LEVERAGING USER PREFERENCES FOR ADAPTIVE DECISION-MAKING IN HUMAN-AGENT INTERACTION

ACTOR-NETWORK THEORY

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Engineering

By
Jinhong Zhao

November 8, 2024

Technical Team Members:
Sujan Sarker

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISORS

Ben Laugelli, Department of Engineering and Society

Tariq Iqbal, Department of Computer Science

Introduction

In recent years, research and conferences on human-robot interaction have increased. Humans and robots working together to achieve a shared goal have gained much attention (Gao, et al. 2024). Humans can change their behavior and decisions based on surrounding circumstances, such as changes in preferences, performance, and the emotional state of others, to improve successful communication and collaboration. Creating such a flexible artificial intelligence (AI) environment is difficult because robots can only mimic human cognitive processes. Still, they cannot feel human emotions (Gillet, Vazquez Andrist Leite & Sebo 2024).

In this research, we propose the development of how compliant robot behavior can affect team performance and teamwork when humans interact with a robot in creative tasks, such as collaborative storytelling. It can offer "considerate" interactions that meet human expectations but are impossible at scale with a live agent. In the technical aspect of this project, we will work to improve robot decisions using Large Language Models (LLM) to estimate human preferences and integrate them into dynamic Markov Decision Processes (MDP) with these learned features. These technological developments raise significant ethical and societal issues about autonomy, trust, and emotional involvement in human-robot relationships. The broader implications of these advancements must be addressed. This study will also investigate the challenges and failures of humanoid robots, like SoftBank's Pepper, in sensitive domains such as healthcare through the lens of Actor-Network Theory (ANT).

I will draw on the STS framework of actor-network theory to examine how robots built by SoftBank, which have no preferences or emotions, influence managerial decisions and design choices in the healthcare industry. The STS actor-network theory paradigm will allow me to investigate the impact of SoftBank-style robots on healthcare management and product development choices despite their lack of feelings and preferences. If SoftBank focuses only on the functionality of the humanoid robot but does not consider human preference or the ability to feel others' emotions, then it could fail to sustain the key social factors necessary to provide valuable and convenient products for the healthcare industry. Because the challenge of improving human-robot collaboration is sociotechnical, attention must be paid to both technical and social aspects to be successfully achieved. In what follows, I set out two related research proposals: a technical project proposal for developing this improved decision-making framework by integrating human preferences into robotic decision-making processes (Choudhury, Swamy, Hadfield-Menell, & Dragan, 2019), and an STS project proposal for examining the impact of trust and autonomy in human-robot collaboration, particularly in creative tasks like collaborative storytelling (Leite et al., 2015). In the present study, we investigated a storytelling task for cooperation between human and robot teams (one robot and two humans).

Technical Project Proposal

The core goal of human-robot interaction is to create robots that can coexist, work with, and work alongside humans in a safe, productive, and dual-beneficial way. Our research resolves explicitly the difficulty of letting robots adapt their real-world behavior to human preferences in collaborative tasks. This technical design focuses on developing a structure and algorithm that allows robots to estimate human preference and incorporate it into decision-making processes dynamically. In the traditional robotics approach, robots follow pre-programmed or pre-defined tasks without thinking about latent factors. Engineers are considering using machine learning to allow robots to make decisions based on different scenarios. However, they need more consideration for predicting human behavior and taking further action (Yang et al. 2024). In human-robot collaborative environments, robots cannot accurately understand and respond to human preferences, which reduces work efficiency, team satisfaction, and alignment (Koppula, Jain, & Saxena, 2016). Traditional methods of decision-making in robots are limited when it comes to dynamic, real-time collaborative environments since these systems either ignore human preferences or need extensive prior knowledge and data to represent them accurately, which seems achievable until now. This research aims to minimize the discrepancy between what humans expect robots to do and how they actually behave. This technical solution finds use in industry, education, and healthcare. This approach makes it relevant for tasks requiring real-time modification, including teaching assistance or partners in industrial settings, because it helps robots modify their behavior with human preferences. Estimating human preferences helps the robot organize its decision-making process as a Markov decision process. It guides the change of the transition probabilities. The robot solves the Markov decision process to add to expected benefits, including team performance and narrative coherence.

Our research claims that combining large language models with Markov Decision Processes enables robots to address these limitations for communication and improve mutual coherence. Robots can change their behavior to better match what people expect them to do by guessing those choices and using them in probabilistic decision-making systems. This reduces the difference between what people think robots should do and what they do. The large language model is a key component in this project, which a large language model is a kind of AI that analyzes and comprehends text using deep learning. This project uses a computational approach of large language models, such as GPT-4o, to predict what people will do regarding collaborative activities, thereby structuring the robot decision-making on the Markov decision process (Shi, Bürkner, & Bulling, 2023). The study posits that human knowledge and preferences are co-variable in decision management. The large language model tackles this task as a kind of question-answering (Shi et al., 2024). This probability of distribution of human behaviors, generated by the large language model, is then incorporated into a Markov decision process model to determine the best robot behavior at each decision point. As the work progresses, this fact helps the robot better understand human preferences. It then uses these preference estimates to methodically change its "mind" invariant properties, adjusting to human preferences so that teams work better and users are happier. For example, in a joint storytelling task, the robot may adjust the narratives chosen as it sees what stories are preferred by its human partner and in which language they communicate, thus increasing mutual coherence and engagement with storytelling.

The proposed approach is evaluated in a human-robot collaborative storytelling challenge where the robot interacts with two human participants to co-develop a narrative. A storyboard will contain words related to the users' chosen story themes. All players are expected to incorporate one word from the storyboard when it is their turn, while the robot will try to make a narrative that leads the next player to choose the predicted word. Participant happiness, narrative quality, and team cohesiveness help evaluate the effectiveness of the approach. Initial results show that this structure leads to significant improvements compared to a baseline system that does not incorporate human preferences. Decision-making on the level of a robot is modeled using the Markov decision process, which is probabilistic; thus, results remain unpredictable but controllable. The Markov decision process is based on the concept that no prior state influences the forthcoming action; only the current state determines what to do next (Unhelkar, Li & Shah 2020). The robot learns how the user has historically interacted to find out what the user wants, which is the "current state" of the robot. This allows the robot to change its behavior in real time to make it easier to work with people.

STS Project Proposal

This research examines the case of Pepper, a humanoid robot that aims to help in various environments like retail, customer service, or healthcare. It was introduced by SoftBank back in 2014 (Calcetero, 2019). Pepper was programmed to provide companionship in hospitals and at senior care facilities, often reminding people about their medication or chats. It all had a good start in healthcare until this went wrong, too. The robot was criticized for being unable to perform simple tasks and repeatedly providing the same programmed responses whenever it interacted with patients or their guardians (Stommel et al., 2022). SoftBank claimed Pepper could understand emotions but failed to live up to these expectations in medical situations. In this research, I will focus on how the interactions between human and non-human actors in Pepper's healthcare network explain its failure to meet user needs. This question shows how the sociotechnical relationships among actors—patients, caregivers, hospital managers, algorithms, and the robot—shaped the outcome of Pepper's integration into healthcare.

Developing meaningful connections with patients, particularly the elderly, who often suffer from loneliness and isolation, was a primary goal of implementing Pepper in healthcare facilities. It is a microcosm for the challenges of robotics in sensitive areas — such as healthcare, where empathy and understanding are key. The research from Stommel et al. involved 36 video recordings of Pepper robots interacting with older adults. The result appears that miscommunication occurs in most of the data. Rest homes turned Pepper down due to an inability to keep its emotions under control, or it was unable to take directions from patients. Faith in Pepper's skills was swiftly eroded as patients and caregivers saw that the robot's replies were repetitious, often irrelevant, and lacked genuine empathy (de Graaf, Ben Allouch, & van Dijk, 2017). Unlike human caretakers, who can adapt their replies to a patient's emotional state and personal history, Pepper's exchanges were almost all written and incapable of meeting its users' complex demands. Previous writers have examined the Pepper robot agent's lack of fluent communication with older people. It has functional deficiencies, such as not recognizing human input in the case. Still, they have not yet adequately addressed whether the robot agent can incorporate human agent preference into account for better communication and support (de Graaf, Ben Allouch, & van Dijk, 2017).

This research offers a deeper understanding of the technology-society relationship illustrated by Pepper's deployment. By analyzing Pepper's failure through the Actor-Network Theory (ANT) lens, this study will reveal how human and non-human actors within the healthcare network collectively influenced the outcome. This method can show how sociotechnical factors affect people's decisions to accept or reject humanoid robots. This analysis provides insights for understanding the design of technologies that align with cultural expectations. With this point of view, stakeholders might get the relationship between technical and social factors right, which could prevent similar operations from failing repeatedly. I can examine Pepper's problem in health care through the lens of actor-network theory (ANT), developed by philosophers such as Bruno Latour and Michel Callon. By applying Actor-Network Theory to the case of Pepper in healthcare, this research argues that the robot's failure was not solely due to its technical limitations but also due to weak network-building among human and non-human actors. This analysis can demonstrate that trust, cultural context, and emotional expectations are key factors of sociotechnical systems and must be integrated into future robotic design.

Actor-network theory (ANT) provides the theoretical basis for this investigation. ANT delves into problem-solving and goal attainment by constructing diverse networks, including human and non-human agents. It highlights the importance of human and non-human actors in determining results, including patients, caregivers, and hospital management, as well as algorithms, institutional rules, and robots. In the case of Pepper's use in healthcare, ANT helps to find the social and technological factors that led to the robot failing to meet user needs. ANT will examine how things in the network relate to and affect each other (Latour, 2005). It looks at how the strength or weakness of these links affected trust, emotional involvement, and how useful people thought the robot was. This theory helps us understand why Pepper didn't catch on in sensitive areas like healthcare by looking at the work of network builders and how they tried to bring together the interests of different players.

This study will use first-hand and second-hand sources to answer the research question. It will look at 36 videos of Pepper robots talking with older adults. These videos were shared by Stommel et al. The goal is to find patterns where communication breaks down. The study will also check how these problems affect trust and emotional connection. Detailed case studies of Pepper's deployment in healthcare facilities will provide contextual insights into how the robot was implemented and received by stakeholders. Additionally, reports and interviews with hospital managers, caregivers, and patients will be reviewed, where available, to understand their expectations, experiences, and feedback regarding Pepper's performance. This study will also use books and articles about Actor-Network Theory (ANT) and its application to healthcare tools. These will help build the main ideas and guide studying how people and technology interact. Technical details about Pepper's design and features will show how its limits affected what users expected. These sources will work together to explain how humans and robots interact in Pepper's healthcare system. They will clearly show the social and technical factors influencing the results.

Conclusion

This work aims to improve human-robot collaboration by enabling an adaptable robotic decision-making process that incorporates preferences from humans in real time. In the technical project, we propose a dynamic framework based on Markov decision processes and large language models for anticipating and responding to human behavior in joint tasks. This approach reduces the action gap between humans and robots while improving team performance, narrative consistency, and user satisfaction. The technical research provides an applied solution to human-robot interaction that can be employed across healthcare, education, and entertainment sectors where adaptive interaction is essential.

The STS study, framed through Actor-Network Theory (ANT), investigates the socio-technical processes influencing the acceptance or rejection of humanoid robots, using Pepper as a case study. This research provides insights into the complex relationships between human and non-human actors in healthcare networks, highlighting trust, cultural expectations, and emotional engagement as critical socio-technical factors. These findings deepen our theoretical understanding of how robots succeed or fail in sensitive environments and offer actionable insights for improving robot design.

Collectively, these projects tackle the technical and societal aspects of human-robot collaboration. The technical research presents tools and frameworks to make robots more adaptive. In contrast, STS research informs the design of more socially aware technologies by identifying the cultural and emotional conditions necessary for success. For instance, the STS findings underscore the importance of integrating trust and emotional engagement into robotic systems. This can guide the technical project in prioritizing features that enhance acceptance and usability. By incorporating these domains, this work advances human-robot interaction as both a technological capability and a socially embedded phenomenon.

Reference

Calcetero, M. A. (2019, June 19). A mass-produced sociable humanoid robot: Pepper: the first machine of its kind. https://www.robotlab.com/research-papers/a-mass-produced-sociable-humanoid-robot-pepper-the-first-machine-of-its-kind?srsltid=AfmBOorEDJwK8v0uFXS0xQpyVaZujxd_UPG7GSGXGxbO0G1ENvsDxHKp

D. Dragan, A., Hadfield-Menell, D., Swamy, G., & Choudhury, R. (2019). On the utility of model learning in HRI. *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 317–325. https://ieeexplore.ieee.org/document/8673256/

De Graaf, M. M. A., Ben Allouch, S., & van Dijk, J. A. G. M. (2017). Why would I use this in my home? A model of domestic social robot acceptance. *Human–Computer Interaction, 34*(2)*, 115–173*. https://doi.org/10.1080/07370024.2017.1312406

Gao, Y., Liu, W., Huang, Y., Yang, W., Fu, Q., Dai, J., Chen, W., Wang, X., Li, S., Yang, W., Wang, W., Lian, Z., Zheng, D., Wang, L., & Liu, F. (2024, January). Enhancing Human Experience in Human-Agent Collaboration: A Human-Centered Modeling Approach Based on Positive Human Gain. Retrieved from https://arxiv.org/html/2401.16444v1

Gillet, S., Vázquez, M., Andrist, S., Leite, I., & Sebo, S. (2024). Interaction-shaping robotics: Robots that influence interactions between other agents. *ACM Transactions on Human-Robot Interaction*, *13*(1)*, 1–23.* https://doi.org/10.1145/3643803

Latour, B. (2005). *Reassembling the social: An introduction to actor-network-theory*. Oxford University Press.

Leite, I., McCoy, M., Lohani, M., Ullman, D., Salomons, N., Stokes, C., Rivers, S., &

Scassellati, B. (2015, March 2). Emotional storytelling in the classroom. *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. *75–82*. https://doi.org/10.1145/2696454.2696481

Saxena, A., Jain, A., & Koppula, H. S. (2016, January). Anticipatory planning for human-robot teams. *In Experimental robotics 453–470*. https://www.researchgate.net/publication/286292497_Anticipatory_Planning_for_Human-Robot_Teams

Shi, L., Bürkner, P.-C., & Bulling, A. (2024, March 25). Inferring human intentions from predicted action probabilities. Retrieved from https://arxiv.org/abs/2308.12194

Shi, L. X., Hu, Z., Zhao, T. Z., Sharma, A., Pertsch, K., Luo, J., Levine, S., & Finn, C. (2024, March 19). Yell at your robot: Improving on-the-fly from language corrections. Retrieved from https://arxiv.org/abs/2403.12910

Stommel, W., de Rijk, L., & Boumans, R. (2022). "Pepper, what do you mean?" miscommunication and repair in robot-led survey interaction. *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), 385–392*. https://ieeexplore.ieee.org/document/9900528/

Unhelkar, V. V., Li, S., & Shah, J. A. (2020, March 9). Decision-making for bidirectional communication in sequential human-robot collaborative tasks. *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*. *329–341* https://doi.org/10.1145/3319502.3374779

Yang, D., Kleinman, E., Troiano, G. M., Tochilnikova, E., & Harteveld, C. (2024, April 11). Snake story: Exploring game mechanics for mixed-initiative co-creative storytelling games. Retrieved from https://arxiv.org/abs/2404.07901