Personalizing Humanoid Robot Behavior through Incremental Learning from Natural Interactions

Timo Weberruß, Leonard Bärmann, Fabian Peller-Konrad, Alex Waibel, Tamim Asfour

***Abstract*— A key aspect of intuitive and natural Human-Robot-Interaction (HRI) is personalization, i. e., adapting the robot's behavior to align with the current user's needs and preferences. While recent approaches leverage foundation models to enable natural HRI, personalization remains an underexplored challenge. In this extended abstract, we propose a system that incrementally learns personalized humanoid robot behavior from natural-language interactions. Specifically, we build upon our existing dialog system that uses a large language model (LLM) to generate high-level code to steer the robot's behavior, and compare two approaches for personalization: (i) explicit user profiles by storing structured user-specific facts and allowing the LLM to determine when to access or update these memories, and (ii) incremental interactive learning by extending our existing learning mechanism to handle multiple user-specific interaction memories over time. We present preliminary experimental results in simulation and further demonstrate our system on the real-world humanoid robot ARMAR-DE. Our findings highlight that personalization remains an important but complex and non-trivial challenge for future research.**

I. Introduction

Personalization is a key aspect of intuitive Human-Robot-Interaction (HRI). Specifically, a robot should recognize who it is currently interacting with, and adapt its dialog and behavior to the current user's needs and preferences. For instance, consider the simple instruction "prepare some coffee for me". Given a set of skills needed to execute the task such as navigation, object localization, grasping, pouring, etc., the robot needs to plan an action sequence to achieve this goal, e. g., getting a cup, putting it into the coffee machine, pressing a button, and finally serving it. However, to fulfill the task to the user's satisfaction, it is essential to respect personal preferences and needs — while some want to have their coffee black, others might want it with milk, or require specific products such as lactose-free or vegan milk.

Recent research increasingly leverages Large Language Models (LLMs) to facilitate flexible and intuitive HRI as described above. A common approach is to define a set of actions or skills that the robot provides, and then use an LLM to map the high-level user instructions to the given skills [1], [2], [3], [4]. Several recent works have explored incremental learning to dynamically expand a robot's capabilities through user feedback: For example, in our previous work [5], we introduce an approach for incremental learning from natural interactions, enabling the robot to acquire new high-level behavior, rules or user preferences via natural language dialog. This approach involves storing refined versions of past interactions in a memory of interaction examples, which are retrieved for future comparable requests. Similarly, HELPER [6] and DROC [7] propose memory-augmented LLM systems to learn new behavior from interactions.

While these approaches theoretically enable some degree of customization for individual users, they overlook a critical challenge: a single humanoid robot may interact with multiple users, each with distinct –and potentially conflicting– preferences. Effectively managing and adapting to these variations remains an open problem in personalized human-robot interaction. Studies in Natural Language Processing address personalization at different levels [8], [9], [10], [11]. However, these approaches are designed for domain-specific applications and are not directly applicable to robotics, where personalization must account for embodied interaction, real-time adaptation, and multimodal sensory input.

In this paper, we extend our dialog system [5] to enable personalized robot behavior with multiple users. Specifically, we compare two approaches: (i) **explicit user profiles** by storing and retrieving explicit user-specific attributes in/from the robot's semantic memory, and (ii) **incremental interaction-based learning** as described in [5] to associate learned interaction experiences with specific users. Our preliminary experiments show that both approaches have their strengths and weaknesses indicating the need for future research to combine their advantages. Our real-world demonstration on the humanoid robot ARMAR-DE [12] showcases how the system can be integrated with automatic speech and speaker recognition [13], [14] to enable personalized dialog in real-world interactions with a humanoid robot.

II. Method

Our previous work [5] uses an Interaction LLM $L_{interact}$ in a closed-loop to produce Python statements based on the interaction history $\mathcal{H}_t = ((f_0, r_0), \ldots, (f_t, r_t))$ up to the current time t. Here, each f_i represents an invocation of one of the available robot functions $\mathcal{F} = \{F_1, \ldots, F_n\}$, while the r_i corresponds to execution results, including action success state, perceptive outputs, or errors. $L_{interact}$ is few-shot prompted with entries $\mathcal{H}$ selected from a memory $\mathcal{M}$ of interaction histories based on their similarity to the current user request. If the user gives feedback, $L_{interact}$ can call a function $F_{improve}$ that invokes an Improvement LLM $L_{improve}$ leading to an

This work has been supported by the Baden-Württemberg Ministry of Science, Research and the Arts (MWK) as part of the state's "digital@bw" digitization strategy in the context of the Real-World Lab "Robotics AI", by the Carl Zeiss Foundation through the JuBot project, and the German Federal Ministry of Education and Research (BMBF) under the Robotics Institute Germany (RIG).

The authors are with the Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Germany `{baermann,asfour}@kit.edu`

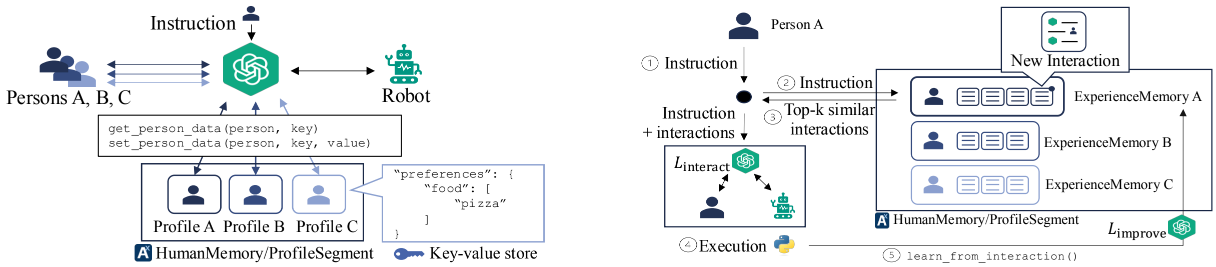

(a) explicit user profile attributes (b) incremental interaction-based learning

Fig. 1: The two approaches for personalizing humanoid robot behavior

improved version $\mathcal{H}_t^*$ of $\mathcal{H}_t$, incorporating the feedback. $\mathcal{H}_t^*$ is then saved to $\mathcal{M}$. This leads to improved behavior the next time the system receives a similar request.

To achieve personalized behavior, we expand the existing generic memory $\mathcal{M}$ with *user-specific profiles*. These profiles reside in the *ArmarX* memory [15] and are structured as dictionary-like representations with nested key-value pairs. Each profile contains both (i) explicit user attributes, such as name, height, or handedness, and (ii) a collection of personal interaction experiences with the robot. The system is provided with a unique identifier u corresponding to the current user.

Explicit user profiles uses the profiles described above as explicit key-value stores for user-specific attributes (Fig. 1a). We propose a hybrid profile structure containing some predefined keys as well as an unstructured part that can store information associated with arbitrary, LLM-generated keys. This allows other applications to reuse predefined parts of the profiles while keeping them flexible enough to match any use case. To enable the LLM to write to and read from these profiles, two special functions F_{set} and F_{get} are added to $\mathcal{F}$. Using prompt-engineering methods, the LLM is then biased to use both methods to query and store relevant user information. They also yield detailed, specifically engineered error messages when used incorrectly, guiding the LLM to explore the profiles and use the correct keys and entry types.

Incremental interaction-based learning (Fig. 1b) builds upon the interactive incremental learning mechanism described above. To enable personalization, $L_{improve}$ is additionally prompted to determine whether the given feedback is generic or user-specific (preference-based), leveraging Chain-of-Thought prompting techniques [16]. The generated improved interaction experience is then stored in the corresponding user-specific profile (if preference-based) or in a global store (if generic). When a new instruction is given to the robot, only the generic histories and those from experiences made with the current user u are used for few-shot-prompting $L_{interact}$. This results in different behavior of $L_{interact}$ depending on u and the feedback u has given in previous situations.

III. Preliminary Experiments

Both approaches were evaluated using a series of wizard-of-oz experiments, where a human operator simulated the environment, in order to focus on the system's ability to learn and recall individual user preferences over multiple interactions. For this, four test scenarios were designed along with a varying number (3 to 12) of user personas for each scenario. Each persona has specific characteristics resulting in different optimal behaviors of the robot. For example, one persona has diabetes, so when asking the robot to help them make a coffee, it is expected to bring artificial sweetener instead of sugar. Each scenario was executed in two iterations, to allow the system to pick up the different preferences in the first and recall them during the second iteration. For each user interaction conducted this way, we manually assessed whether the user preferences were respected (p) and whether the task was generally successful (t). Then, we averaged the success rates for each condition to obtain s_p, s_t, $s_{p\wedge t}$. As a baseline, we used the existing system without personalization [5].

Results Overall, both methods proved effective for personalizing the robot's task execution, although the performance of both approaches differs substantially across the four scenarios, depending on the complexity of the task and the nature of the personalized characteristics. The *explicit user profiles* approach was able to raise $s_{p\wedge t}$ from a baseline result of *32%* to *80%* while being aware of the user's preferences in $s_p =$ *92%* of cases (micro-average over each interaction of the second iteration). The *incremental interaction-based learning* approach, on the other hand, was able to generate more focused interactions with an overall higher s_t, although sometimes not properly respecting the user preferences (lower s_p). Across all scenarios, s_t was never below *80%*, while for the first approach, s_t fell to values as low as *20%* in one scenario. The system was also demonstrated qualitatively on the humanoid robot ARMAR-DE [17] in a personalized dialog scenario, integrating it with automatic speech and speaker recognition [14], [13] (see the supplementary video).

IV. Discussion

While our preliminary results demonstrate the potential of personalized humanoid robot behavior, they also show that personalization remains an open challenge. In the future, we will work on combining the strengths of our two approaches to build a more flexible system. In particular, decoupling profile attributes access from the interaction LLM could lead to less distraction from the actual task. In addition, access control mechanisms are needed to ensure that user-specific data remains private. While we support different users in different interactions, future work should also address settings with multiple persons interacting with the robot simultaneously in a group setting, a common but so far underexplored scenario.

REFERENCES

[1] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, A. Herzog, D. Ho, *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” in *Annu. Conf. Rob. Learn.*, 2022.

[2] W. Huang, F. Xia, T. Xiao, H. Chan, J. Liang, P. Florence, A. Zeng, J. Tompson, I. Mordatch, Y. Chebotar, P. Sermanet, T. Jackson, *et al.*, “Inner monologue: Embodied reasoning through planning with language models,” in *Annu. Conf. Rob. Learn.*, 2022.

[3] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, “Code As Policies: Language Model Programs for Embodied Control,” in *IEEE Int. Conf. Robot. Automat.*, 2023, pp. 9493–9500.

[4] I. Singh, V. Blukis, A. Mousavian, A. Goyal, D. Xu, J. Tremblay, D. Fox, J. Thomason, and A. Garg, “ProgPrompt: Generating situated robot task plans using large language models,” in *IEEE Int. Conf. Robot. Automat.*, 2023, pp. 11 523–11 530.

[5] L. Bärmann, R. Kartmann, F. Peller-Konrad, J. Niehues, A. Waibel, and T. Asfour, “Incremental learning of humanoid robot behavior from natural interaction and large language models,” *Frontiers in Robotics and AI*, vol. 11, 2024. [Online]. Available: https://www.frontiersin.org/journals/robotics-and-ai/articles/10.3389/frobt.2024.1455375

[6] G. Sarch, Y. Wu, M. Tarr, and K. Fragkiadaki, “Open-ended instructable embodied agents with memory-augmented large language models,” in *Conf. Emp. Meth. Nat. Lang. Proc.*, 2023, pp. 3468–3500.

[7] L. Zha, Y. Cui, L.-H. Lin, M. Kwon, M. G. Arenas, A. Zeng, F. Xia, and D. Sadigh, “Distilling and retrieving generalizable knowledge for robot manipulation via language corrections,” in *Work. Lang. Robot Learn., CoRL*, 2023.

[8] K. Zhang, F. Zhao, Y. Kang, and X. Liu, “Memory-Augmented LLM Personalization with Short- and Long-Term Memory Coordination,” Sept. 2023.

[9] Y. Wu, X. Ma, and D. Yang, “Personalized Response Generation via Generative Split Memory Network,” in *NAACL-HLT 2021*, K. Toutanova, A. Rumshisky, L. Zettlemoyer, D. Hakkani-Tur, I. Beltagy, S. Bethard, R. Cotterell, T. Chakraborty, and Y. Zhou, Eds. Online: Association for Computational Linguistics, June 2021, pp. 1956–1970.

[10] A. Salemi, S. Mysore, M. Bendersky, and H. Zamani, “LaMP: When Large Language Models Meet Personalization,” May 2023.

[11] M. Abbasian, I. Azimi, A. M. Rahmani, and R. Jain, “Conversational Health Agents: A Personalized LLM-Powered Agent Framework,” Oct. 2023.

[12] T. Asfour, M. Wächter, L. Kaul, S. Rader, P. Weiner, S. Ottenhaus, R. Grimm, Y. Zhou, M. Grotz, and F. Paus, “Armar-6: A high-performance humanoid for human-robot collaboration in real world scenarios,” *IEEE Robotics & Automation Magazine*, vol. 26, no. 4, pp. 108–121, 2019.

[13] T.-B. Nguyen and A. Waibel, “Synthetic conversations improve multi-talker asr,” in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 10 461–10 465.

[14] ——, “MSA-ASR: Efficient multilingual speaker attribution with frozen asr models,” 2025. [Online]. Available: https://arxiv.org/abs/2411.18152

[15] F. Peller-Konrad, R. Kartmann, C. R. G. Dreher, A. Meixner, F. Reister, M. Grotz, and T. Asfour, “A memory system of a robot cognitive architecture and its implementation in armarx,” *Robotics and Autonomous Systems*, vol. 164, pp. 1–20, 2023.

[16] J. Wei, X. Wang, D. Schuurmans, M. Bosma, b. ichter, F. Xia, E. H. Chi, Q. V. Le, and D. Zhou, “Chain-of-thought prompting elicits reasoning in large language models,” in *Int. Conf. Neural Inf. Process. Syst.*, 2022.

[17] T. Asfour, L. Kaul, M. Wächter, S. Ottenhaus, P. Weiner, S. Rader, R. Grimm, Y. Zhou, M. Grotz, F. Paus, D. Shingarey, and H. Haubert, “ARMAR-6: A Collaborative Humanoid Robot for Industrial Environments,” in *IEEE-RAS Int. Conf. Humanoid Robots*, 2018, pp. 447–454.