Longitudinal Robot Learning from Demonstration with Care Providers in a Home Environment

Nina Moorman*, Julianna Schalkwyk, Vriksha Srihari, Qingyu Xiao, Kamel Alrashedy,
Hongseok Jeong, Kiersten Lange, Matthew B. Luebbers, Matthew Gombolay

***Abstract*—Learning from demonstration (LfD) methods enable non-expert end users to teach robots novel skills without explicit programming. However most evaluations of the usability of LfD with non-experts has been conducted in controlled laboratory environments with a robotics experimenter present. In this work we identify non-expert end users' key barriers when teaching robots via demonstration without live robotics expert feedback in a home environment. In our human subjects experiment we support the non-expert end users through two forms of demonstrator guidance developed in prior work: pre-training and adaptive feedback. Towards the ecological validity of the evaluation, we conduct this experimentation over multiple visits, with a population of care providers. Finally, we propose to open source the resulting LfD dataset of care providers teaching a robot assistive tasks over multiple visits to a home environment.**

***Index Terms*—learning from demonstration, translational research, assistive robotics**

I. INTRODUCTION

Household robots will need to interact with and adapt to unstructured, dynamic human environments, presenting a challenge for robot policies trained in controlled, laboratory-like settings [1]. While increasing the scale of training data could help address this challenge, collecting such real-world data is expensive, time consuming, and risks exposing private information. Alternatively, in-home supplemental learning from end users would enable personalization of robot behavior to those users' needs and environments, without requiring the same prohibitive scale of outside training data. Learning from Demonstration (LfD) methods enable robots to learn skills from user demonstrations. Various prior works have developed instructional materials [2], [3] and feedback mechanisms [4], [5] to enable non-experts to teach robots via LfD.

This work seeks to ascertain the remaining challenges non-expert users will face when teaching robots via LfD. As prior work has shown that demonstrators improve at providing kinesthetic demonstrations over the course of multiple sessions [6], we propose a multi-visit, multi-task evaluation. To improve the ecological validity of our evaluation, we conduct our study in a realistic home environment where non-expert end users will define and demonstrate their own custom unimanual and bimanual manipulation tasks. We will recruit participants with no prior robotics or computer science experience and, over the course of three visits, provide them with automated training and feedback, to determine, with such support, what challenges they face when teaching a robot assistive household tasks.

In this work we propose to contribute the following:

1) We design a human subjects experiment to characterize current challenges and opportunities for deploying LfD systems with non-expert end users provided with pre-training and adaptive feedback, in a home environment.
2) We will open source the subsequent multi-visit dataset of non-experts teaching robots a series of assistive tasks in a home environment.

II. DEMONSTRATION SYSTEM

To detect task-relevant objects and produce segmentation masks, we process an RGB-D stream from an Intel RealSense camera with YOLOE [7]. The corresponding masked depth measurements are back-projected into object-specific 3D point clouds. To estimate object pose, point clouds are registered against predefined object shapes using Iterative Closest Point (ICP) [8], producing a 6D pose estimate for tracked objects.

For robot skill learning, we use Cartesian Probabilistic Motion Primitives (ProMPs) [9] over end effector and tracked object poses. We choose to employ ProMPs as they can learn from limited demonstrations, can be trained quickly, and are robust to changing object poses allowing for generalization to different environment setups, compared to Behavior Cloning approaches which often struggle to generalize [10] or Vision-Language Action models that may struggle to adapt to people's specific preferences without time consuming finetuning.

We develop an interactive interface for providing kinesthetic demonstrations. The interface allows participants to plan a set of low-level skills to teach to accomplish tasks in a given domain (called a task decomposition), provide demonstrations for each skill, and create recipes of skills to accomplish tasks. The interface additionally supports the participant through pre-training and through adaptive feedback.

Pre-Training (PT)

We provide participants with PT developed in prior work [2] where, for a set of pre-defined domains, the participant first attempts to break down a task in that domain, then is shown videos of an optimal task decomposition and how those skills are taught by a robotics expert via video recordings. PT allows the participant to improve upon their task decompositions *on novel domains* without the need for live feedback from an expert demonstrator [2].

---

*Corresponding author: nmoorman3@gatech.edu
Authors affiliated with Georgia Institute of Technology, Atlanta, GA, USA.
This research was supported by NSF grants IIS-2340177 and IIS-2112633.

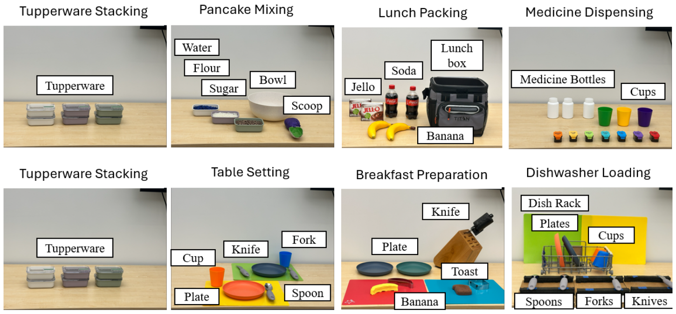

Fig. 1. Domains for visit 1 (first row) and visit 2 (second row). Tupperware stacking serves as a kinesthetic on-boarding task for both visits. These domains were chosen to represent a range of activities of daily living. Additionally, we balance task complexity and distribution of unimanual to bimanual tasks between visits.

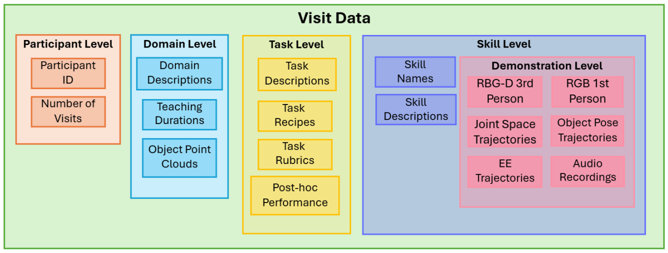

Fig. 2. This figure depicts the dataset structure where each datapoint represents one visit, with information about the participant, domains, tasks, and demonstrated skills.

Adaptive Feedback (AF)

Foundation Model (FM) Feedback: FM feedback provides guidance on whether a participant's planned task decomposition can adequately generalize to downstream tasks [4]. As input, the FM feedback module takes in descriptions of the domain, environment, and known tasks. To support novice users in defining and describing novel domains and tasks, we develop an additional FM feedback module, beyond that of prior work, which iteratively asks clarification questions until the domain description has the necessary amount of detail for the aforementioned task decomposition feedback to be helpful.

Real Robot Replay (RRR) Feedback: The robot rolls out the learned policy, allowing the user to directly observe its performance in-situ [4].

Augmented Reality (AR) Feedback: The robot rolls out the learned policy in AR to serve as a visual debugging tool [11], allowing the user to directly observe policy performance and generalizability without time consuming or dangerous environmental interactions [4].

III. DATA COLLECTION & HUMAN EVALUATION

We conduct our evaluation in the Georgia Tech Aware Home [12], an authentic home environment. We recruit from a population of local care providers of older adults, both formal (doctors, nurses, physical and occupational therapists), and informal [13] (family, friends, or neighbors who interface with the older adult on at least a monthly basis).

Procedure: In this work, participants teach tasks to a pair of JACO 2 arms over the course of three visits. In the first visit, participants experience the PT [2] where they acquire experience teaching the robot a series of tasks via LfD in various domains (depicted in the first row of Figure 1). In the second visit, participants teach the robot additional tasks in various domains (depicted in the second row of Figure 1) with all three forms of AF available (FM, AR, and RRR feedback). In the third and final visit, participants are tasked to define their own custom domains and tasks, with AF available.

Conditions and Research Questions (RQ): Our between-subjects experiment has two conditions. In the **PT+AF** condition, participants obtain PT in visit 1, then experience AF feedback for visit 2 and visit 3. In the **AF** condition, participants experience AF feedback for visit 2 and visit 3, but do not participate in visit 1. Our RQs include:

- **RQ1:** *Are there differences between the PT+AF condition and the AF condition in terms of task performance?* After each visit, we evaluate task completion percentage using a pre-defined rubric.
- **RQ2:** *Are there differences between the PT+AF condition and the AF condition in terms of alignment between predicted and actual robot performance?* Alignment is calculated as the absolute difference between task performance and normalized perceived performance [4].
- **RQ3:** *Are there differences between the PT+AF condition and the AF condition in terms of user experience?* The user experience metrics of interest include usability and acceptance [14] and learned trust [15], and workload [16].

Statistical comparisons between conditions will employ multiple linear regressions with mixed effects, controlling for time-on-task, number of demonstrations, and visit count.

IV. OUTCOMES

The outcomes from this work will be twofold. We will open source the data collected, supplementing existing datasets of demonstrated robot skills such as MIME [17] (which focuses on unimanual, multi-step tasks), or RoboPro [18] (focusing on bimanual, short-horizon tasks) with realistic assistive tasks trained in-situ over multiple visits by care providers. This will benefit the field by enabling standardized benchmarking for interactive robot learning algorithmic development and evaluation. The dataset structure is depicted in Figure 2.

For each of the eleven domains (or eight domains in the AF condition), participants employ their skill library to accomplish six tasks. Natural language descriptions of domains, tasks, and skills will be included in the dataset, along with the point-clouds of all objects in each domain. Demonstration data will include end-effector and joint trajectories, participant audio, and egocentric and exocentric video footage.

Furthermore, in addition to reporting upon the human-factors findings related to our three research questions, we plan to report on lessons learned and propose design guidelines for researchers and practitioners deploying LfD systems in home environments for assistive tasks, designing for care providers, and collecting multi-visit demonstration datasets.

REFERENCES

[1] S. A. Olatunji, Y.-L. Tsai, S. A. Gowrishankar, M. A. Bayles, and W. A. Rogers, "Robots for older adults: A scoping review," *ACM Transactions on Human-Robot Interaction*, vol. 15, no. 3, pp. 1–35, 2026.

[2] N. M. Moorman, N. Gopalan, A. Singh, E. Hedlund-Botti, M. L. Schrum, C. Yang, L. Seelam, and M. Gombolay, "Investigating the impact of experience on a user's ability to perform hierarchical abstraction," in *Robotics: Science and Systems*, 2023.

[3] M. Cakmak and L. Takayama, "Teaching people how to teach robots: The effect of instructional materials and dialog design," in *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*, 2014, pp. 431–438.

[4] N. M. Moorman, M. Luebbers, Z. Xi-Jia, Y. C. Lau, Y. Yao, M. Langwasser, Z. Zaidi, L. Chen, S. van Waveren, and M. Gombolay, "Teaching the teacher: Live foundation model and augmented reality feedback for human-to-robot skill transfer," in *Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction*, 2026, pp. 544–552.

[5] M. Srivastava, N. Goodman, and D. Sadigh, "Generating language corrections for teaching physical control tasks," in *International Conference on Machine Learning*. PMLR, 2023, pp. 32 561–32 574.

[6] P. Aliasghari, M. Ghafurian, C. L. Nehaniv, and K. Dautenhahn, "How non-experts kinesthetically teach a robot over multiple sessions: diversity in teaching styles and effects on performance," *International Journal of Social Robotics*, vol. 16, no. 11, pp. 2079–2105, 2024.

[7] A. Wang, L. Liu, H. Chen, Z. Lin, J. Han, and G. Ding, "Yoloe: Real-time seeing anything," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025, pp. 24 591–24 602.

[8] P. J. Besl and N. D. McKay, "Method for registration of 3-d shapes," in *Sensor fusion IV: control paradigms and data structures*, vol. 1611. Spie, 1992, pp. 586–606.

[9] A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann, "Probabilistic movement primitives," *Advances in neural information processing systems*, vol. 26, 2013.

[10] F. Codevilla, E. Santana, A. M. López, and A. Gaidon, "Exploring the limitations of behavior cloning for autonomous driving," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9329–9338.

[11] M. B. Luebbers, C. Brooks, C. L. Mueller, D. Szafir, and B. Hayes, "Arc-lfd: Using augmented reality for interactive long-term robot skill maintenance via constrained learning from demonstration," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 3794–3800.

[12] C. D. Kidd, R. Orr, G. D. Abowd, C. G. Atkeson, I. A. Essa, B. MacIntyre, E. Mynatt, T. E. Starner, and W. Newstetter, "The aware home: A living laboratory for ubiquitous computing research," in *International workshop on cooperative buildings*. Springer, 1999, pp. 191–198.

[13] M. H. Cantor, "Neighbors and friends: An overlooked resource in the informal support system," *Research on aging*, vol. 1, no. 4, pp. 434–463, 1979.

[14] D. Belanche, L. V. Casaló, and C. Flavián, "Integrating trust and personal values into the technology acceptance model: The case of e-government services adoption," *Cuadernos de Economía y Dirección de la Empresa*, vol. 15, no. 4, pp. 192–204, 2012.

[15] J.-Y. Jian, A. M. Bisantz, and C. G. Drury, "Foundations for an empirically determined scale of trust in automated systems," *International journal of cognitive ergonomics*, vol. 4, no. 1, pp. 53–71, 2000.

[16] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," in *Advances in psychology*. Elsevier, 1988, vol. 52, pp. 139–183.

[17] P. Sharma, L. Mohan, L. Pinto, and A. Gupta, "Multiple interactions made easy (mime): Large scale demonstrations data for imitation," in *Conference on robot learning*. PMLR, 2018, pp. 906–915.

[18] Z. Li, "Robopro: 80-task bimanual manipulation demonstrations on aloha-agilex," 2026. [Online]. Available: https://huggingface.co/datasets/Hoshipu/RoboPro