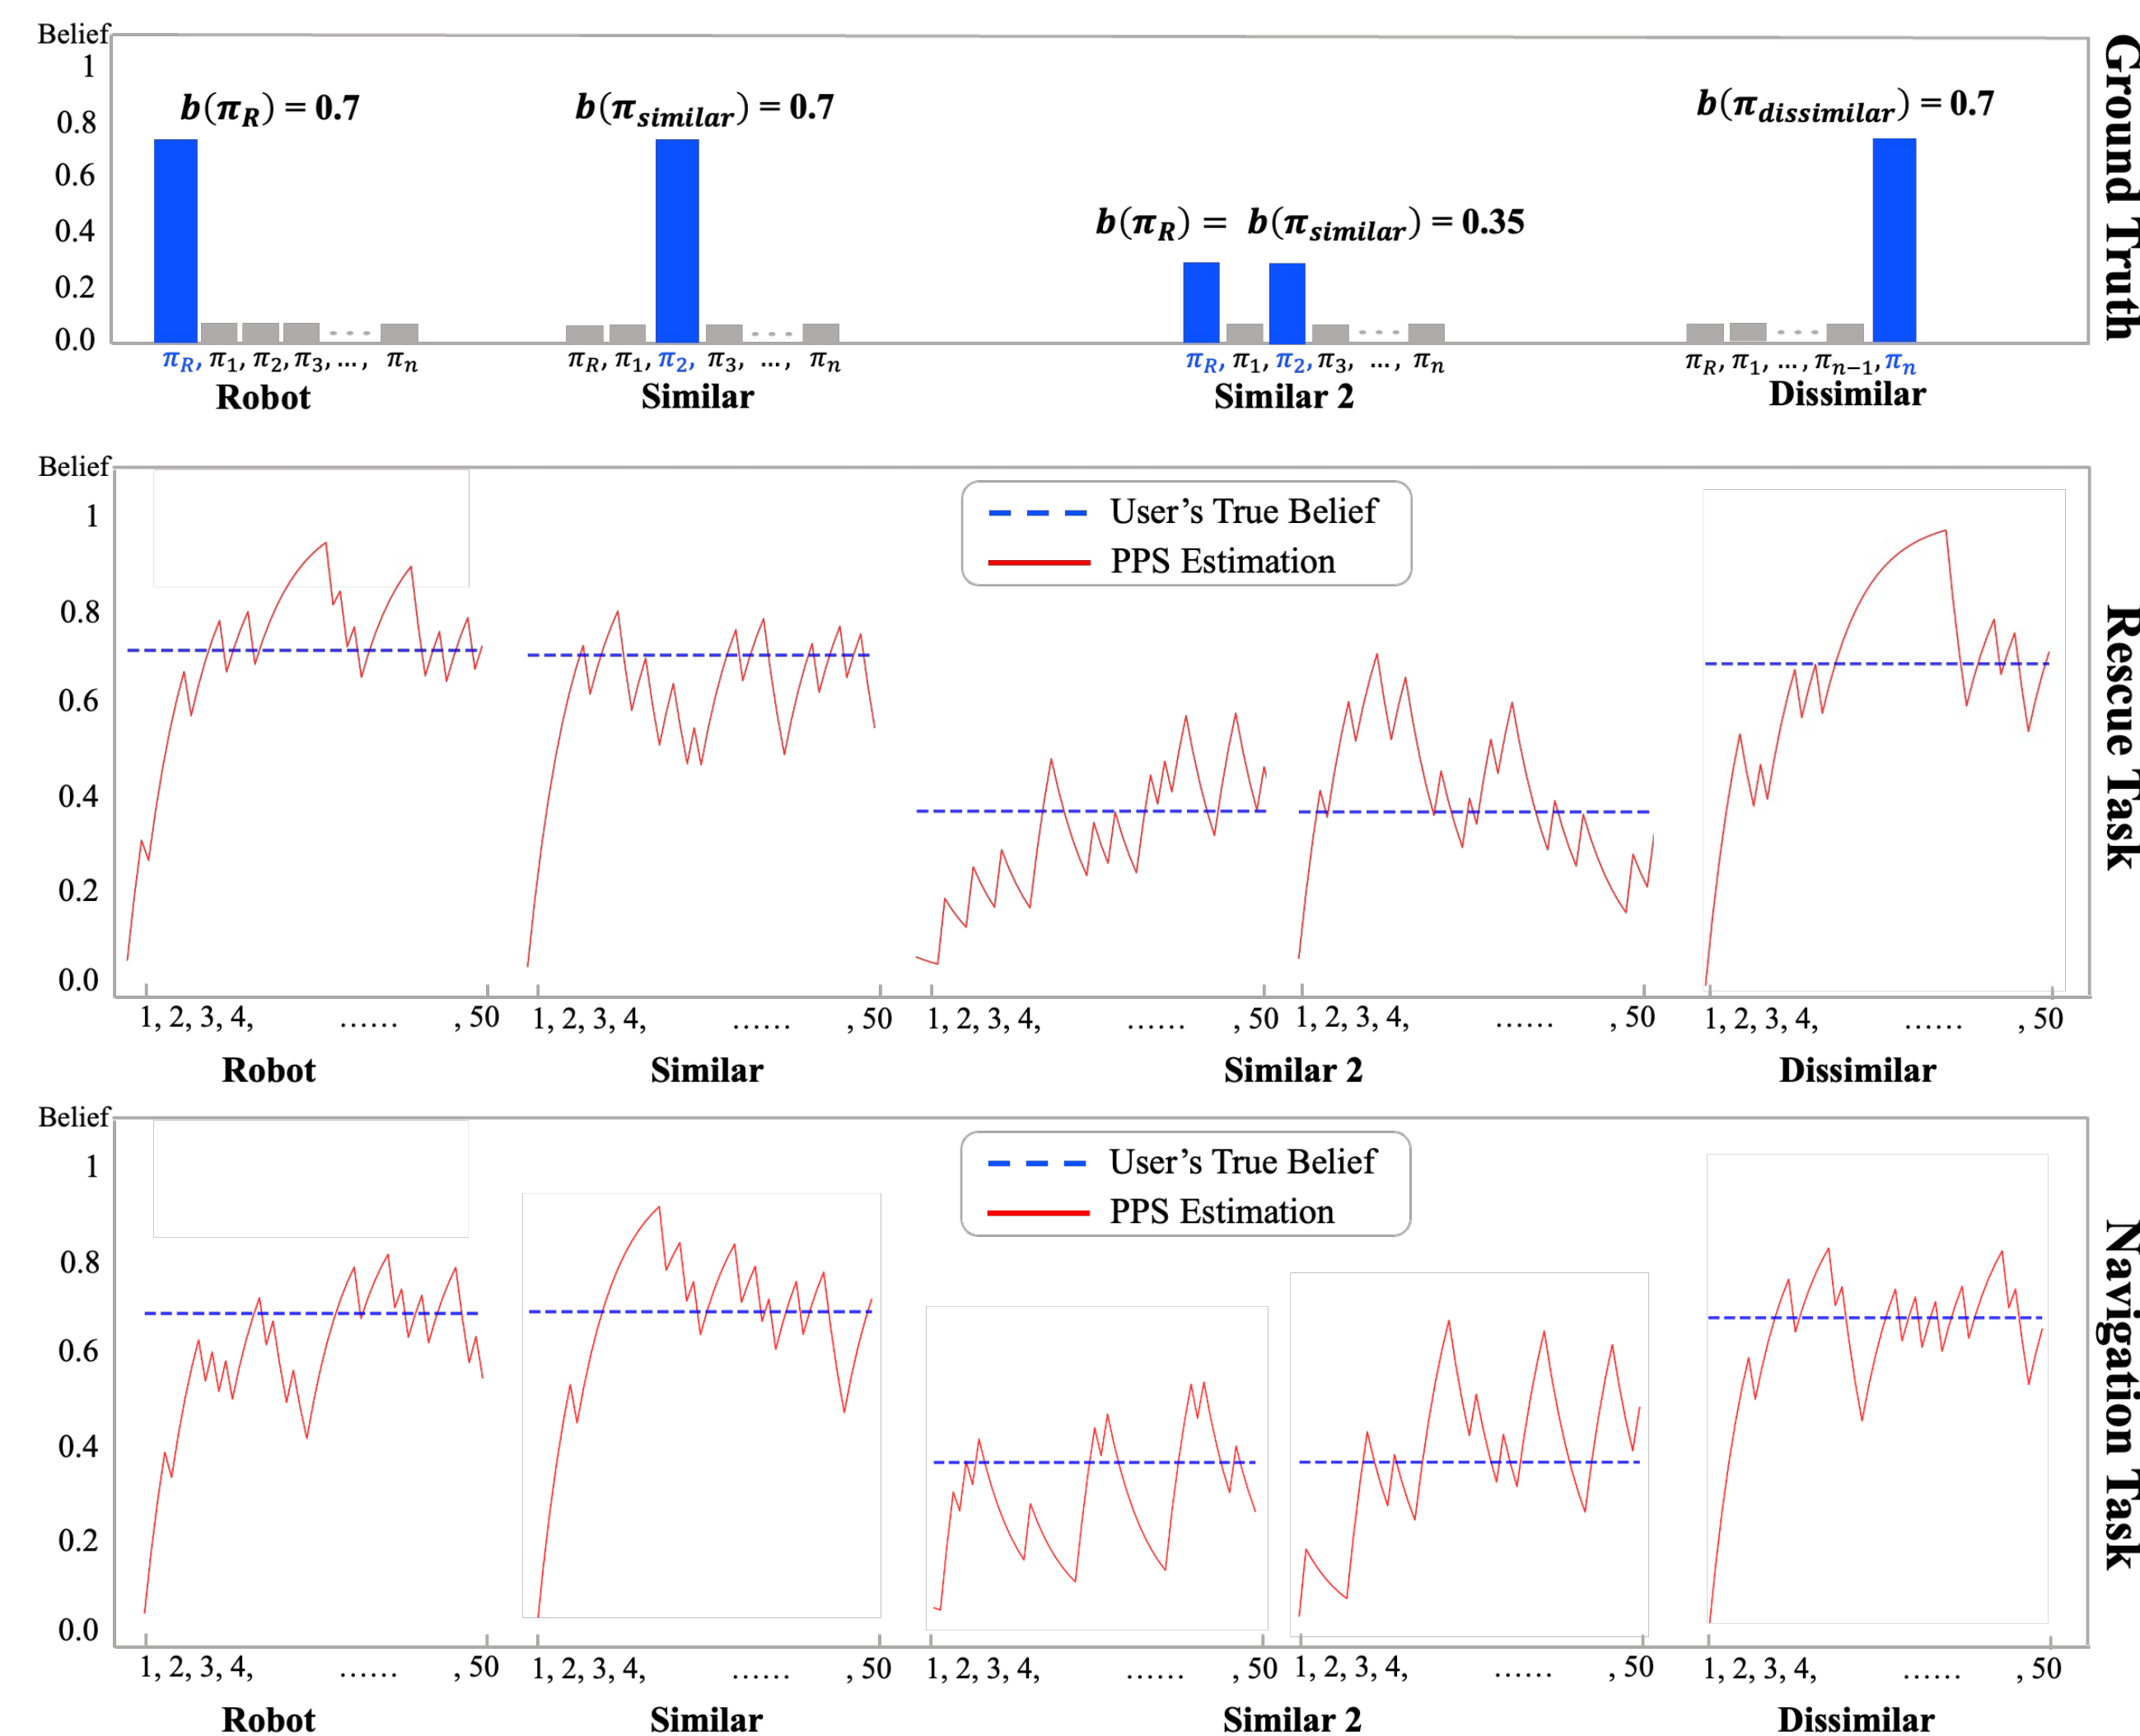
Limitations of Prior Work

	Prior	Ours
Explanations	One-fits-all	Personalized
Estimating user learning curve	Bayesian Theory of Mind (BToM)	BToM + user comprehension
Generating questions for user belief evaluation	⊗	Dirichlet estimation & entropy minimization

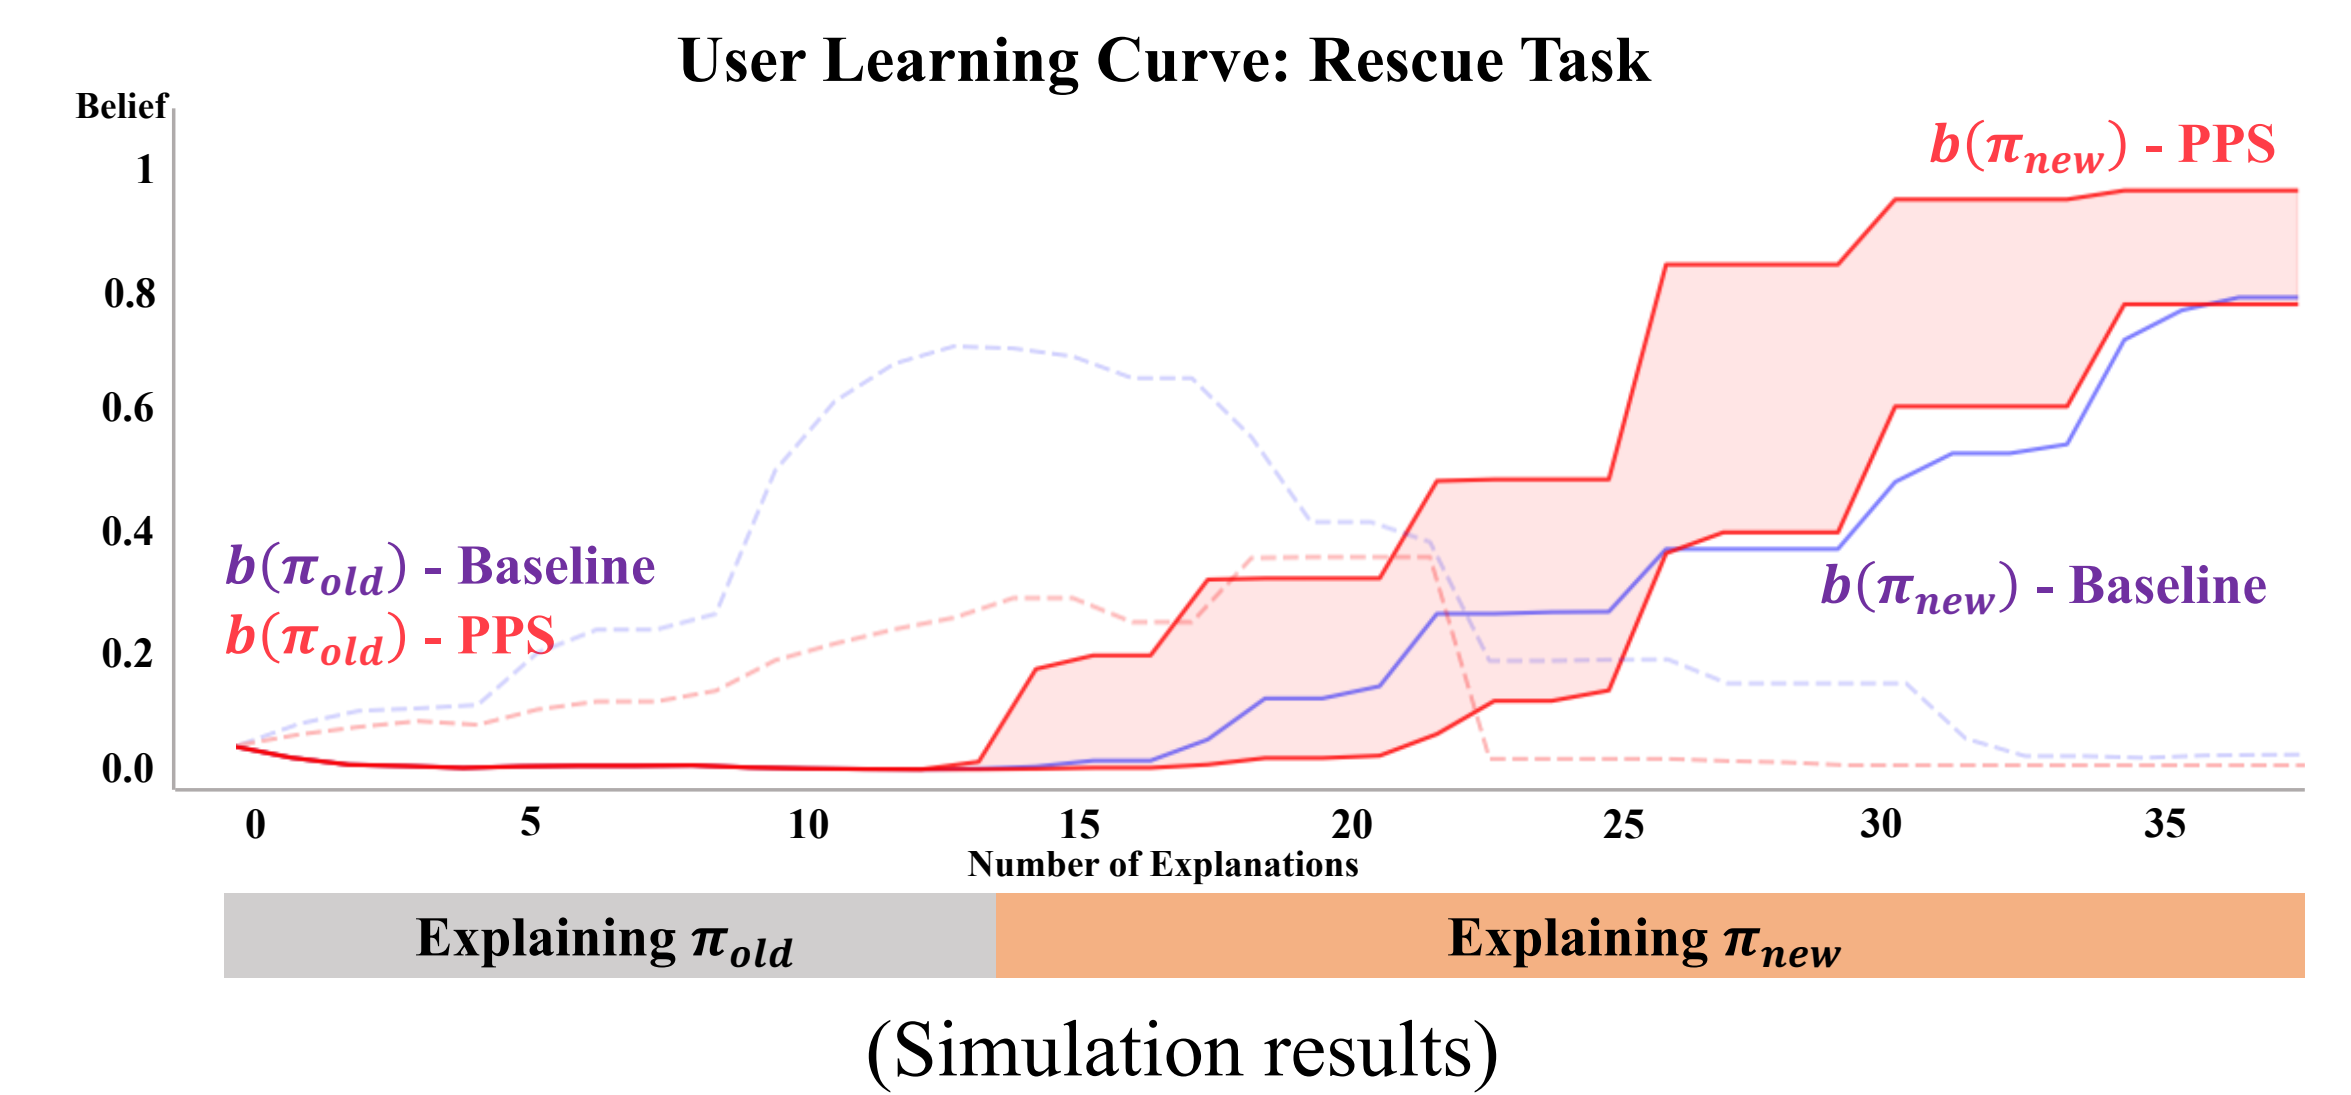
How PPS Works



Key Result 1: Belief Estimation



Key Result 2: User Learning Curve



Key Result 3: Human Experiment Scores (N=30)

