

# Expanding Robot Utility via User-Centered Adaptation

Nathaniel Dennler

University of Southern California; dennler@usc.edu

## I. RESEARCH VISION

Robots have the potential to complete a variety of tasks to aid users, from retrieving lost items in cluttered home environments to encouraging users to perform physical rehabilitation. A robot’s ability to *functionally* perform these tasks is a necessary but not sufficient property that drives robot adoption [46]. Across use cases, a robot’s utility depends on its ability to align with the individual needs and preferences of each user, but users can have conflicting preferences that prevent a single robot policy from benefiting all users [2, 41, 45]. A foundational goal of human-robot interaction is to develop robots that can be readily used by *any* user. However, current systems that take a one-size-fits all approach to robot deployment cannot satisfy users that have diametrically opposed preferences.

**Prior work** in robot learning has demonstrated the efficacy of user-collected data to allow robots to computationally reason over user preferences [4]. Users can teach robots desired behaviors through many channels such as trajectory comparisons [3, 39, 32], trajectory rankings [6, 10, 34], demonstration [27, 35, 49], corrections [1, 11, 31], language [30, 37, 44, 47], or by specifying reward values [9, 25]. An overarching limitation of these approaches is that they require conscious effort from the user to communicate over these channels. However, users naturally communicate preferences through many *embodied* channels (e.g., gaze, proxemics, or gestures) that can be computationally leveraged [29, 40].

**My research** establishes theoretical and algorithmic foundations to allow robots to adapt to users’ individual needs and preferences through natural *embodied* interfaces. I pursue this goal by (1) learning representations from data generated during robot use and (2) encoding cognitive and structural biases into algorithms to efficiently adapt robot policies with limited data.

## II. LEARNING FROM NATURALLY GENERATED DATA

Internet-scale robot datasets are scarce, but users naturally emit data through several embodied channels that can be leveraged. My first research direction addresses this gap by learning robot representations from these embodied channels.

**Robot Embodiment.** Users automatically infer expectations from the physical design of a robot [7, 36], highlighting design as a communication channel. To understand this channel, We introduced the concept of *design metaphors* [8] as a representation of users’ mental models of expected robot function [17]. We found that using these metaphors can reliably describe population-level social and functional expectations over a dataset of 165 socially interactive robots observed by over 1800 people. The semantic relationships between metaphors additionally induced corresponding measurable changes in users’ expectations of robots, e.g., robots

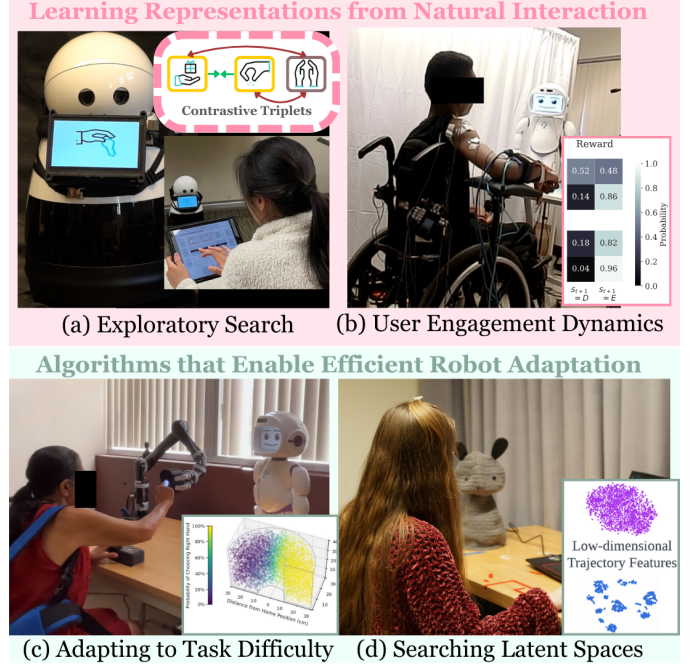


Fig. 1: Research overview. My research models natural interactions from users such as (a) exploratory search and (b) user engagement dynamics to adapt robot policies to individual users. These models facilitate adaptation in limited data domains such as (c) adjusting rehabilitation exercises and (d) finding preferred trajectories in learned representation spaces.

that were described as “a person” were perceived as more competent than robots described as “a toddler”. We also showed that design metaphors can be used as a predictive measurement to understand expectations of unseen robot embodiments [19]. Extending design as a communication channel, we created a generative framework to modulate robot expectations through clothing design [22] and screen-based faces [23].

**Exploratory Search.** Design metaphors capture users’ mental models of static robot designs, but preferences can also emerge through interaction. Traditional preference-learning assumes fixed preferences, but users often engage in exploratory search to jointly discover what they *prefer* the robot to do and what it *can* do [33]. We created an interface that allows users to perform exploratory search in high-dimensional robot action spaces [18]. This work learns from users’ *exploratory actions*, where users selectively tested embodied actions on a physical robot that assisted users with locating and fetching items, shown in Fig 1a. The users’ exploratory actions when selecting behaviors suggested their true preferences. We formulated the problem of *contrastive learning from exploratory actions* (CLEA) to learn human-aligned representations from users’

natural interactions [24], without requiring explicit user labels. We found that CLEA representations captured preferences using 80% less data than self-supervised representations in down-stream preference learning tasks.

**User Engagement Dynamics.** Both Design Metaphors and CLEA learned population-level representations of robot behaviors before users engaged in routine interaction, however, robots can learn about users continually during use. We identified engagement is a salient embodied channel that can be used to train robot social policies. We formulated successful social interaction policies as optimizing user engagement [15]. To achieve this, we learned personal *engagement dynamics* models from facial cues,  $p(s_{t+1} | s_t, a_t)$ , rather than directly predicting reward or failure as in prior works [5, 12, 43]. We evaluated this approach in a rehabilitation exercise domain for users with cerebral palsy, shown in Fig. 1b. Robots that select actions to maximize engagement according to these dynamics models resulted in significantly increased volume of interaction with robots in our user simulations.

### III. ALGORITHMS TO EFFICIENTLY ADAPT ROBOTS

Though users emit data from many channels, they also have limited time to spend adapting robots. My second research direction creates techniques to *efficiently* adapt robot behavior.

**Physical Interaction.** Robots have the potential to collect longitudinal and objective physical interaction data that can be used for several purposes, increasing their utility. We investigated how robots can learn personalized movement models of post-stroke participants for adaptive rehabilitation by creating the bimanual arm reaching test with robots (BARTR) [16, 20]. BARTR, shown in Fig. 1c, consisted of a physically assistive robot to position a reaching target and a socially assistive robot that provided feedback and test instructions. Using time to reach and arm choice data, we trained Gaussian Processes to model users from limited system use data. These models allowed us to develop a causal framework to calculate personalized task difficulty to facilitate rehabilitation [20]. We also used these models to derive an objective metric for *arm nonuse*, which is traditionally difficult to monitor in neurorehabilitation [26]. The BARTR metric was correlated with existing clinical standards [42], but achieved a higher test-retest reliability, enabling repeated nonuse assessment.

**Generating Trajectories for Hair Care.** The BARTR interaction clinically assessed post-stroke users, but the same robot can be used to augment users’ ability to perform activities of daily living. To expand the robot’s usefulness, we created the first robotic hair combing interface [14]. Specifying these paths manually requires a high degree of user effort to “draw” trajectories that appropriately comb hair, requiring high levels of skill and dexterity. We estimated the structural gradient of the user’s hair from images to plan trajectories that follow the hair’s structural gradients, requiring only one click from the user to generate combing trajectories. This technique generalizes across several hair lengths, colors, and styles. We found through a user study that this algorithm produces trajectories that are comparable to human-drawn

paths and exert minimal forces. Through a collaboration, we extended this interaction to robots with dexterous soft end effectors, extending hair care to three new tasks that leverage the structural compliance of soft robots [48].

**Intuitively Ranking Trajectories.** Robot hair combing leverages structural priors to reduce user effort. Many robot behaviors exist in learned representation spaces that do not have such intuitive structures for users to efficiently navigate [4]. Previous approaches to eliciting preferences from rankings maximize an information gain objective to generate sets of trajectories that are easy for users to rank. Maximizing this objective does not produce trajectory sets that users perceive as improving over time. I defined a novel algorithm called CMA-ES-IG that combines an information gain objective with CMA-ES optimization to produce trajectory sets that *efficiently* converge to users’ preferred trajectories in as few as ten rankings [21], compared to hundreds of pairwise comparisons. Our user studies evaluated this algorithm across robots performing social gestures and functional handover tasks, shown in Fig. 1d. Users found that CMA-ES-IG learned their preferences more quickly and easily than state-of-the-art baselines [3, 32].

### IV. FUTURE WORK

**Jointly Learning Function and Preference.** Many preference-learning approaches in robotics conflate *preference* and *function*. Users generally rate successful policies as better than unsuccessful policies, but preferences are more likely to vary between users. In dialogue generation, functional policies are often trained from data before preference learning techniques are applied to better collect user feedback [28]. My prior work assumed access to policies that *functionally* accomplish tasks from manually engineered robot systems. My future work will leverage *quality-diversity* (QD) optimization [13] to autonomously learn sets of robot policies that maximize an objective,  $f(\pi)$ , that functionally performs tasks (quality), while varying along several measure functions,  $m_i(\pi)$ , that capture differences in user preferences (diversity). I aim to evaluate QD algorithms across several user-interactive domains, such as hair care, feeding, item retrieval, and cooking. These domains enable the collection of embodied user data and will highlight the generalization of QD-based policy learning.

**Transferring Preferences Across Tasks.** My work demonstrated that the same robot can perform more than one task, such as hair combing and stroke assessment, but current preference learning techniques require robots to relearn user models for each new task. I hypothesize that some elements of users’ preferences may transfer across robot tasks, such as how quickly the robot moves during task execution [2] or the robot’s personality [38]. My future work will learn task representations to measure preference distribution similarity across robot tasks. By computing task similarity, we can improve the efficiency of preference learning in novel tasks. This approach allows robots to first learn preferences in low-risk tasks like item retrieval, and transfer those preferences to higher-risk tasks like cooking, where preference data can be more difficult to collect, improving robots’ utility for users.

## REFERENCES

- [1] Andrea Bajcsy, Dylan P Losey, Marcia K O'Malley, and Anca D Dragan. Learning from physical human corrections, one feature at a time. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 141–149, 2018.
- [2] Tapomayukh Bhattacharjee, Ethan K Gordon, Rosario Scalise, Maria E Cabrera, Anat Caspi, Maya Cakmak, and Siddhartha S Srinivasa. Is more autonomy always better? exploring preferences of users with mobility impairments in robot-assisted feeding. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*, pages 181–190, 2020.
- [3] Erdem Bıyık, Nicolas Huynh, Mykel J Kochenderfer, and Dorsa Sadigh. Active preference-based gaussian process regression for reward learning and optimization. *The International Journal of Robotics Research*, 43(5):665–684, 2024.
- [4] Andreea Bobu, Andi Peng, Pulkit Agrawal, Julie A Shah, and Anca D Dragan. Aligning human and robot representations. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, pages 42–54, 2024.
- [5] Joost Broekens. Emotion and reinforcement: affective facial expressions facilitate robot learning. In *Artificial Intelligence for Human Computing: ICMi 2006 and IJCAI 2007 International Workshops, Banff, Canada, November 3, 2006, Hyderabad, India, January 6, 2007, Revised Selected and Invited Papers*, pages 113–132. Springer, 2007.
- [6] Daniel Brown, Wonjoon Goo, Prabhat Nagarajan, and Scott Niekum. Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations. In *International conference on machine learning*, pages 783–792. PMLR, 2019.
- [7] Colleen M Carpinella, Alisa B Wyman, Michael A Perez, and Steven J Stroessner. The robotic social attributes scale (rosas) development and validation. In *Proceedings of the 2017 ACM/IEEE International Conference on human-robot interaction*, pages 254–262, 2017.
- [8] John M Carroll, Robert L Mack, and Wendy A Kellogg. Interface metaphors and user interface design. In *Handbook of human-computer interaction*, pages 67–85. Elsevier, 1988.
- [9] Carlos Celemin and Javier Ruiz-del Solar. An interactive framework for learning continuous actions policies based on corrective feedback. *Journal of Intelligent & Robotic Systems*, 95:77–97, 2019.
- [10] Letian Chen, Rohan Paleja, and Matthew Gombolay. Learning from suboptimal demonstration via self-supervised reward regression. In *Conference on robot learning*, pages 1262–1277. PMLR, 2021.
- [11] Yuchen Cui and Scott Niekum. Active reward learning from critiques. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 6907–6914. IEEE, 2018.
- [12] Yuchen Cui, Qiping Zhang, Brad Knox, Alessandro Allievi, Peter Stone, and Scott Niekum. The empathic framework for task learning from implicit human feedback. In *Conference on robot learning*, pages 604–626. PMLR, 2021.
- [13] Antoine Cully and Yiannis Demiris. Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation*, 22(2):245–259, 2017.
- [14] Nathaniel Dennler, Eura Shin, Maja Matarić, and Stefanos Nikolaidis. Design and evaluation of a hair combing system using a general-purpose robotic arm. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3739–3746. IEEE, 2021.
- [15] Nathaniel Dennler, Catherine Yunis, Jonathan Realmuto, Terence Sanger, Stefanos Nikolaidis, and Maja Matarić. Personalizing user engagement dynamics in a non-verbal communication game for cerebral palsy. In *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*, pages 873–879. IEEE, 2021.
- [16] Nathaniel Dennler, Amelia Cain, Erica De Guzman, Claudia Chiu, Carolee J Winstein, Stefanos Nikolaidis, and Maja J Matarić. A metric for characterizing the arm nonuse workspace in poststroke individuals using a robot arm. *Science Robotics*, 8(84):eadf7723, 2023.
- [17] Nathaniel Dennler, Changxiao Ruan, Jessica Hadiwijoyo, Brenna Chen, Stefanos Nikolaidis, and Maja Matarić. Design metaphors for understanding user expectations of socially interactive robot embodiments. *ACM Transactions on Human-Robot Interaction (T-HRI)*, 12(2):1–41, 2023.
- [18] Nathaniel Dennler, David Delgado, Daniel Zeng, Stefanos Nikolaidis, and Maja Matarić. The rosid tool: Empowering users to design multimodal signals for human-robot collaboration. *International Symposium of Experimental Robotics (ISER)*, 2024.
- [19] Nathaniel Dennler, Stefanos Nikolaidis, and Maja Matarić. Singing the body electric: The impact of robot embodiment on user expectations. *Robotics: Science and Systems (RSS) Workshop on Social Intelligence in Humans and Robots*, 2024.
- [20] Nathaniel Dennler, Stefanos Nikolaidis, and Maja Matarić. Using causal trees to estimate personalized task difficulty in post-stroke individuals. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) Workshop on Assistive Robotics for Citizens*, 2024.
- [21] Nathaniel Dennler, Zhonghao Shi, Stefanos Nikolaidis, and Maja Matarić. Improving user experience in preference-based optimization of reward functions for assistive robots. *International Symposium of Robotics Research (ISRR)*, 2024.
- [22] Nathaniel Dennler, Mina Kian, Stefanos Nikolaidis, and Maja Matarić. Designing robot identity: The role of

- voice, clothing, and task on robot gender perception. *International Journal of Social Robotics (IJSR)*. In Press, 2025.
- [23] Nathaniel Steele Dennler, Evan Torrence, Uksang Yoo, Stefanos Nikolaidis, and Maja Mataric. Pylips: an open-source python package to expand participation in embodied interaction. In *Adjunct Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, pages 1–4, 2024.
- [24] Nathaniel Steele Dennler, Stefanos Nikolaidis, and Maja Mataric. Contrastive learning from exploratory actions: Leveraging natural interactions for preference elicitation. In *2025 ACM/IEEE International Conference on Human-Robot Interaction*, 2025.
- [25] Shane Griffith, Kaushik Subramanian, Jonathan Scholz, Charles L Isbell, and Andrea L Thomaz. Policy shaping: Integrating human feedback with reinforcement learning. *Advances in neural information processing systems*, 26, 2013.
- [26] Cheol E Han, Sujin Kim, Shuya Chen, Yi-Hsuan Lai, Jeong-Yoon Lee, Rieko Osu, Carolee J Winstein, and Nicolas Schweighofer. Quantifying arm nonuse in individuals poststroke. *Neurorehabilitation and neural repair*, 27(5):439–447, 2013.
- [27] Dana Hughes, Akshat Agarwal, Yue Guo, and Katia Sycara. Inferring non-stationary human preferences for human-agent teams. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 1178–1185. IEEE, 2020.
- [28] Natasha Jaques, Asma Ghandeharioun, Judy Hanwen Shen, Craig Ferguson, Agata Lapedriza, Noah Jones, Shixiang Gu, and Rosalind Picard. Way off-policy batch deep reinforcement learning of implicit human preferences in dialog. *arXiv preprint arXiv:1907.00456*, 2019.
- [29] Hong Jun Jeon, Smitha Milli, and Anca Dragan. Reward-rational (implicit) choice: A unifying formalism for reward learning. *Advances in Neural Information Processing Systems*, 33:4415–4426, 2020.
- [30] Jason Xinyu Liu, Ziyi Yang, Ifrah Idrees, Sam Liang, Benjamin Schornstein, Stefanie Tellex, and Ankit Shah. Grounding complex natural language commands for temporal tasks in unseen environments. In *Conference on Robot Learning*, pages 1084–1110. PMLR, 2023.
- [31] Dylan P Losey and Marcia K O’Malley. Including uncertainty when learning from human corrections. In *Conference on Robot Learning*, pages 123–132. PMLR, 2018.
- [32] Shihan Lu, Mianlun Zheng, Matthew C Fontaine, Stefanos Nikolaidis, and Heather Culbertson. Preference-driven texture modeling through interactive generation and search. *IEEE transactions on haptics*, 15(3):508–520, 2022.
- [33] Gary Marchionini. Exploratory search: from finding to understanding. *Communications of the ACM*, 49(4):41–46, 2006.
- [34] Vivek Myers, Erdem Biyik, Nima Anari, and Dorsa Sadigh. Learning multimodal rewards from rankings. In *Conference on robot learning*, pages 342–352. PMLR, 2022.
- [35] Stefanos Nikolaidis, Ramya Ramakrishnan, Keren Gu, and Julie Shah. Efficient model learning from joint-action demonstrations for human-robot collaborative tasks. In *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction*, pages 189–196, 2015.
- [36] Steffi Paepcke and Leila Takayama. Judging a bot by its cover: An experiment on expectation setting for personal robots. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 45–52. IEEE, 2010.
- [37] Shreyas Sundara Raman, Vanya Cohen, Eric Rosen, Ifrah Idrees, David Paulius, and Stefanie Tellex. Planning with large language models via corrective re-prompting. In *NeurIPS 2022 Foundation Models for Decision Making Workshop*, 2022.
- [38] Silvia Rossi, François Ferland, and Adriana Tapus. User profiling and behavioral adaptation for hri: A survey. *Pattern Recognition Letters*, 99:3–12, 2017.
- [39] Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. *Active preference-based learning of reward functions*. 2017.
- [40] Albrecht Schmidt. Implicit human computer interaction through context. *Personal technologies*, 4:191–199, 2000.
- [41] Zhonghao Shi, Han Chen, Anna-Maria Velentza, Siqi Liu, Nathaniel Dennler, Allison O’Connell, and Maja Mataric. Evaluating and personalizing user-perceived quality of text-to-speech voices for delivering mindfulness meditation with different physical embodiments. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pages 516–524, 2023.
- [42] Annette Sterr, Susanna Freivogel, and Dieter Schmalohr. Neurobehavioral aspects of recovery: assessment of the learned nonuse phenomenon in hemiparetic adolescents. *Archives of physical medicine and rehabilitation*, 83(12): 1726–1731, 2002.
- [43] Maia Stiber, Russell Taylor, and Chien-Ming Huang. Modeling human response to robot errors for timely error detection. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 676–683. IEEE, 2022.
- [44] Aaqib Tabrez, Jack Kawell, and Bradley Hayes. Asking the right questions: Facilitating semantic constraint specification for robot skill learning and repair. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6217–6224. IEEE, 2021.
- [45] Adriana Tapus, Cristian Țăpuș, and Maja J Matarić. User—robot personality matching and assistive robot behavior adaptation for post-stroke rehabilitation therapy. *Intelligent Service Robotics*, 1:169–183, 2008.

- [46] Viswanath Venkatesh and Fred D Davis. A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management science*, 46(2): 186–204, 2000.
- [47] Jimmy Wu, Rika Antonova, Adam Kan, Marion Lepert, Andy Zeng, Shuran Song, Jeannette Bohg, Szymon Rusinkiewicz, and Thomas Funkhouser. Tidybot: Personalized robot assistance with large language models. *Autonomous Robots*, 47(8):1087–1102, 2023.
- [48] Uksang Yoo, Nathaniel Steele Dennler, Eliot Xing, Maja Matarić, Stefanos Nikolaidis, Jeffrey Ichnowski, and Jean Oh. Soft and compliant contact-rich hair manipulation and care. In *2025 ACM/IEEE International Conference on Human-Robot Interaction*, 2025.
- [49] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.