

Intelligent Agents for Virtual Simulation of Human-Robot Interaction

Ning Wang¹, David V. Pynadath¹ Unnikrishnan K.V.¹, Santosh Shankar², and Chirag Merchant¹

¹ Institute for Creative Technologies, University of Southern California
12015 Waterfront Drive, Playa Vista, CA 90094-2536, USA,

nwang@ict.usc.edu,

WWW home page: <http://people.ict.usc.edu/~nwang/>

² Apple, Inc., Cupertino, CA 95014, USA

Abstract. To study how robots can work better with humans as a team, we have designed an agent-based online testbed that supports virtual simulation of domain-independent human-robot interaction. The simulation is implemented as an online game where humans and virtual robots work together in simulated scenarios. This testbed allows researchers to carry out human-robot interaction studies and gain better understanding of, for example, how a robot's communication can improve human-robot team performance by fostering better trust relationships among humans and their robot teammates. In this paper, we discuss the requirements, challenges and the design of such human-robot simulation. We illustrate its operation with an example human-robot joint reconnaissance task.

Keywords: human robot interaction, intelligent virtual agent, social simulation

1 Introduction

Robots have become increasingly prevalent and no doubt will become an integral part of future human society. From factory robotic arms, to expressive humanoids, robots have evolved from machines operated by humans to autonomous intelligent entities that operate with humans. As robots gain complexity and autonomy, it is important yet increasingly challenging for humans to understand their decision process. Research has shown that people will more accurately trust an autonomous system, such as a robot, if they have a more accurate understanding of its decision-making process [21]. Trust is a critical element to how humans and robots perform together [22]. For example, if robots are more suited than humans for a certain task, then we want the humans to trust the robots to perform that task. If the robots are less suited, then we want the humans to appropriately gauge the robots' ability and have people perform the task manually. Failure to do so results in disuse of robots in the former case and misuse in the latter [28]. Real-world case studies and laboratory experiments show that failures in both cases are common [22].

Successful human-robot interaction (HRI) therefore relies on the robot’s ability to make its decision-making process transparent to the people it works with. However, while hand-crafted explanations have been effective in providing such transparency [10], we are interested here in pursuing a more general approach to explanation that can be reused across domains. As a first step toward that goal, we need an experimental testbed that will allow us to quantify the effectiveness of different explanation algorithms in terms of their ability to make a robot’s decision-making process transparent to humans.

There are several challenges and requirements in the design and implementation of such a testbed. The first challenge is how to model a HRI scenario that facilitates the research of robot communication. Section 2 surveys the literature on HRI simulations, and Section 3 presents the requirements that we extracted from that survey with respect to studying human-robot trust relationships. A second challenge for such a simulation is the generation of the autonomous behaviors of the robots within that scenario. The robot’s decision-making must account for the complex planning, noisy sensors, faulty effectors, etc. that complicate even single-robot execution and that are often the root of trust failures in HRI. Section 4 describes how we use a multiagent social simulation framework, PsychSim [24, 31], as the agent-based platform for our testbed. Importantly for our purposes, PsychSim includes sensitivity analysis algorithms for explanations [30] that are based in a general decision-theoretic agent framework [16].

The resulting virtual simulation thus provides an experimental testbed that allows researchers to carry out online human-subject studies and gain better understanding of how a robot’s communication can improve human-robot team performance by fostering better trust relationships among humans and their robot teammates. In this paper, we discuss the design decisions in the implementation of the agent-based online testbed that supports virtual simulation of domain-independent HRI.

2 Related Work

There is a large body of work on simulating HRI. In the review presented here, we take the perspective of the needs of a testbed specifically for studying human-robot trust. Many HRI simulations seek a high-fidelity re-creation of the physical capabilities of a robot and the physical environment it operates in. For example, Gazebo Player is a 3D simulation framework that contains several models of real robots with a variety of sensors (e.g., camera, laser scanner) [19]. Although this framework supports various kinds of dynamic interaction, dynamic objects (especially humans) are not integrated in the framework. Another high-fidelity simulation environment, USARSim, models urban search-and-rescue robots to provide a research tool for the study of HRI and multirobot coordination [23]. USARSim includes realistic simulations of the physical environment and the physical robots, focusing on tasks like maneuvering through rubble, fallen buildings, etc.

While simulation of physical interaction is important for HRI, the emphasis of our human-robot trust testbed is more on the social interaction. Thus for the time being, we instead focus on simulations that use lower-fidelity models of the physical environment, and use an agent-based simulation to highlight the robot’s decision-making challenges (e.g., planning, coordination). For example, Military Operations in Urban Terrain (MOUT) have been modeled within multiagent simulations that capture both team coordination and HRI [11]. However, these particular agents generate the behavior for both the robots and the humans. While such a simulation can provide useful insight into the impact of different coordination strategies on team performance, we instead need an *interactive* simulation to gather behavior data from human participants.

A variety of interactive simulations have modeled scenarios in which people work with a simulated robot subordinate. One environment used the ADAPT framework [37] to build a simulated marketplace, which a semi-autonomous robot navigates based on multimodal directions from human soldiers [6]. The Mixed Initiative eXperimental (MIX) testbed [1] supported a simulation of a generic military crew station to study the differential impact of autonomous systems that are teleoperated, semi-autonomous, or adaptive between the two [8]. Human operators worked with unmanned vehicles under their direction to perform reconnaissance tasks in a hostile environment. This testbed has been successful in measuring the impact of the level of the robot’s autonomy on the cognitive load and, in turn, task performance of those operators.

The cooperative nature of this joint reconnaissance task and the complementary responsibilities of humans and robots represent two critical features for our human-robot trust scenario. However, we first need to adapt the task to move the robots away from being directly supervised by a human operator and instead give them full autonomy. In other words, we wish to elevate the robot to the status of teammate, rather than subordinate. By removing the human from the supervisory role, we allow for the possibility for both misuse and disuse of the robot, which is critical in being able to induce trust failures.

Fully autonomous robots have shared a simulated space with people in scenarios like emergency response [34], assisted living [3], and joint cooking tasks [39]. The platforms used for these scenarios, like SIGVerse [39] and the HRI extension of SimVis3D [14], do provide an environment for creating simulations of joint tasks between people and autonomous robots where we could induce the needed trust failures. However, to systematically vary the robot’s domain-level and communication-level capabilities, we also require an underlying agent platform on which we can explore general-purpose algorithms for both decision-making and explanation generation.

3 HRI Design

The examination of relevant HRI simulations with respect to the needs of trust exercises leads to the following list of requirements for our testbed:

1. The simulation should encourage the human and the robot to work together as a team (as in [6, 8, 11]). The mission should require joint effort, so that neither the person nor the robot can achieve the objective by working in isolation. We thus design a joint reconnaissance mission, where the robot scouts out potential dangers to its human teammates, who are responsible for conducting a detailed search to locate a hostage and gather other important intelligence. Thus, the robot cannot achieve the search objective itself, while the human teammates run the risk of being harmed if they ignore the robot's scouting reports.
2. The simulation should encourage people to work along-side the robot, instead of just being its tele-operators (as in real-world scenarios like bomb disposal or disaster response, or in simulated scenarios [6, 8]). This means that, in the scenario, the robot should be able to complete its tasks fully autonomously without the human teammate's input. The human teammate is not required to monitor the robot's progress and give commands to the robot on what to do next at every step.
3. With the robot's being capable of acting without supervision, we must also assign the humans their own tasks; otherwise, they may revert to passively monitoring the robot's actions. Thus, we designed the simulation so that the human is also moving through the simulated environment, instead of being a stationary observer/operator of the robot. Surveys have shown that one role that robot teammates might be expected to play is that of a reconnaissance scout, on the lookout for potential threats to their human teammates [38]. We therefore designed our task so that the robot serves as exactly such an advanced scout to sniff out danger, as its human teammates follow up with their own reconnaissance tasks (e.g., searching buildings to locate a hostage). By placing a time limit on completing the joint mission, we incentivize the human (and the robot) to continually pursue their own tasks in parallel.
4. The simulation should encourage communication between humans and the robot, so that the robot can take an active role in establishing trust. To achieve this goal, we took away the interface elements that would provide users with constant situational awareness about the environment and the robots. For example, after the robot scouts a building, we could simply mark the building on the map as red or green to signal whether it is safe or unsafe for human teammates to enter. However, this would take away the opportunity for the robot to directly communicate with its teammates. Similarly, we could directly show the human team members the robot's "raw" sensor readings, but again, this would take away an opportunity for the robot to explain its decisions based on those readings (not to mention potentially creating cognitive overload for the humans). Instead, the human teammates receive information (e.g., assessments, explanations) from the robot only when it actively communicates to them.
5. The task performed by the human and the robot in the simulation should introduce sources of distrust (e.g. robot malfunction, uncertainty in the environment, etc.). While people may occasionally distrust even a robot that never makes a mistake, we would rather ensure the occurrence of mistakes

that threaten the trust relationship. Controlling these potential trust failures gives us a better opportunity to research ways to use explanation to (re)establish trust. We therefore design the robot so that it will not have perfect knowledge of the environment and add variable limitations to the robot’s sensors (e.g., varying error rates in its detection of dangerous chemicals). Studies have shown that the frequency and significance of errors can greatly impact user trust in an online system (e.g., a series of small mistakes is worse than one big one [7]). We therefore expect that controlling the error dimensions will be essential in isolating these trust failures and identifying the best explanation algorithms for repairing them.

6. While surveys can provide insight into the human-robot trust relationship, we also want more objective measures of trust in the form of behavioral data. Prior studies have used the “take-over” and “hand-over” behavior a human supervisor does to a robot worker (e.g. takes over a task the robot is currently performing and does it by himself instead) as a measure of the trust or distrust he had in the robot [41]. We follow a similar model in constructing our scenario to include behavioral indicators of disuse and misuse deriving from lack of trust and too much trust, respectively, in the robot. For example, if the human teammate follows the robot’s recommendation (e.g., avoids going into a building that the robot said was unsafe), this behavior would be an objective indicator of trust. In contrast, we might infer a lack of trust if the human asks the robot to re-search an area that the robot has already searched. Additionally, our user interface allows the human to choose to directly view the camera feed of the robot. Using this function can be an indication that the human teammates wishes to oversee the robot’s behavior and thus, a lack of trust.
7. To ensure that the human teammate’s behavior can be indicative of his trust in the robot (e.g., following the robot’s recommendation), the robot’s mistakes (e.g., incorrectly identifying a building as safe) should have an inherent cost to its human teammates. Otherwise, there will be no reason for the human teammates to *not* act based on the robot’s communication. Studies have shown that people will follow the requests of even an incompetent robot if the negative consequences are somewhat trivial [36]. We therefore design our game so that inappropriate trust of the robot can potentially lead to failure to complete the mission and to even “death” of the player.

4 Agent-Based Simulation of HRI

To meet these requirements, we have implemented an agent-based online testbed that supports virtual simulation of domain-independent HRI. Our agent framework, PsychSim [24, 31], combines two established agent technologies—decision-theoretic planning [16] and recursive modeling [12]. Decision-theoretic planning provides an agent with quantitative utility calculations that allow agents to assess tradeoffs between alternative decisions under uncertainty. Recursive modeling gives the agents a *theory of mind* [40], allowing them to form beliefs about

the human users’ preferences, factor those preferences into the agent’s own decisions, and update its beliefs in response to observations of the user’s decisions. The combination of decision theory and theory of mind within a PsychSim agent has proven to be very rich for modeling human decision-making across a wide variety of social and psychological phenomena [32]. This modeling richness has in turn enabled PsychSim agents to operate in a variety of human-agent interaction scenarios [15, 17, 18, 26, 27].

PsychSim agents generate their beliefs and behaviors by solving partially observable Markov decision problems (POMDPs) [9, 16]. The POMDP model’s quantitative transition probabilities, observation probabilities, and reward functions are a natural fit for our application domain, and they have proven successful in both robot navigation [4, 20] and HRI [29]. In our own work, we have used POMDPs to implement agents that acted as 24/7 personal assistants that teamed with researchers to handle a variety of their daily tasks [5, 33]. In precise terms, a POMDP is a tuple, $\langle S, A, T, \Omega, O, R \rangle$, that we describe in terms of our human-robot team.

The *state*, S , consists of objective facts about the world, some of which may be hidden from the robot itself. By using a *factored* state representation [2, 13], the model maintains separate labels and values of each feature of the state, such as the separate locations of the robot, its human teammate, the hostage, and the dangerous chemicals. The state also includes feature-value pairs that represent the respective health levels of the teammate and hostage, any current commands from the teammate, and the accumulated time cost so far. Again, while this state represents the true value of all of these features, the robot cannot directly access this true state.

The robot’s available *actions*, A , correspond to the possible decisions it can make. Given its search mission, the robot’s primary decision is where to move to next. We divide the environment into a set of discrete waypoints, so the robot’s action set includes potentially moving to any of them. The robot also makes a decision as to whether to declare a location as safe or unsafe for its human teammate. For example, if the robot believes that dangerous chemicals are at its current location, then it will want its teammate to take adequate preparations before entering. Because there is a time cost to such preparations, the robot may instead decide to declare the location safe, so that its teammates can more quickly complete their own reconnaissance tasks in the building.

The state of the world changes in response to the actions performed by the robot. We model these dynamics using a *transition probability*, T function that captures the possibly uncertain effects of these actions on the subsequent state. We simplify the robot’s navigation task by assuming that a decision to move to a specific waypoint succeeds deterministically. However, we could relax this assumption to decrease the robot’s movement ability, as is done in more realistic robot navigation models [4, 20]. The robot’s recommendation decision affects the health of its teammate and the hostage, although only stochastically, as there is no guarantee that the teammate will follow the recommendation. Instead, a

recommendation that a building is safe (unsafe) has a high (low) probability of decreasing the teammate’s health if there are, in fact, chemicals present.

As already mentioned, the robot and human teammate have only indirect information about the true state of the world. Within the POMDP model, this information comes through a subset of possible *observations*, Ω , that are probabilistically dependent (through the *observation function*, O) on the true values of the corresponding state features. We make some simplifying assumptions, namely that the robot can observe the location of itself and its teammate with no error (e.g., via GPS).

However, it cannot directly observe the locations of the hostage or dangerous chemicals. Instead, it receives a local reading about their presence (or absence) at its current location. For example, if dangerous chemicals are present, then the robot’s chemical sensor will detect them with a high probability. However, there is also a lower, but nonzero, probability that the sensor will *not* detect them. In addition to such a false negative, there is also a potential false positive reading, where there is a low, but nonzero, probability that it will detect chemicals even if there are none present.

Partial observability gives the robot only a subjective view of the world, where it forms beliefs about what it *thinks* is the state of the world, computed via standard POMDP *state estimation* algorithms. For example, the robot’s beliefs include its subjective view on the location of the hostage, potentially capturing statements like: “There is an 80% probability that the hostage is being held at my current location.” or “If you visit this waypoint, there is a 60% chance that you will be exposed to dangerous chemicals.” By varying the accuracy of the robot’s observation models, we will decrease the accuracy of its beliefs and, subsequently, its recommendations to its human teammates.

On the other hand, the structured dependency structure of the observation function gives the robot explicit knowledge of the uncertainty in its own observations. It can thus communicate its noisy sensor model to its human teammates, potentially making statements like, “My chemical sensor has a 20% chance of generating a false negative.” Therefore, even though a less capable robot’s recommendations may be less reliable to its teammate, the robot will be able to explicitly explain that inaccuracy in a way that mitigates the impact to the trust relationship.

PsychSim’s POMDP framework instantiates the human-robot team’s mission objectives as a *reward*, R , that maps the state of the world into a real-valued evaluation of benefit for the agent. The highest reward is earned in states where the hostage is rescued and all buildings have been explored by the human teammate. This reward component incentivizes the robot to pursue the overall mission objective. There is also an increasingly positive reward associated with level of the human teammate’s health. This reward component punishes the robot if it fails to warn its teammate of dangerous buildings. Finally, there is a negative reward that increases with the time cost of the current state. This motivates the robot to complete the mission as quickly as possible. By providing different weights to these goals, we can change the priorities that the robot assigns to them. For

example, by lowering the weight of the teammate’s health reward, the robot may allow its teammate to search waypoints that are potentially dangerous, in the hope of finding the hostage sooner. Alternatively, lowering the weight on the time cost reward might motivate the robot to wait until being almost certain of a building’s threat level (e.g., by repeated observations) before recommending that its teammate visit anywhere.

The robot can arrive at such policies based on its POMDP model of the world by determining the optimal action based on its current beliefs about the state of the world [16]. Rather than perform an offline computation of a complete optimal policy over all possible beliefs, we instead take an online approach so that the robot makes optimal decisions with respect to only its current beliefs [35]. The robot uses a bounded lookahead procedure that seeks to maximize expected reward by simulating the dynamics of the world from its current belief state. In particular, the robot first uses the transition function to project the immediate effect of a candidate action, and then projects a finite number of steps into the future, weighing each state against its reward function. Following such an online algorithm, the robot can thus choose the optimal action with respect to its current beliefs.

On top of this POMDP layer, PsychSim provides a suite of algorithms that are useful for studying domain-independent explanation. By exploring variations of these algorithms within PsychSim’s scenario-independent language, we ensure that the results can be re-used by other researchers studying other HRI domains, especially those using POMDP-based agents or robots. To begin with, PsychSim agents provide support for transparent reasoning that is a requirement for our testbed. PsychSim’s original purpose was human-in-the-loop social simulation. To identify and repair errors in a social-simulation model, the human user must be able to understand the POMDP reasoning process that the agents went through in generating their simulation behavior. In other words, the agents’ reasoning must be transparent to the user. To this end, PsychSim’s interface made the agent’s reasoning process available to the user, in the form of a branching tree representing its expected value calculation. The user could expand branches as needed to drill down into the agent’s considerations across possible decisions and outcomes.

This tree provided a maximum amount of transparency, but it also provided a high volume of data, often obscuring the most salient features from the analyst. Therefore, PsychSim imposes a *piecewise* linear structure on the underlying agent models that allows it to quantify the degree to which state features, observations, and goals are salient to a given decision [30]. PsychSim exploits this capability to augment the agent’s reasoning trace by highlighting points of possible interest to the user. For example, the interface can identify the belief that the decision is most sensitive to (e.g., quantifying how saving time along a particular route outweighs the increased threat level). We have some anecdotal evidence that the identification of such critical points was useful in previous applications like human-in-the-loop modeling and tutor recommendations.

In this work, we apply this capability to the robot’s explanations to its human teammate. In explaining its recommendation that a certain building is safe, the robot can use this sensitivity analysis to decide whether the most salient reward component is the minimization of time cost or the maximization of teammate health. It can then easily map the identified motivation into natural-language expression. Similarly, it can use its lookahead process to generate a natural-language expression of the anticipated consequences to its teammates who violate its recommendation—e.g., “If you visit this location, you will be exposed to the toxic chemicals that are here, and your health will suffer.” By implementing robots that use different explanations of its decision-making process, we can quantify the differential impact that they have on human-robot trust and team performance.

5 Discussion

During the design process of an interactive simulation, there is a delicate balance between simulation and game. We learned to maintain this balance to make sure that the simulation serves our purpose of a testbed for studying human-robot trust. For example, we leave out common game elements like scoring, using mission success/failure as a performance indicator. This encourages the human teammate to focus on the mission with the robot, instead of trying to maximize a score that’s indicative of his personal performance. We also omitted the usual game elements that help game players’ situational awareness, but discourage communication between players and robot, as we observed in our early playtesting.

Our immediate next step is to use the testbed to gather data on how a robot’s explanations of its decision process impact human-robot trust and team performance. The explanations are currently provided by the robot during the mission. We are planning to extend the robot’s explanation to continue after the mission is completed. This offers the robot an opportunity to “repair” the trust relationship with its teammate, particularly when the mission ends in failure.

The current robot only interacts with people who are its teammates. However, robots in the real world will often have to interact with people who do not share its same mission objective. A future variation of our scenario can include, for example, civilian bystanders in the town where the mission is carried out. The relationships between the robot and people in these different roles will call for different explanation strategies used by the robot. For example, the robot may not want to offer explanations of its decisions to civilians in order to maintain social distance and relative power. The need to maintain social distance will likely engender additional considerations of communication tactics like politeness.

Finally, we are exploring the transition of the scenario from a simulated robot to a physical one. Compared to virtual simulations, teaming up with a physical robot that operates in the same space as a human can potentially increase the stakes of trusting the robot. Additionally, we expect this physical testbed to elevate certain dimensions (e.g., robot embodiment) in importance, as well as

providing a higher-fidelity testbed for studying the factors that impact human-robot trust.

References

1. D. Barber, L. Davis, D. Nicholson, N. Finkelstein, and J. Y. Chen. The mixed initiative experimental (MIX) testbed for human robot interactions with varied levels of automation. Technical report, DTIC Document, 2008.
2. C. Boutilier, R. Dearden, and M. Goldszmidt. Stochastic dynamic programming with factored representations. *Artificial Intelligence*, 121(1):49–107, 2000.
3. T. Braun, J. Wettach, and K. Berns. A customizable, multi-host simulation and visualization framework for robot applications. In *Recent Progress in Robotics: Viable Robotic Service to Human*, pages 357–369. Springer, 2008.
4. A. R. Cassandra, L. P. Kaelbling, and J. A. Kurien. Acting under uncertainty: Discrete Bayesian models for mobile-robot navigation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, pages 963–972, 1996.
5. H. Chalupsky, Y. Gil, C. A. Knoblock, K. Lerman, J. Oh, D. V. Pynadath, T. A. Russ, and M. Tambe. Electric Elves: Agent technology for supporting human organizations. *AI Magazine*, 23(2):11–24, 2002.
6. J. Cockburn, Y. Solomon, M. Kapadia, and N. Badler. Multi-modal human robot interaction in a simulation environment. Technical report, University of Pennsylvania, 2013.
7. C. L. Corritore, B. Kracher, and S. Wiedenbeck. On-line trust: Concepts, evolving themes, a model. *Int. J. Hum.-Comput. Stud.*, 58(6):737–758, June 2003.
8. K. Cosenzo, J. Chen, L. Reinerman-Jones, M. Barnes, and D. Nicholson. Adaptive automation effects on operator performance during a reconnaissance mission with an unmanned ground vehicle. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 54, pages 2135–2139. SAGE Publications, 2010.
9. P. Doshi and D. Perez. Generalized point based value iteration for interactive POMDPs. In *Proceedings of the Conference on Artificial Intelligence*, pages 63–68, 2008.
10. M. T. Dzindolet, S. A. Peterson, R. A. Pomranky, L. G. Pierce, and H. P. Beck. The role of trust in automation reliance. *International Journal of Human-Computer Studies*, 58:697–718, 2003.
11. R. E. Giachetti, V. Marcelli, J. Cifuentes, and J. A. Rojas. An agent-based simulation model of human-robot team performance in military environments. *Systems Engineering*, 16(1):15–28, 2013.
12. P. J. Gmytrasiewicz and E. H. Durfee. A rigorous, operational formalization of recursive modeling. In *Proceedings of the International Conference on Multi-Agent Systems*, pages 125–132, 1995.
13. C. Guestrin, D. Koller, R. Parr, and S. Venkataraman. Efficient solution algorithms for factored mdps. *Journal of Artificial Intelligence Research*, 19:399–468, 2003.
14. J. Hirth, S. A. Mehdi, N. Schmitz, and K. Berns. Development of a simulated environment for human-robot interaction. *TELKOMNIKA (Telecommunication Computing Electronics and Control)*, 9(3):465–472, 2011.
15. W. L. Johnson and A. Valente. Tactical language and culture training systems: Using AI to teach foreign languages and cultures. *Artificial Intelligence Magazine*, 30(2), 2009.

16. L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134, 1998.
17. J. M. Kim, J. Randall W. Hill, P. J. Durlach, H. C. Lane, E. Forbell, M. Core, S. Marsella, D. Pynadath, and J. Hart. BiLAT: A game-based environment for practicing negotiation in a cultural context. *International Journal on Artificial Intelligence in Education: Special Issue on Ill-Defined Domains*, 19(3):289–308, 2009.
18. J. Klatt, S. Marsella, and N. Krämer. Negotiations in the context of AIDS prevention: An agent-based model using theory of mind. In *Proceedings of the International Conference on Intelligent Virtual Agents*, 2011.
19. N. Koenig and A. Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pages 2149–2154. IEEE, 2004.
20. S. Koenig and R. Simmons. Xavier: A robot navigation architecture based on partially observable Markov decision process models. In D. Kortenkamp, R. P. Bonasso, and R. R. Murphy, editors, *Artificial Intelligence Based Mobile Robotics: Case Studies of Successful Robot Systems*, pages 91–122. MIT Press, 1998.
21. J. Lee and N. Moray. Trust, control strategies and allocation of function in human-machine systems. *Ergonomics*, 35(10):1243–1270, 1992.
22. J. D. Lee and K. A. See. Trust in automation: Designing for appropriate reliance. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 46(1):50–80, 2004.
23. M. Lewis, J. Wang, and S. Hughes. USARSim: Simulation for the study of human-robot interaction. *Journal of Cognitive Engineering and Decision Making*, 1(1):98–120, 2007.
24. S. C. Marsella, D. V. Pynadath, and S. J. Read. PsychSim: Agent-based modeling of social interactions and influence. In *Proceedings of the International Conference on Cognitive Modeling*, pages 243–248, 2004.
25. R. C. Mayer, J. H. Davis, and F. D. Schoorman. An integrative model of organizational trust. *The Academy of Management Review*, 20(3):709–734, 1995.
26. R. McAlinden, A. Gordon, H. C. Lane, and D. Pynadath. UrbanSim: A game-based simulation for counterinsurgency and stability-focused operations. In *Proceedings of the AIED Workshop on Intelligent Educational Games*, 2009.
27. L. C. Miller, S. Marsella, T. Dey, P. R. Appleby, J. L. Christensen, J. Klatt, and S. J. Read. Socially optimized learning in virtual environments (SOLVE). In *Proceedings of the International Conference on Interactive Digital Storytelling*, 2011.
28. R. Parasuraman and V. Riley. Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2):230–253, 1997.
29. J. Pineau, M. Montemerlo, M. Pollack, N. Roy, and S. Thrun. Towards robotic assistants in nursing homes: Challenges and results. *Robotics and Autonomous Systems*, 42(3):271–281, 2003.
30. D. V. Pynadath and S. C. Marsella. Fitting and compilation of multiagent models through piecewise linear functions. In *Proceedings of the International Conference on Autonomous Agents and Multi Agent Systems*, pages 1197–1204, 2004.
31. D. V. Pynadath and S. C. Marsella. PsychSim: Modeling theory of mind with decision-theoretic agents. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1181–1186, 2005.
32. D. V. Pynadath, M. Si, and S. C. Marsella. Modeling theory of mind and cognitive appraisal with decision-theoretic agents. In J. Gratch and S. Marsella, editors,

- Social emotions in nature and artifact: Emotions in human and human-computer interaction*, pages 70–87. Oxford University Press, 2014.
33. D. V. Pynadath and M. Tambe. Electric Elves: Adjustable autonomy in real-world multi-agent environments. In K. Dautenhahn, A. Bond, D. Canamero, and B. Edmonds, editors, *Socially Intelligent Agents: Creating Relationships with Computers and Robots*, chapter 12, pages 101–108. Kluwer, 2002.
 34. P. Robinette and A. M. Howard. Trust in emergency evacuation robots. In *Proceedings of the IEEE International Symposium on Safety, Security, and Rescue Robotics*, pages 1–6. IEEE, 2012.
 35. S. Ross, J. Pineau, S. Paquet, and B. Chaib-Draa. Online planning algorithms for POMDPs. *Journal of Artificial Intelligence Research*, 32:663–704, 2008.
 36. M. Salem, G. Lakatos, F. Amirabdollahian, and K. Dautenhahn. Would you trust a (faulty) robot?: Effects of error, task type and personality on human-robot cooperation and trust. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI '15*, pages 141–148, New York, NY, USA, 2015. ACM.
 37. A. Shoulson, N. Marshak, M. Kapadia, and N. I. Badler. Adapt: The agent development and prototyping testbed. *IEEE Transactions on Visualization and Computer Graphics*, 20(7):1035–1047, 2014.
 38. C. C. Swiecicki, L. R. Elliott, and R. Wooldridge. Squad-level soldier-robot dynamics: Exploring future concepts involving intelligent autonomous robots. Technical Report ARL-TR-7215, Army Research Laboratory, 2015.
 39. J. T. C. Tan and T. Inamura. Sigverse: A cloud computing architecture simulation platform for social human-robot interaction. In *Proceedings of the International Conference on Robotics and Automation*, pages 1310–1315. IEEE, 2012.
 40. A. Whiten, editor. *Natural Theories of Mind*. Basil Blackwell, Oxford, UK, 1991.
 41. A. Xu and G. Dudek. Optimo: Online probabilistic trust inference model for asymmetric human-robot collaborations. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI '15*, pages 221–228, New York, NY, USA, 2015. ACM.