

Large Language Models Enable Automated Formative Feedback in Human-Robot Interaction Tasks

Emily Jensen
University of Colorado Boulder
Boulder, Colorado, USA
emily.jensen@colorado.edu

Sriram Sankaranarayanan
University of Colorado Boulder
Boulder, Colorado, USA
srirams@colorado.edu

Bradley Hayes
University of Colorado Boulder
Boulder, Colorado, USA
bradley.hayes@colorado.edu




Figure 1: Example pipeline for generating automated feedback for HRI tasks

1 INTRODUCTION

Rapid advances in robotic systems will significantly impact the future of work. Workers will need to be able to quickly adapt to using new systems as new capabilities and application domains emerge. A recent report estimates that one third of job requirements will require technological skills that are not yet considered crucial [9]. It is also estimated that 50% of existing employees will need to be retrained or upskilled by 2025 to keep up with technological advancement, placing significant pressure on both employers and employees to meet these demands [7, 9]. As such, we need scalable methods to train users for human-robot interaction (HRI) tasks.

Automated feedback is a promising approach to scale up training for HRI tasks. By pairing domain knowledge representations with effective assessment, automated feedback systems can identify a learner’s current strengths and weaknesses and suggest future actions that will help the learner master the target task. In this paper, we discuss how large language models (LLMs) can be used as a tool for providing automated feedback for learning HRI tasks alongside illustrative examples.

2 POSITION STATEMENT

Representing knowledge and assessing someone’s ability in an HRI task is difficult, due to complex objectives and high variability in human performance. In [5], we begin to address this question by breaking down HRI tasks into objective primitives that can be combined sequentially and concurrently (e.g., maintain slow speed and reach waypoints). We then show that signal temporal logic specifications, paired with a robustness metric, are a useful tool for assessing performance along each primitive. These formal methods allow designers to precisely represent ideal trajectories.

This formulation admits explainability, as one can identify and elaborate upon specific objectives that learners did not accomplish.

We claim that *LLMs can be paired with formal analysis methods to provide accessible, relevant feedback for HRI tasks*. While logic specifications are useful for defining and assessing a task, these representations are not easily interpreted by non-experts. Luckily, LLMs are adept at generating easy-to-understand text that explains difficult concepts. By integrating task assessment outcomes and other contextual information into an LLM prompt, we can effectively synthesize a useful set of recommendations for the learner to improve their performance (see Figure 1 as one example).

3 CHALLENGES AND OPPORTUNITIES

To the best of our knowledge, there has been no systematic work addressing how to train humans to effectively work with robotic systems. If we consider training motor skills more broadly, many approaches focus on augmenting sensory input to provide control-level feedback during a task [10]. Some examples include visualizing the predicted future trajectory of a drone [13] or generating a haptic response to bias the operator to an ideal course [1, 8]. These approaches to feedback fall short in several respects: (1) just-in-time feedback does not promote intentional reflection from the learner on how to improve their performance, (2) feedback is not integrated with an established training curriculum, and (3) feedback does not adapt to a learner’s learning trajectory over time.

LLMs are a promising technology poised to tackle these challenges in generating feedback. To address challenge (1), we can develop feedback templates that include elements of effective formative feedback. Formative feedback is an established approach in education that focuses on motivating learners, having them reflect on their performance, and providing a manageable amount of feedback they can use on their next attempt. For example, we generated the following feedback for the safe trajectory shown in Figure 2:




Figure 2: Example of three trajectories in a 2D drone landing task. The pilot must keep the drone within the black boundary and safely reach the gray landing pad with a velocity ≤ 5 m/s and attitude $\leq 10^\circ$.

Great job maintaining the safety boundaries throughout your flight! While your landing was accurate with respect to y-axis, it seems there’s room for improvement in maintaining the x-axis positioning within the landing pad. To refine this, consider adjusting the roll input slightly when approaching the landing pad. Reflect on the trajectory taken during the descent and think about how you could make it more centered over the pad. I’m confident that with a few tweaks, you’ll nail the perfect landing next time!

LLMs can also be integrated with broader training systems that document domain knowledge and skill structures, addressing challenge (2). Finally, LLMs can access historical records of the learner’s performance to understand repeated mistakes and opportunities for growth. The greatest strength of LLMs is generating friendly, approachable text, making it ideal for providing feedback that learners perceive positively.

4 POTENTIAL IMPACT

LLM feedback is easy to iterate on and integrate into existing technical workflows. For example, if a task involves multiple robots, we can quickly modify a template prompt for feedback by including a short description of each robot’s dynamics or the specific part of the task it is used for. If we need a new feedback format to test a new learning theory, we can swap out the part of the prompt that tells the LLM how to frame its response. Changing the feedback presentation is also straightforward; the LLM can act as an intermediate interface between the task and feedback, generating appropriate low-level code to display on a virtual reality headset or to be spoken by a social robot.

The social implications of the proposed approach are also worth considering. As more people are required to interact with robots, training needs to be scalable and personalized to each learner. LLMs can help us reach this goal by making training a friendlier and more appealing process. When paired with commercially available products like virtual reality headsets, training for robotic systems can be accessible to a more diverse group of learners and enable

them to train for technical jobs. Additionally, automated feedback generation lessens the burden on human instructors, who will not need to provide as much direct oversight during training.

5 CHALLENGES

As with any new technology that is not fully understood, there are many questions to consider before integrating LLM feedback into high-stakes or high-impact HRI domains.

The first consideration is anticipating and handling unexpected outputs from the feedback system. It is widely known that LLMs can produce incorrect and harmful responses due to the stochastic nature of the model, biases in the training data, and nuances in system prompts. An automated feedback system should have internal processes to moderate potential harmful outputs, which can be built into the prompt. For example, Tree-of-Thought prompting [4] can be used to emulate experts giving multiple feedback variations and having them reach a consensus based on internal knowledge (recent performance, historical errors, possible emotional states). This approach allows the LLM to recognize and discard inaccurate or poorly phrased feedback. Additional research can collaborate with natural language processing efforts to develop safety alignment when training new models [2].

The second consideration is implementing the feedback system in a robust and sustainable manner. A feedback system requires a thorough yet flexible knowledge representation of the target domain. Using principles from participatory design [11], system creators and domain experts can work together to identify key learning outcomes and assessment criteria. These core learning concepts can then be codified in a formal framework such the Knowledge-Learning-Instruction Framework [6], which associates each practice item with one or more knowledge components. Finally, feedback systems should operationalize theory-driven intervention strategies for providing feedback, such as the Zone of Proximal Development [12] or Deliberate Practice [3]. This would bring more rigor to HRI studies while also contributing to interdisciplinary discourse on training for real-world domains.

6 CALL TO ACTION

In this paper we have presented concepts suggesting that LLMs are a promising tool for automatically generating feedback for learning HRI tasks. By using theory-driven approaches to assessing learner performance and providing feedback, the HRI community can develop robust systems that both contribute to a larger discussion of learning and training and fill critical gaps in domains with immediately relevant applications and societal benefits.

As this area develops, we propose a few starting research directions. First, we must understand how learning for HRI tasks (which commonly involve physical manipulation or inter-agent communication) differs from the traditional classroom context and how this impacts theories of learning. Second, we can investigate how learner agency fits into an automated feedback system; how much control do learners want over the tone, content, and delivery of their feedback? Do these choices make the feedback less effective? Finally, we need to consider the temporal aspect of learning. How can feedback address both the most recent learning attempt and glean long-term trends from historical data? These directions can prompt more nuanced questions as the research develops.

REFERENCES

- [1] Sunil K. Agrawal, Xi Chen, Christina Ragonese, and James C. Galloway. 2012. Training Toddlers Seated on Mobile Robots to Steer Using Force-Feedback Joystick. *IEEE Transactions on Haptics* 5, 4 (2012), 376–383. <https://doi.org/10.1109/TOH.2011.67>
- [2] Zhichen Dong, Zhanhui Zhou, Chao Yang, Jing Shao, and Yu Qiao. 2024. Attacks, Defenses and Evaluations for LLM Conversation Safety: A Survey. arXiv:2402.09283 [cs.CL]
- [3] K. Anders Ericsson, Ralf Th. Krampe, and Clemens Tesch-Römer. 1993. The Role of Deliberate Practice in the Acquisition of Expert Performance. *Psychological Review* 100, 3 (1993), 363–406. <https://doi.org/10.1037/0033-295x.100.3.363>
- [4] Dave Hulbert. 2023. Using Tree-of-Thought Prompting to boost ChatGPT's reasoning. <https://github.com/dave1010/tree-of-thought-prompting>. <https://doi.org/10.5281/ZENODO.10323452>
- [5] Emily Jensen, Bradley Hayes, and Sriram Sankaranarayanan. 2023. More Than a Number: A Multi-dimensional Framework For Automatically Assessing Human Teleoperation Skill. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, Stockholm Sweden, 653–657. <https://doi.org/10.1145/3568294.3580167>
- [6] Kenneth R. Koedinger, Albert T. Corbett, and Charles Perfetti. 2012. The Knowledge-Learning-Instruction Framework: Bridging the Science-Practice Chasm to Enhance Robust Student Learning. *Cognitive Science* 36, 5 (July 2012), 757–798. <https://doi.org/10.1111/j.1551-6709.2012.01245.x>
- [7] Ling Li. 2022. Reskilling and Upskilling the Future-ready Workforce for Industry 4.0 and Beyond. *Inf Syst Front* (July 2022). <https://doi.org/10.1007/s10796-022-10308-y>
- [8] Carine Rognon, Vivek Ramachandran, Amy R Wu, Auke J Ijspeert, and Dario Floreano. 2019. Haptic Feedback Perception and Learning With Cable-Driven Guidance in Exosuit Teleoperation of a Simulated Drone. *IEEE Transactions on Haptics* 12, 3 (2019), 375–385. <https://doi.org/10.1109/TOH.2019.2925612>
- [9] Klaus Schwab and Saadia Zahidi. 2020. *The Future of Jobs Report 2020*. Technical Report. World Economic Forum. https://www3.weforum.org/docs/WEF_Future_of_Jobs_2020.pdf
- [10] Roland Sigrüst, Georg Rauter, Robert Riener, and Peter Wolf. 2013. Augmented visual, auditory, haptic, and multimodal feedback in motor learning: A review. *Psychon Bull Rev* 20, 1 (Feb. 2013), 21–53. <https://doi.org/10.3758/s13423-012-0333-8>
- [11] Clay Spinuzzi. 2005. The methodology of participatory design. *Technical communication* 52, 2 (2005), 163–174.
- [12] L. S. Vygotsky. 1978. *Mind in Society: The Development of Higher Psychological Processes*. Harvard University Press, Cambridge, MA. <https://doi.org/10.2307/1421493> ISSN: 00029556.
- [13] Markus Wilde, May Chan, and Brian Kish. 2020. Predictive Human-Machine Interface for Teleoperation of Air and Space Vehicles over Time Delay. In *2020 IEEE Aerospace Conference*. 1–14. <https://doi.org/10.1109/AERO47225.2020.9172297>