Algorithmic Robot Teachers for Healthcare Skill Training

Peizhu Qian
Rice University, Houston, Texas USA
pqian@rice.edu

I. Motivation

The need to train new workers effectively and upskill the existing workforce is a challenge faced by almost every industry across the globe. The healthcare industry, in particular, is confronting a crisis. The World Health Organization (WHO) projects a shortage of 10 million healthcare workers by 2030 [1]. Although no country is exempt from this growing problem, the greatest gaps are found in countries in Africa, Southeast Asia, the Mediterranean Region, and parts of Latin America [2]. This problem is further compounded by workers leaving their home countries to pursue opportunities elsewhere. A shortage of experienced healthcare workers and faculty to teach is a limiting factor that leads to enrollment limitations [3]. To protect the health of the world's population, we must investigate transformative solutions to achieve efficient, resilient, and sustainable local and global healthcare systems.

Recognizing this pressing need, I create robotic teachers that assist human learners in the acquisition of new skills, with applications primarily in healthcare.

II. Current and Past Work

My work has demonstrated the richness of interesting research questions that arise from the area of skill acquisition as well as the intellectual, technical, and societal contributions.

A. Cognitive Modeling for Personalized Explanations

Research Gap: Work in explainable AI (XAI) [4, 5, 6, 7] primarily uses machine learning measures such as the Q-value or Shapley value to generate explanations, which overlook human factors such as mental modeling. However, creating a computational model of humans that accurately captures their cognitive processes is extremely challenging [8]. Methods like Inverse Reinforcement Learning (IRL) require a large number of user inputs, which is not feasible in real-world applications [9, 10, 11, 12].

Contribution: My work, AI Teacher [13] focused on generating explanations for human users of autonomous robots. It was the first to incorporate interactivity in XAI and to address individual users' confusion. AI Teacher creates a Bayesian-Theory-of-Mind probabilistic model of the human user then algorithmically generates explanations for the user by estimating their learning curve. AI Teacher uses an interactive user interface that enables users to ask specific questions regarding a robot's behavior. To further capture individual user's experience and prior knowledge, Personalized Policy Summarization (PPS) [14] proposed a Dirichlet-based model that can quickly learn the user's mental model through a few questions. PPS provided the first solution to algorithmically evaluate a user's knowledge of a robot policy online and computationally generate personalized explanations. This is achieved through rigorous mathematical modeling, entropy-based question selection, and robust online explanation generation. Human experiments showed that learners performed 32% better with my methods and found my approach more engaging ($p<0.05$) than competing algorithms.

B. Virtual + Physical Robot Demonstration

Research Gap: XAI work relies on virtual methods to demonstrate explanations,e.g., figures or animations. Virtual systems provide demonstrations quickly but are limited by the fidelity of the simulation systems. Work in robotics uses physical systems but they are resource- and time-intensive [15, 16].

Contribution: I developed a hybrid system that consisted of virtual and physical robots [17]. Users used the virtual robot to gain a quick understanding of the robot's behavior in perfect settings, then switched to the physical robot to learn the robot's behavior in more general, real-world settings. While this work demonstrated the benefits of both virtual and physical systems, the physical robot demonstration still posed challenges in cost and time. Thus, my team created STOREE [18], a virtual system that includes real-world scenarios. Experiment results indicated statistically significant improvement in learner performance ($p<0.001$) using our approach.

C. Intelligent Robot Tutor for Nursing Education

Research Gap: Automated skill evaluation is a long-standing research question in Intelligent Tutoring Systems (ITS) and related fields such as medical robots. The biggest challenge lies in extending methods such as Knowledge Tracing from discrete observation (e.g., math questions) to long-horizon and continuous observations (e.g., surgery) [19, 20].

Contribution: To observe a learner's movements, I used deep-learning-based computer vision methods. I collaborated with Houston Methodist Hospital to develop Astrid [21], a robotic tutoring system for nursing students. Astrid tracks a nursing student's body and hand positions (Fig.1-left) and alerts the student if a mistake has been detected in real-time (Fig.1-middle). After the practice session, Astrid provides the student with a summary of their performance, as well as screenshots of their mistakes to help the nursing student

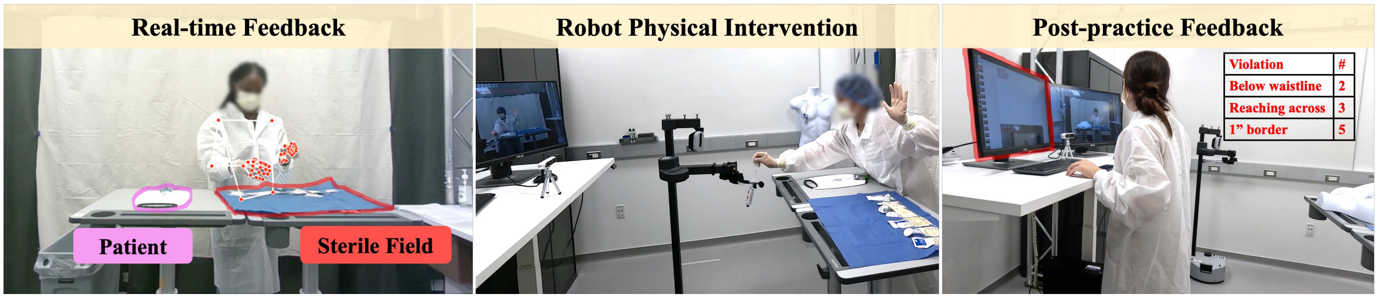

Fig. 1. Nurses practicing with the robotic tutor in a training environment. The tutor offers real-time guidance, interventions, and post-practice feedback.

quickly review the practice (Fig.1-right). My system also creates realistic scenarios through physical robot intervention to help nursing students prepare for real-life practices.

ASTRID is the first robotic tutor for nursing education. To understand the nursing discipline, I attended nursing classes, observed on hospital floors, and worked closely with diverse stakeholders – nursing faculty, students, nurses, nursing scientists, and hospital executives. My effort was commended by both computer science and nursing communities, and led to successful cross-disciplinary and cross-institution collaboration.

III. Future Directions

I am interested in developing human-centered robotic systems that assist workers in acquiring new skills. I plan to achieve this goal by investigating the three research threads outlined in this section, ranging from algorithmic advances to real-world deployment.

A. Extend XAI to Diverse Task Models

Research Gap: In existing XAI literature, tasks are most commonly represented using Markovian models or neural networks. However, many human tasks, like those in healthcare, may not be best represented using these models. How to extend existing XAI methods to explain diverse task models is an open and interesting research question.

Proposal: First, we need to understand how humans represent tasks. In healthcare, step-by-step checklists are frequently used. However, converting checklists into computational models is not trivial because steps in checklists include complex environment observations, states of medical tools, and human-object interaction. Some steps are composed of multiple actions (e.g., use sterile techniques to open the dressing kit) – sterile techniques are not specified and may vary based on the tasks and environment. To explain tasks that involve complicated environmental observations and human-object interactions, I will first extend existing task models in robotics, such as Planning Domain Definition Language (PDDL) [22, 23, 24] and Hierarchical Task Networks (HTN) [25], as useful starting points. In the long run, I will investigate novel techniques that capture the diversity in healthcare tasks and hospital-specific practices. I will explore the use of Large Language Models (LLM) to enable healthcare experts to translate their domain knowledge into robot-interpretable computational models, and then generate high-quality explanations. Leveraging the benefits of personalized learning, these explanations will be tailored to individual learners.

B. Automate Skill Evaluation

Research Gap: Existing Knowledge Tracing methods, traditionally used for discrete observations such as multiple choices, struggle with continuous observations. Unlike discrete domains where the observation-to-skill mapping is relatively straightforward, evidence of skill mastery in continuous, real-world domains is implicit and largely depends on the context.

Proposal: To tackle this fundamental question, it is important to look at how human experts evaluate skills in such domains. For rule-based skills, such as sterile techniques, we can track the number of times a student breaks each rule using multimodal data – vision, language, and interaction. For knowledge-based skills that depend on context, such as treatment for dropping blood sugar levels, I propose to use physical robot intervention to create realistic scenarios for students to practice. Simulating realistic scenarios is a common practice that experienced nurses use to train new nurses but it varies from nurse to nurse and from hospital to hospital. A standardized simulation can be realized through robot teachers. I also propose to extend existing methods to create new Knowledge Tracing methods for broad real-world domains.

C. Real-world Implementation and Societal Impact

Challenges: Implementing robots in hospitals poses significant challenges. On the technical side, hospital environments are extremely complex. On the non-technical side, patient privacy, data security, and legal concerns are important considerations.

Proposal: To start, I will look for local collaborators who oversee education programs in health-related fields, such as nursing schools. I will build robot teachers to be used in simulation labs and classrooms for training purposes. Over time I will expand my systems to be used for orientation, internship, quarterly skill checkoffs, and retraining. At large hospitals, thousands of nurses go through training and evaluation each year. This process requires a large amount of human effort and financial resources. I will collaborate with hospitals to automate the training and evaluation processes. **As the leader in this research area, I will work with international collaborators and lead a global movement to produce a sustainable healthcare workforce.**

References

[1] WHO, "Health workforce."

[2] M. Boniol, T. Kunjumen, T. S. Nair, A. Siyam, J. Campbell, and K. Diallo, "The global health workforce stock and distribution in 2020 and 2030," *BMJ Global Health*, 2022.

[3] M. Marć, A. Bartosiewicz, J. Burzyńska, Z. Chmiel, and P. Januszewicz, "A nursing shortage–a prospect of global and local policies," *International nursing review*, vol. 66, no. 1, pp. 9–16, 2019.

[4] Y. Zhan, A. Fachantidis, I. Vlahavas, and M. E. Taylor, "Agents teaching humans in reinforcement learning tasks," in *Proceedings of International Conference on Autonomous Agents and Multiagent Systems*, 2014.

[5] D. Amir and O. Amir, "Highlights: Summarizing agent behavior to people," in *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 2018.

[6] S. H. Huang, D. Held, P. Abbeel, and A. D. Dragan, "Enabling robots to communicate their objectives," *Autonomous Robots*, vol. 43, February 2019.

[7] M. S. Lee, H. Admoni, and R. Simmons, "Machine teaching for human inverse reinforcement learning," *Frontiers in Robotics and AI*, vol. 8, p. 188, 2021.

[8] Y. Rong, T. Leemann, T.-T. Nguyen, L. Fiedler, P. Qian, V. Unhelkar, T. Seidel, G. Kasneci, and E. Kasneci, "Towards human-centered explainable ai: A survey of user studies for model explanations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 4, pp. 2104–2122, 2024.

[9] A. Y. Ng and S. J. Russell, "Algorithms for inverse reinforcement learning," in *Proceedings of the Seventeenth International Conference on Machine Learning (ICML)*, pp. 663–670, 2000.

[10] D. Ramachandran and E. Amir, "Bayesian inverse reinforcement learning," in *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2586–2591, 2007.

[11] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *AAAI Conference on Artificial Intelligence*, pp. 1433–1438, 2008.

[12] D. Hadfield-Menell, S. J. Russell, P. Abbeel, and A. Dragan, "Cooperative inverse reinforcement learning," in *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3909–3917, 2016.

[13] P. Qian and V. Unhelkar, "Evaluating the role of interactivity on improving transparency in autonomous agents," in *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '22, p. 1083–1091, International Foundation for Autonomous Agents and Multiagent Systems, 2022.

[14] P. Qian, H. Huang, and V. Unhelkar, "Pps: Personalized policy summarization for explaining sequential behavior of autonomous agents," *Proceedings of the 2024 AAAI/ACM Conference on AI, Ethics, and Society*, 2024.

[15] S. Wallkötter, S. Tulli, G. Castellano, A. Paiva, and M. Chetouani, "Explainable embodied agents through social cues: a review," *ACM Transactions on Human-Robot Interaction (THRI)*, 2021.

[16] T. Sakai and T. Nagai, "Explainable autonomous robots: a survey and perspective," *Advanced Robotics*, 2022.

[17] P. Qian and V. V. Unhelkar, "Interactively explaining robot policies to humans in integrated virtual and physical training environments," in *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 847–851, 2024.

[18] H. Huang, P. Qian, A. P. B. Campo, and V. Unhelkar, "Sim-to-real gaps in explainable ai for robot policy summarization," *Preprint*, 2025.

[19] C. Quintero-Pena, P. Qian, N. M. Fontenot, H.-M. Chen, S. K. Hamlin, L. E. Kavraki, and V. Unhelkar, "Robotic tutors for nurse training: Opportunities for hri researchers," in *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pp. 220–225, IEEE, 2023.

[20] P. Qian, Z. Cai, and X. Hu, "Life in the semantic space: Structures of the language network," *Society for Computers in Psychology 48th Annual Meeting (SCiP)*, 2018.

[21] P. Qian, F. Bajraktari, C. Quintero-Peña, Q. Meng, S. Hamlin, L. E. Kavraki, and V. Unhelkar, "Astrid: A robotic tutor for nurse training to reduce healthcare-associated infections," *Proceedings of Robotics: Science and Systems (RSS)*, 2025.

[22] D. McDermott, M. Ghallab, A. Howe, C. Knoblock, A. Ram, M. Veloso, D. Weld, and D. Wilkins, "PDDL — the planning domain definition language," tech. rep., Yale Center for Computational Vision and Control, 1998.

[23] M. Abuazizeh, T. Kirste, and K. Yordanova, "Computational state space model for intelligent tutoring of students in nursing subjects," in *Proceedings of the 13th ACM International Conference on PErvasive Technologies Related to Assistive Environments*, pp. 1–7, 2020.

[24] M. Abuazizeh, K. Yordanova, and T. Kirste, "Affect-aware conversational agent for intelligent tutoring of students in nursing subjects," in *Intelligent Tutoring Systems* (A. I. Cristea and C. Troussas, eds.), (Cham), pp. 497–502, Springer International Publishing, 2021.

[25] K. Erol, J. Hendler, and D. S. Nau, "Htn planning: Complexity and expressivity," in *AAAI*, vol. 94, pp. 1123–1128, 1994.