

Investigating the Impact of Experience on a User’s Ability to Perform Hierarchical Abstraction

Nina Moorman*

Georgia Institute of Technology
nmoorman3@gatech.edu

Nakul Gopalan*

Arizona State University
nakul.gopalan@asu.edu

Aman Singh

Georgia Institute of Technology
asingh609@gatech.edu

Erin Hedlund-Botti

Georgia Institute of Technology
erin.botti@gatech.edu

Mariah Schrum

Georgia Institute of Technology
mschrum3@gatech.edu

Chuxuan Yang

Georgia Institute of Technology
soyang@gatech.edu

Lakshmi Seelam

Georgia Institute of Technology
lseelam3@gatech.edu

Matthew C. Gombolay

Georgia Institute of Technology
matthew.gombolay@cc.gatech.edu

Abstract—The field of Learning from Demonstration enables end-users, who are not robotics experts, to shape robot behavior. However, using human demonstrations to teach robots to solve long-horizon problems by leveraging the hierarchical structure of the task is still an unsolved problem. Prior work has yet to show that human users can provide sufficient demonstrations in novel domains without showing the demonstrators explicit teaching strategies for each domain. In this work, we investigate whether non-expert demonstrators can generalize robot teaching strategies to provide necessary and sufficient demonstrations to robots *zero-shot in novel domains*. We find that increasing participant experience with providing demonstrations improves their demonstration’s degree of sub-task abstraction ($p < .001$), teaching efficiency ($p < .001$), and sub-task redundancy ($p < .05$) in novel domains, allowing generalization in robot teaching. Our findings demonstrate for the first time that non-expert demonstrators can transfer knowledge from a series of training experiences to novel domains without the need for explicit instruction, such that they can provide necessary and sufficient demonstrations when programming robots to complete task and motion planning problems.

I. INTRODUCTION

Due to the diversity of end users and deployment settings, it is intractable for a robot to be pre-programmed to do any task in any environment. One solution is to allow robots to learn new skills in situ, from end-users. Prior work in Learning from Demonstration (LfD) has investigated how to allow non-roboticist end-users to operate in the role of the robot teacher [8, 15, 46, 47, 49] in order to communicate personal preferences and leverage their domain knowledge [37, 44, 56]. However, many previous approaches require human demonstrators to learn how to teach the robot tasks in each domain, using videos or demonstrations from robot experts [3, 22, 40]. This approach does not scale up in enabling human users to teach a variety of tasks to a robot, as demonstrator training is domain-dependent, and an expert is still required to be in the training loop. In this work, we develop a series of demonstrator training tasks through which participants obtain knowledge about providing sufficient

and necessary demonstrations that can generalize to novel domains.

Substantial emphasis in LfD has been placed on teaching robots single, short-horizon *skills*, such as picking up or making contact with an object [23, 29, 34, 38]. However, there is a lack of work enabling robots to learn long-horizon *tasks*, such as learning in-home assistive tasks or manufacturing process assembly operations, from human demonstrations. Such tasks can be considered multi-task problems. For example, setting a dinner table would require a robot to set multiple place settings, dependent on the number of guests, where each table setting consists of multiple objects that each require a different manipulation procedure. A demonstrator cannot be expected to provide demonstrations for each task specification of these multi-tasks, such as for each possible number of plate settings. Since long-horizon tasks require a robot to solve repetitive multi-task problems, demonstrators need to break up their demonstrations into shorter abstractions that a robot can reuse. Prior work has shown that demonstrators are capable of teaching abstractions when explicitly instructed on how to do so. Akgun et al., for instance, demonstrated that users can teach Keyframe-based abstractions for robot movement [3]. Likewise, Mohseni-Kabir et al., demonstrated that, for hierarchical tasks, novel sub-tasks can be taught to a robot using combinations of previously taught sub-tasks [40]. However, if not prescribed how to teach the robot, prior work has found that participants struggle to provide demonstrations that exhibit abstractions sufficient for a hierarchical task [32]. Gopalan et al. similarly find that the majority of participants did not naturally teach the robot tasks using abstractions [22]. The authors compare various modes of demonstrator training and find that only videos where the experimenter demonstrates how to provide demonstrations using sub-task abstractions in the current domain enabled participants to use abstractions (in the same domain) that would be robust to novel task specifications [22]. In this work, we investigate whether, given enough practice in previous domains, participants can use

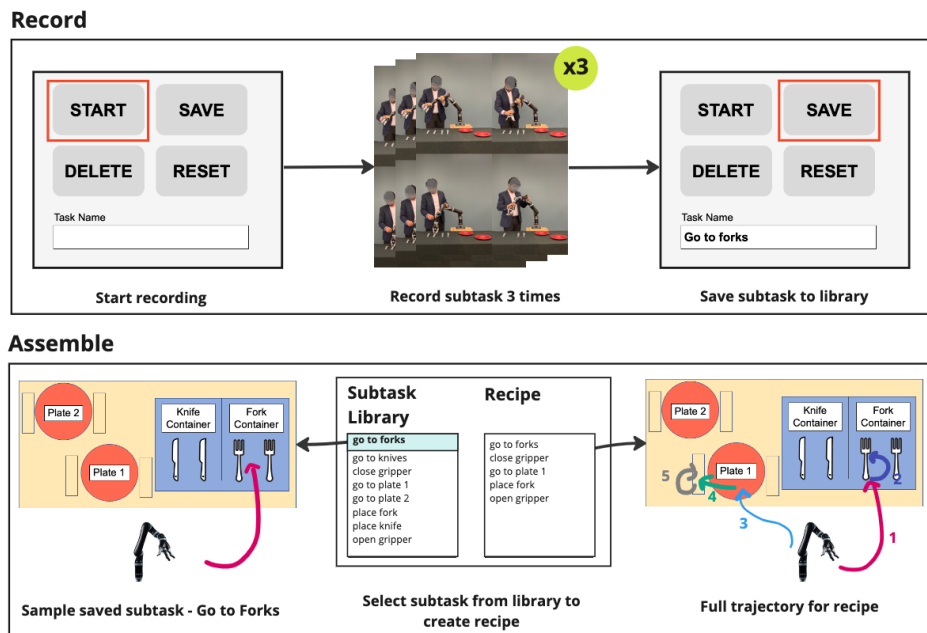


Fig. 1: Via the interface and kinesthetic teaching, participants record three demonstrations to save a sub-task. Saved sub-tasks are then available in the interface library. To execute a task, participants assemble a recipe from the set of recorded sub-tasks.

sub-task abstractions to provide sufficient demonstrations in a novel domain without being told how to do so.

Instead of training demonstrators from scratch in each domain they encounter, we want the knowledge the demonstrator learns about providing demonstrations in one domain to be transferable to novel domains. We develop a three-hour training procedure with five domains and corresponding unpersonalized expert videos. This procedure enables demonstrators to transfer acquired knowledge about providing sufficient and necessary demonstrations – gained through trial, error, and expert feedback – to novel domains where no expert feedback is available. In this work, we study the impact of domain experience on users’ ability to provide demonstrations when teaching a robot. Here, we define domain experience as the number of training domains for which (i) the participant has taught the robot a task via demonstration and (ii) the demonstrator has subsequently watched a training video showing a robotics expert providing the optimal teaching strategy. We employ a standard way to formulate multi-modal, multi-task problems in robotics [20], namely Task and Motion Planning (TAMP), for our sub-task representations.

To the best of our knowledge, our findings show for the first time that humans can transfer knowledge from a few training experiences to provide sufficient TAMP demonstrations in a novel domain. In this work, we contribute the following.

- 1) We design a novel user study and set of training domains to investigate whether participants improve their ability to teach the robot via demonstrations over time, without the use of a curriculum.
- 2) Our results demonstrate that participants are able to generalize knowledge about task abstraction ($p < .001$),

teaching efficiency ($p < .001$), and redundancy ($p < .05$) zero shot to novel domains.

- 3) We additionally find that prior teaching experience impacts sub-task count ($p = .011$), that sub-task count impacts perceived workload ($p = .005$) as well as robot likeability ($p = .007$), and that participant agreeableness impacts teaching duration ($p = .006$).

II. METHODS

A. Study Design

We conducted a 1×4 within-subjects experiment with twenty-eight participants, seven per ordering condition (see Appendix for the domain ordering of each condition). Participants experience five domains in this study, a practice domain that all participants experience first, and the four ordered domains. We control for the ordering of the remaining four domains using a Latin square, ensuring that the participant count per condition is balanced. The independent variable in this study is the number of domains encountered thus far. The robot employed in this study is the JACO arm [12] attached to a hand-crafted base located next to the experiment’s table. We additionally designed a user interface that allows users to record and save sub-tasks they demonstrate to the robot. Participants can then use the interface to combine different sub-tasks to accomplish a task. We require the participant to record three demonstrations for each sub-task in order to capture variability in the way the participant moves the robot for robustness to noise.

Research Questions

RQ1: What is the impact of domain experience on the quality of demonstrations? We investigate whether participants can

perform zero-shot transfer to novel domains of any acquired knowledge as measured by sub-task abstraction score, teaching efficiency, and sub-task redundancy.

RQ2: *What is the effect of demonstration abstraction on participants' perceived workload?* We hypothesize that higher abstraction scores will reduce the repetitiveness of participant demonstrations, thereby reducing perceived workload.

RQ3: *Do participant demographics impact the quality of demonstrations?* We investigate whether participant demographics, such as prior robotics experience and prior teaching experience, impact the quality of their demonstrations. We posit that participants with robotics or teaching experience will teach the robot, via demonstration, more effectively and efficiently.

RQ4: *Does domain type impact the quality of demonstrations?* We hypothesize that the domain type will impact the sub-task count and redundancy, abstraction score, and teaching duration.

B. Metrics

The objective metrics we collect in our user study are as follows. These metrics are collected per domain, for each participant. We employ the abstraction scoring method validated in [22]. In this scoring method, one point is allotted for each instance of a sufficient sub-task employed to accomplish a task, and one point is awarded for sufficient sub-tasks that could be constructed by composing other sufficient sub-tasks. The latter ensures that finer-grain abstractions are scored higher than coarser-grain abstractions, within reason (no points are awarded for gratuitously low-level abstractions such as *move left*). The abstraction scoring rubrics employed in each domain can be found in the Appendix. This abstraction score metric allows us to evaluate the sufficiency of demonstration sub-tasks. We additionally count the number of redundant sub-tasks taught, i.e., sub-tasks whose function can be fulfilled by another existing sub-task or a combination of existing sub-tasks. This metric allows us to evaluate the necessity, independently from the sufficiency, of demonstration sub-tasks. Furthermore, we count the total number of sub-tasks taught to the robot and the total time the participant taught the robot, including time using the interface and time spent providing the kinesthetic demonstrations.

The subjective metrics in our user study are as follows (the details of hand-crafted surveys, Cronbach's alpha, qualitative results, and quotes from interview questions are in the Appendix). In the pre-study questionnaire, we collect participants' age, gender, education, and race/ethnicity. We additionally administer the Big Five Personality survey [21], the Negative Attitudes Towards Robotics (NARS) Scale [51], and hand-crafted prior robotics experience and prior teaching experience surveys. In the post-domain questionnaire, we ask participants to "please explain your strategy and thought process when teaching the robot in this domain." Finally, in the post-study questionnaire, we administer the NASA Task Load Index (NASA TLX) [24] and the Perceived Intelligence and Likeability sub-scales of the Godspeed Questionnaire

Series [7]. We additionally ask participants five post-interview questions.

Procedure

This study was approved by our university's Institutional Review Board (IRB), protocol #H22450. We recruited all participants through advertisements on campus. The study took three hours, and participants were compensated with a \$50 Amazon gift card. The procedure of the study is as follows. Participants first take the pre-study questionnaire. After the pre-study questionnaire, participants observe the introduction video (link and description in Appendix). Next, participants teach the robot to complete the block-touching tasks in the demo (i.e., practice) domain. After this demo domain, the participant explains their teaching strategy, recorded via a voice recording. This block-touching demo domain serves to familiarize the participant with moving the robot and using the interface.

Then, for the testing portion of the study, participants teach the robot how to accomplish tasks in four different domains: box packing, table setting, soil mixing, and medicine dispensing. Each participant experiences one ordering condition, which defines the order in which the domains are encountered. All participants experience each of these domains (within subjects). For each of these four domains, participants are introduced to the domain verbally, then asked to teach the robot how to do three tasks in that domain using the interface, as seen in Figure 1. To teach the sub-tasks, participants provide kinesthetic demonstrations in which participants physically manipulate the robot. After teaching the robot, the participants answer the post-domain interview question and then observe a video showing the optimal way of teaching the robot in that domain (communicating the proper sub-task breakdown) prior to experiencing the next novel domain. The optimal teaching strategy video for each domain was designed to communicate how to optimally teach the robot, listing the optimal sub-tasks, along with how to teach and record those sub-tasks on the robot using the interface. Next, the videos show how to use the sub-tasks to build the recipe for one task in the domain. After experiencing all four domains, participants take the post-survey questionnaire. Finally, participants answer the post-interview questions.

For each demonstration saved via the interface, we record the robot trajectories along with a third-person perspective video of the participant moving the robot, collected using a Kinect camera. While participants record their sub-tasks, the experimenter takes detailed notes on participant behavior, recording which sub-tasks they record. Using notes, three coders scored participant abstraction and redundancy scores, resulting in an intra-class correlation coefficient of 0.998 for abstraction scores and 0.755 for redundancy scores.

III. RESULTS AND DISCUSSION

We conducted our study with 28 participants (39.26% female, mean age = 22.89, standard deviation = 1.63). Before running statistical tests, we first checked that our data met

parametric assumptions via Shapiro-Wilk’s test and Levene’s test. Due to our statistical models not passing tests for normality, we employ non-parametric tests throughout our analysis. We employ Bonferroni correction when applying multiple tests for the same hypothesis to reduce the risk of Type I errors [48]. To test RQ1 and RQ4 we employ the Friedman rank sum test. For follow-up pairwise comparisons, we employ the Nemenyi Wilcoxon-Wilcox all-pairs test, for which we report the p-value. To test RQ2 and RQ3 we employ Spearman’s rank correlation test.

Impact of domain experience on the demonstration sufficiency, necessity, and efficiency (RQ1).

We find that participant abstraction score is positively impacted by the number of domains experienced ($p < .001$), meaning that over time participants provide demonstrations that manifest higher levels of abstraction. We further find that teaching duration is negatively impacted by the number of domains experienced ($p < .001$), and that redundancy score is negatively impacted by the number of domains experienced ($p = .046$). These findings indicate that participants can generalize knowledge gained about providing demonstrations efficiently, using more sub-task abstraction from previously experienced domains to a novel domain. **These findings indicate that demonstrators can be trained via demonstration training to efficiently provide sufficient and necessary demonstrations to new domains, zero-shot.**

Impact of prior teaching experience on sub-task count (RQ3).

We find that prior teaching experience is negatively correlated with sub-task count ($p = .011$), indicating that participants with more teaching experience record fewer sub-tasks. We note that we don’t find significance between prior teaching experience and abstraction score or redundancy score. **This finding indicates that increasing teaching experience will increase sub-task efficiency, decreasing sub-task count, though not at the expense of sub-task sufficiency.** Since general teaching experience does not appear to translate to demonstration quality, our findings highlight the need for a way to teach demonstrators how to provide sufficient and necessary sub-tasks. Our results show that we contribute a scalable and generalizable method for training LfD demonstrators, by exposing demonstrators to multiple domains in which they practice and observe the optimal teaching method.

Impact of sub-task count on perceived workload (RQ2).

We find that participant workload is negatively correlated with their sub-task count ($p = .005$). This indicates that a lower sub-task count correlated with a higher perceived workload. We hypothesize that this is due to the lengthier process of demonstrating and recording under-abstracted sub-tasks. **In addition to abstractions being useful for robust robot learning, this finding suggests that participants find correct abstractions less effort to teach, as observed via decreased perceived workload.**

Impact of robot likeability, participant agreeableness, and negative attitudes on demonstrations (RQ3).

We find that robot likeability is negatively correlated with sub-task count ($p = .007$). **This suggests that people rated the robot as more likeable when the teaching was less involved.** On the other hand, we find that participant agreeableness is positively correlated with teaching duration ($p = .006$). This finding suggests that demonstrators with higher agreeableness take longer when providing demonstrations, though not at the expense of sub-task count, abstraction score, or redundancy. **This finding indicates that more agreeable demonstrators take their time when recording demonstrations.** We posit this is because these participants either wanted to please the experimenter or because they wanted to be thorough in order to be helpful.

Participants that perceived robots as more socially negative additionally took longer to teach the robot ($p = .001$). **Participants that are warier of robots take more time to provide demonstrations,** therefore we posit that addressing negative robot perceptions will reduce the time people take to teach robots.

Limitations and Future Work

A limitation of our work is that participants do not observe the learned robot behavior (resulting from their demonstrations). This absence of observing the subsequent consequences on the environment of their demonstrations may have reduced participant urgency and desire to improve. Additionally, as our abstraction scoring method is taken from prior work which validates it on a real robot system [22], we do not validate the scoring method again in this study. In future work, we propose to further validate our abstraction score and original subjective surveys.

IV. CONCLUSION

Learning from demonstration enables non-expert end-users to be involved in robot learning. However, providing usable demonstrations is not intuitive to most demonstrators. Demonstrators have to be trained in order to provide demonstrations that would be usable, and this training is often domain-specific. Instead of teaching demonstrators how to provide sufficient demonstrations in all possible domains, we propose to teach demonstrators such that what they learn can generalize to novel domains. In this work, we study the impact of experience providing demonstrations across multiple domains on the quality of demonstrations for LfD. We find that as participants gain domain experience they are able to generalize knowledge about sub-task abstraction ($p < .001$), teaching efficiency ($p < .001$), and sub-task redundancy ($p < .05$) zero shot, to novel task domains. We show that with a few hours of training, we can teach human demonstrators to provide sufficient, necessary, and efficient demonstrations in novel domains.

APPENDIX A PRELIMINARIES

In this section, we define terms pertaining to our work. We include a separate related works section preceding our discussion to contextualize our findings.

Multi-task problems – In multi-task problems, the objects that the robot interacts with remain the same. However, the number of objects, their locations, or the order in which the robot interacts with the objects changes between sub-tasks [13]. This is often accomplished by leveraging similarities between the sub-tasks [55].

Multi-modal tasks – A mode is a sub-manifold of robot motion within which the robot’s contact specification, with respect to different objects in the world, remains constant [4, 5, 26, 27]. A multi-modal task is one where the robot transitions between at least two modes to solve a task. For example, to pick up a block, the robot first is restricted to a mode where all its motion is confined to a sub-manifold within which the robot’s gripper is not in contact with any object. After picking up the block, the mode of the robot is the sub-manifold within which it is in continuous contact with the block. Similarly, a *Long Horizon Task* is a task where the robot needs to perform multiple mode switches to solve the task. Thus, per our definition, multi-modal tasks have at least one mode switch (≥ 1), and long horizon tasks have several (> 2) mode switches. In our work, the robot is solving multi-modal tasks.

Task and Motion Planning (TAMP) – Robotics problems require an interplay between symbolic and continuous domains. For example, to pass medicines to a patient, the robot needs to make a high-level symbolic plan to know which boxes of medicines to pick up and pass to a patient. This plan and its corresponding state are symbolic and discrete over the type and quantities of medicine box objects required. However, to pick up a box the robot needs to create a continuous motion plan without collisions such that the box is in the robot’s hand. This motion planning problem occurs over the continuous state of the robot’s joints. Such problems, that exhibit an interplay between symbolic and continuous plans, are TAMP problems. We choose to define our domains as TAMP problems, as they require an interplay between symbolic goal states and continuous motion from the robot.

Sub-task based abstraction – Transitions between the symbolic states of a TAMP problem are called sub-tasks¹ [20]. For example, when a robot moves to pick up a cup, the state of the world transitions symbolically, such that the cup is in the robot’s hand. In TAMP formulations, the sub-tasks are described by preconditions (`pre`) and effects (`eff`), as well as constraints (`con`) that must hold for all continuous actions for the duration of time the action is being taken. We provide a sample mathematical TAMP formulation for the sub-tasks of the medicine dispensing domain in the Appendix.

¹Sub-tasks are referred to as actions in [20]; however, we refer to these actions as sub-tasks to prevent confusion between low-level robot actions and TAMP level actions.

Sufficient sub-tasks – In our domains, a sub-task is deemed *sufficient* if the sub-task changes the symbolic state of the world and results in at most one mode change. For a sub-task to change the symbolic state of the world, the change must go beyond a negligible change in the robot’s pose. Moreover, limiting the sub-task to at most one mode change ensures that the robot can change its interaction with only one object within the sub-task. Such design of sub-tasks ensures that a sub-task transition affects only a small set of symbolic state variables at a time. These sub-tasks can then be sequenced by a task planner to reach a larger set of the symbolic state space, allowing maximal generalizability in the tasks that can be solved within the domain.

Redundant sub-tasks – A sub-task is deemed redundant if its goal can be met by another sufficient sub-task or a combination of sufficient sub-tasks previously taught. Sub-task redundancy is defined with respect to a given set of demonstrations being taught.

Necessary sub-tasks – Similarly, a sub-task is deemed necessary if its goal can *not* be met by another sufficient sub-task or a combination of sufficient sub-tasks previously taught. Sub-task necessity is defined with respect to a given set of demonstrations being taught. A sub-task is deemed necessary if it is not a redundant sub-task.

Domain experience – We define domain experience, a metric for demonstrator training, as the number of domains experienced thus far in the user study. Note that for each domain, this experience entails participants first providing demonstrations to the robot, then observing the optimal teaching sub-task breakdown in the form of a video.

APPENDIX B RELATED WORKS

In this section, we discuss relevant prior work in robot learning from human demonstrators and hierarchical task representations to further contextualize and motivate our results prior to the discussion.

Learning from Human Demonstrators

The field of LfD explores how human demonstrations can be used to teach robots new skills [15, 46, 8, 43, 14, 28]. LfD enables the agent to learn from a small set of examples, i.e., demonstrations provided by a teacher rather than learning from lengthy exploration, i.e., experience collected in an environment [6]. The mode of demonstration collection depends on whether the LfD algorithm aims to model the human feedback [33], latent reward function [1, 19, 57], or unknown robot policy directly [28]. Additional LfD design decisions include accounting for prior robotics experience, demonstrator sub-optimality, and demonstration heterogeneity [50, 49, 14, 10, 41]. **In this paper, we evaluate how well non-experts can teach robots kinesthetically without explicitly being taught by domain experts.**

Much prior work in LfD has focused on enabling robots to perform short-horizon skills [29, 34, 23, 38]. There has been a

lack of approaches using LfD to train robots to perform long-horizon tasks and multi-tasks, which would require the robot to accomplish a series of shorter sub-tasks. LfD approaches to multi-task learning are often expensive since demonstrators need to be taught how to provide demonstrations for each of the tasks required [17]. We show that demonstrators can generalize knowledge about providing sufficient and necessary demonstrations to novel tasks. **These findings suggest that demonstrators do not need to be taught how to teach in each domain explicitly, making long-horizon, multi-task LfD more tractable.**

Hierarchical Task Representations

Prior work has investigated how the hierarchical nature of a task can facilitate long-horizon task completion [2, 36, 39]. Breaking up a long-horizon task hierarchically and abstracting the task components into repurposable sub-tasks reduces planning depth and allows for faster planning [25, 30]. This representation of the task affords the agent the ability to adapt to novel environments that share features with the distribution of environments previously experienced [18].

Currently, LfD demonstration collection requires multiple demonstrations for each possible configuration of a multi-task setup. One way to make this process more tractable is for demonstrators to break up the task into a series of shorter abstractions that the robot could reuse for multiple configurations within the same domain. Prior work has shown that users can teach robots using task abstractions [3, 40]. Akgun et al. found that demonstrators can teach abstractions for robot end-effector movement [3], and Mohseni-Kabir et al. found that participants could teach the robot novel sub-tasks using previously taught sub-tasks as building blocks [40]. However, in most prior work establishing a human demonstrator’s ability to provide usable demonstrations that contain abstractions, the participants are shown precisely how to teach the robot. They are then asked to reproduce the method of robot teaching that was prescribed. Cakmak et al. compare written and video demonstrator instruction, and find that trial and error plays a large role in the learning process [11]; we note that these demonstrators learn and are evaluated on the same task. Teaching a robot using abstractions without this guidance is not intuitive to non-experts [32]. Gopalan et al. find that the majority of participants are unable to provide sufficiently abstracted demonstrations naturally, and find that, even when told to employ abstractions, demonstrators struggle to provide demonstrations robust to minor changes in the task specification, such as item multiplicity or item location [22]. In this work, we investigate whether demonstrators’ ability to provide sufficient sub-task abstractions improves over time, as they practice providing demonstrations in multiple different domains. **Our findings support participants’ ability to learn to provide sufficient sub-task abstractions in novel domains, with enough practice.**

There have been algorithmic approaches to learning tasks from user demonstrations without requiring the demonstrations to specify task abstractions [35, 31, 16]. However, these

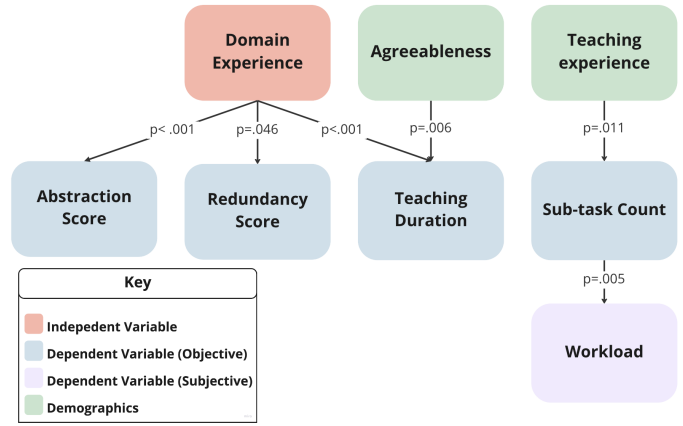


Fig. 2: Depicted is a summary of our significant results.

approaches require the collection of demonstration datasets that would not be scalable for multi-task settings. In order to investigate the scalability of training non-expert demonstrators, we additionally investigate whether participants can generalize knowledge about providing demonstrations zero-shot, to novel domains. **Our findings support that, rather than training demonstrators in each domain they encounter, experimenters could train demonstrators in a handful of training domains, for demonstrators to generalize this training to novel domains.**

APPENDIX C RESULTS

We conducted our study with 28 participants (39.26% female, mean age = 22.89, standard deviation = 1.63). Before running statistical tests, we first checked that our data met parametric assumptions via Shapiro-Wilk’s test and Levene’s test. Due to our statistical models not passing tests for normality, we employ non-parametric tests throughout our analysis. We employ Bonferroni correction when applying multiple tests for the same hypothesis to reduce the risk of Type I errors [48]. To test RQ1 and RQ4 we employ the Friedman rank sum test, where we report χ^2 (degree of freedom) and p-value. For follow-up pairwise comparisons, we employ the Nemenyi Wilcoxon-Wilcox all-pairs test, for which we report the p-value. To test RQ2 and RQ3 we employ Spearman’s rank correlation test, where we report ρ and p-value.

Research Question 1

We first study the impact of domain experience on the quality of demonstrations. This hypothesis investigates whether participants can perform zero-shot transfer of knowledge regarding sub-task abstraction, teaching efficiency, and sub-task redundancy to novel domains.

We note that the block touching domain was the demo task, intended to familiarize the participant with the robot and the interface, to isolate the effect of learning in the actual test rounds. As the participants do not observe the optimal demonstration after this demo task, we do not include the block-touching domain in our domain experience.

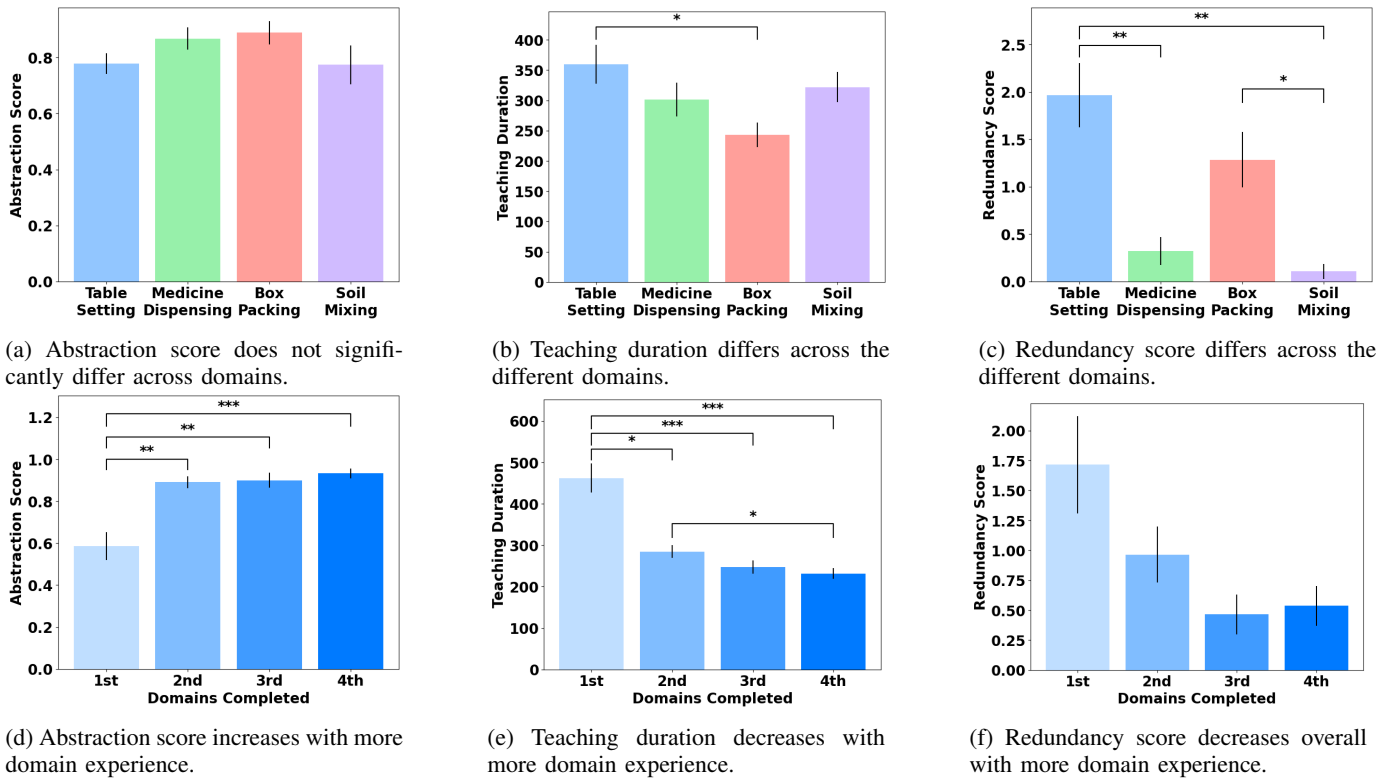


Fig. 3: We depict results with respect to domain type (top row) and domain experience (bottom row).

Abstraction Score – Through a Friedman test, we find a main effect of the participant’s domain experience on the participant’s domain abstraction score ($\chi^2(3) = 28.056, p < .001$). We conduct pairwise comparisons using a Nemenyi Wilcoxon-Wilcox all-pairs test, visualized in Figure 3d, and find significance between the first and second domain ($p = .006$), the first and third domain ($p = .001$), and the first and fourth domain experienced ($p < .001$).

We first observe in Figure 3d that abstraction scores improve between the first and second domains. This finding points to participants’ ability to transfer knowledge about sub-task abstraction zero-shot to a novel domain. We further observe that abstraction scores improve between the first domain and all subsequent domains. This finding supports our hypothesis that participants improve the level of abstraction of their demonstrations as they gain domain experience.

Our results show that abstraction scores, on average, are monotonically increasing. While the statistically significant improvement in abstraction score occurs after the first domain, the results show a positive trend in subsequent rounds. The diminishing but positive improvement is consistent with prior work finding that human task performance improves logarithmically with practice [45].

Teaching Duration – We find significance with respect to teaching duration and domain experience ($\chi^2(3) = 41.796, p < .001$). We find the significant pairs (Figure 3e) to be between the first and second domain ($p = .014$), the first and third domain ($p < .001$), and the first and fourth

domain ($p < .001$), as well as between the second and fourth domain ($p = .041$). This finding indicates that participants provide demonstrations more efficiently over time.

Sub-task Redundancy– We find a main effect with respect to learning experience and sub-task redundancy ($\chi^2(3) = 8.018, p = .046$), but find no pairwise significance, (Figure 3f). This finding suggests that there may be a trend between domain experience and sub-task redundancy, but more data are needed.

Research Question 2

We investigate the effect of demonstration sub-task abstraction on participants’ perceived workload.

Sub-task Count – We perform a Spearman’s correlation test and find significance between sub-task count and perceived workload ($\rho = -.519, p = .005$). These findings imply that sub-task count is negatively correlated with perceived workload. High sub-task count means breaking up the task into many smaller sub-tasks, each of which can be reused to avoid redundant demonstrations. One possible explanation of this finding is that fewer sub-tasks for a task indicate more repetitive demonstrations.

Research Question 3

We investigate whether participant demographics impact the quality of demonstrations.

Teaching Experience – We find significance between prior teaching experience and sub-task count ($\rho = -.473, p = .011$). This finding is evidence that increased prior teaching

experience is negatively correlated with sub-task count. This gained understanding of the impact of prior teaching experience on sub-task count could be used to improve the existing curriculum designed to teach demonstrators how to provide sufficient demonstrations.

Likeability – We find significance between sub-task count and robot likeability ($\rho = -.501, p = .007$). This finding is evidence that increased robot likeability is negatively correlated with sub-task count.

Agreeableness – Next, we find significance between teaching duration and the agreeableness sub-scale of the Big Five Personality survey ($\rho = .503, p = .006$). This finding is evidence that participant agreeableness is negatively correlated with the efficiency with which they provide demonstrations, namely that more agreeable participants utilize more time to provide demonstrations.

Negative Social Influence – Finally, we find significance between teaching duration and the negative social influence sub-scale of the Negative Attitude towards Robotics survey ($\rho = .577, p = .001$). This finding is evidence that higher teaching duration is correlated to perceptions of negative robot social influence, i.e., participants that are warier of robots take more time to provide demonstrations.

Research Question 4

We now investigate whether domain type impacts the quality of demonstrations.

Teaching Duration – Through a Friedman rank sum test, we find significance in the teaching duration among domains ($\chi^2(3) = 8.656, p = .034$). We find one significant pair between table setting and box packing domains ($p = .036$). We plot domain type against teaching time, as seen in Figure 3b.

Sub-task Redundancy – Through a Friedman rank sum test, we find significance in the redundancy score among domains ($\chi^2(3) = 33.836, p < .001$). We find the significant pairs to be between table setting and medicine dispensing ($p = .006$), table setting and soil mixing ($p < .001$), and box packing and soil mixing ($p = .023$) (Figure 3c).

Sub-task Count – Through a Friedman rank sum test, we find significance in the unique sub-task count among domains ($\chi(3)^2 = 45.87, p < .001$). A Nemenyi-Wilcoxon-Wilcox all-pairs test yields significant pairs for table setting and box packing ($p < .001$), table setting and medicine dispensing ($p = .006$), and table setting and soil mixing ($p < .001$).

Abstraction Score – Finally, we note that we find no significance between the abstraction score and domain, as seen in Figure 3a.

APPENDIX D SURVEYS

First, we provide Cronbach’s alpha scores for the surveys we employ to test for internal reliability.

We next provide details regarding the interview questions and hand-crafted questionnaires we administer.

Scale	α
Agreeableness	.787
Likeability	.928
Negative Situation	.749
Negative Social	.654
Negative Emotions	.766

TABLE I: Cronbach’s alpha.

A. Post-Interview Questions

We ask participants five post-interview questions.

- 1) “Describe your experience using the interface.”
- 2) “Did you feel any information was missing from the explanations provided to you in this study?”
- 3) “What did you think about the videos you watched in this study?”
- 4) “Do you think you were able to teach the robot better as the experiment progressed?”
- 5) “Do you have any other comments or suggestions?”

B. Additional Subjective Metrics

We collect the following subjective metrics in the post-survey questionnaire, but find no significance, and therefore did not describe them in the paper.

- **System Usability:** We use the System Usability Scale (SUS), 10 questions rated on a seven-point scale (Strongly Disagree=1 to Strongly Agree=7) to measure robot and system usability [9].
- **Anthropomorphism:** We use the Anthropomorphism sub-scale of the Godspeed Questionnaire Series, consisting of five questions rated on a five-point scale [7].
- **Methods of Teaching the Robot:** We employ a hand-crafted questionnaire consisting of 17 questions rated on a seven-point scale (Strongly Disagree=1 to Strongly Agree=7) that measured the perceived usefulness and effectiveness of the method of teaching (see Appendix).
- **Trust:** We employ the Multi-Dimensional Measure of Trust (MDMT) questionnaire [52] consisting of 16 questions rated on an eight-point scale (Not at all=0 to Very=7, with an option for Does Not Fit). We measure the overall scale score, as well as participants’ score on the Capacity Trust (with Reliable and Capable sub-scales) and Moral Trust (with Ethical and Sincere sub-scales).

C. Hand-Crafted Questionnaires

1) **Robotics Prior Experience Survey:** Please provide your years of experience with robots (0 to 10+). You must make a selection, even if it is to keep the slider at 0.



Fig. 4: Depicted is the single item robotics prior experience question.

2) *Teacher Prior Experience Survey*: This section of the survey is a 7-point Likert scale ranging from strongly disagree to strongly agree.

- I have experience teaching.
- I have experience mentoring.
- I have experience tutoring.
- I have experience coaching.
- I have experience training others.

3) *Methods of Teaching the Robot*: This section of the survey is a 7-point Likert scale (Strongly Disagree to Strongly Agree).

- 1) Using this method of teaching robots is useful for me
- 2) Using this method of teaching robots will improve my effectiveness
- 3) Using this method of teaching robots will improve my performance
- 4) This method of teaching robots would make it be easy to teach robots behaviors I need
- 5) Learning this method of teaching robots would be easy
- 6) This method of teaching robots would be easy to use
- 7) Using this method to train agents is an idea I like
- 8) Using this method to train agents would be a pleasant experience
- 9) Using this method to train agents is a good idea
- 10) Using this method to train agents is a wise idea
- 11) I trust this method to teach robots effectively
- 12) This method to teach robots is reliable
- 13) This method to teach robots is trustworthy
- 14) I am concerned about how I use my time to train agents
- 15) When I will need it I would like to use this method to train robots
- 16) When I will need it I will intend to use this method to train robots
- 17) When I will need it I predict I would use this method to train robots

APPENDIX E METHODS

A. Videos

Introduction Video and Optimal Teaching Strategy Videos can be found at the following website.

<https://sites.google.com/view/moormanetal-rss2023>

The introduction video introduces the study, the robot, and the interface used to teach the robot sub-tasks. The video then consists of a conceptual description of how to optimally teach the robot to make an omelet. The optimal sub-tasks described for this example included (1) going to the egg carton, (2) picking up an egg, (3) going to the pan, and (4) breaking an egg into the pan. This portion of the video motivates breaking up the task into sub-tasks that can be called many times, to generalize to an omelet of any quantity of eggs. It also suggests recording the “go to the egg container” sub-task separately from the “pick up the egg” sub-task, allowing the robot to generalize going to an egg carton whose location has been moved. Finally, this video communicates that the sub-tasks

can be called from generalized starting positions so multiple sub-tasks can be chained together without going back to a home position first. We note that this initial omelet domain is experienced entirely virtually, and we do not show the participant how it would be taught on the physical robot.

B. Interface Design

The interface for this study relied on usability heuristics and user experience principles. In order to understand what improvements are needed, we analyzed an existing robot manipulation interface [22] against Nielsen Norman’s 10 heuristics principles to determine which design elements could be improved [42].

The design entails a recording screen (Figure 5a) and an “assemble” screen (Figure 5b). The recording screen allows users to facilitate robot teaching and learning on demand, record their actions three times (to optimize robot learning), and save learned tasks into a library. System status is visible via the recording progress bar, which helps notify users of a robot’s overall learning progress. In the assemble screen, the *library* concept allows users to save, categorize, and visualize all of the robot’s learned behavior; it is created to match the user’s mental model of a knowledge repository and allows users to quickly navigate to learned robot commands and be able to call to action quickly. The *queue* concept allows users to pick existing learned tasks from the library, arrange them in a logical sequence, and form coherent, continuous, longer tasks; it is created to maximize personalization for users while accounting for a robot’s technical capabilities.

The design also incorporated a stop button that will halt all robot movement immediately. This feature aligns with the Nielsen Norman Group’s heuristics principles [42] in ensuring enough user freedom and control in an interactive environment. Though the experiment is designed to be in a controlled setting, this design added another user exit option to increase both confidence in robot teaching as well as user safety.

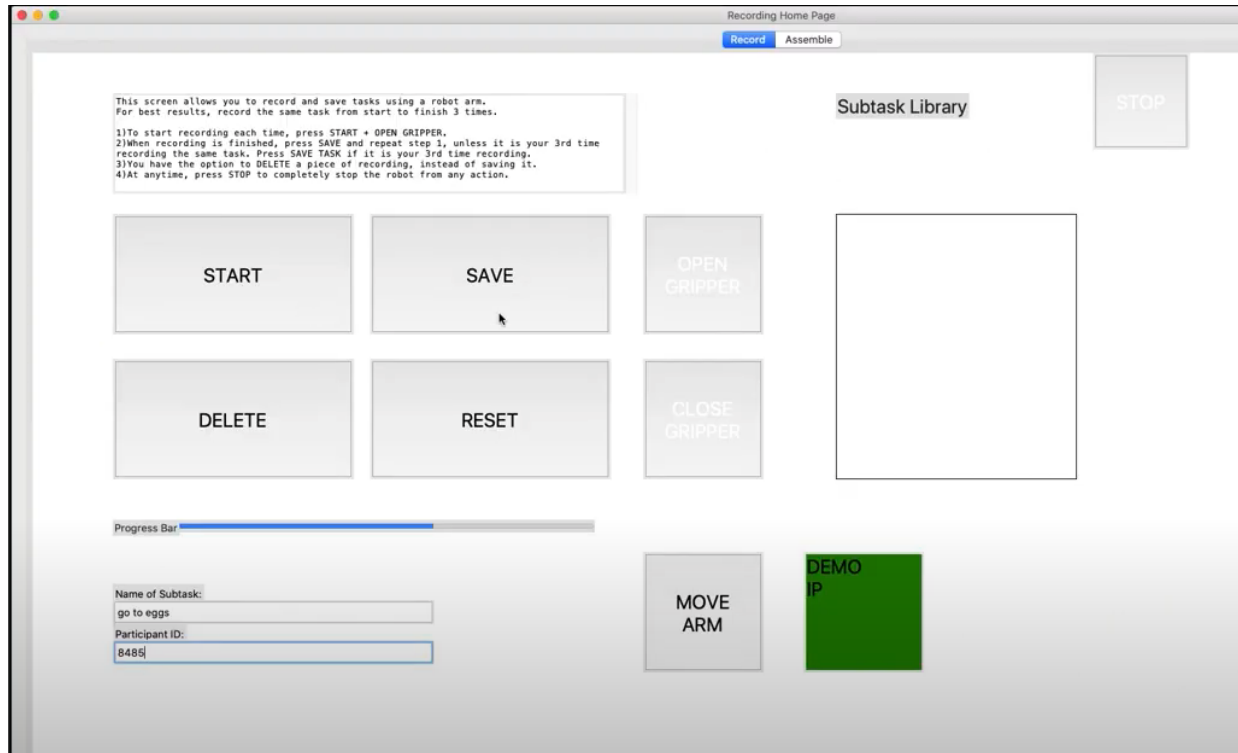
APPENDIX F STUDY DESIGN

A. Domains and Tasks


- **Block Touching:** A blue, green, and red block are placed in front of the robot. Participants are asked to teach the robot to touch the blocks in a particular order using the gripper of the robot. The tasks for the block touching domain are as follows.

- 1) Touch the red block, then touch the blue block.
- 2) Touch the blue block, then touch the green block.
- 3) Touch the green block, then touch the blue block, then touch the red block.

- **Box Packing:** Two plastic bananas, two Jell-O boxes, and two SPAM cans, along with a cardboard box are laid out in front of the robot (these food items are chosen to be in compliance with items commonly used in robotics benchmarking, such as the YCB dataset [53]). Participants are asked to teach the robot to pick up and place a



(a) Depicted is an image of the interface's recording screen.



(b) Depicted is an image of the interface's assemble screen.

Fig. 5: This figure illustrates the screens of the interface.




Fig. 6: Depicted is a flow chart of the record interface tab.




Fig. 7: This figure depicts the five domains in which participants taught the robot.

combination of these food items into the cardboard box. In this domain, bananas served as distractor items, and are not relevant to the tasks the participant must teach the robot to accomplish. The tasks for the box packing domain were as follows.

- 1) Pack 2 SPAM cans and 1 jello box in the box.
 - 2) Pack 1 SPAM can in the box.
 - 3) Pack 2 jello boxes in the box.
- **Table Setting:** Two forks, two knives, and two plates are placed in front of the robot. Participants are asked to teach the robot to pick up the utensils and place them in designated locations around the two plate settings. Duct tape has been wrapped around the utensils in order to facilitate manipulation of the utensils by the robot's gripper. The tasks for the table setting domain were as follows.
 - 1) Place a fork and knife for plate setting 2 and a knife for plate setting 1.
 - 2) Place a fork for plate setting 1.
 - 3) Place a fork and knife for plate setting 1.
 - **Soil Mixing:** A bucket of manure, a bucket of sand, and a bucket of lime are placed in front of the robot, along with a mixing bowl into which scoops of each of these materials are to be poured. Participants are asked to teach the robot to create different soil mixtures for different plants using the scoop, which has been placed in the robot's gripper. The tasks for the soil mixing domain were

as follows.

- 1) Assemble 2 scoops of sand and 1 scoop of lime.
 - 2) Assemble 1 scoop of manure and 1 scoop of lime.
 - 3) Assemble 1 scoop of sand and 1 scoop of manure.
- **Medicine Dispensing:** Four kinds of medicine (red pill cup, green pill cup, yellow pill cup, and TUMS pill cup) along with three trays labeled persons 1, 2, and 3 are placed in front of the robot. Participants are asked to pick and place these medicine cups into the appropriate person's tray to dispense the proper medication to each person. Note that here the TUMS pill cup and person 3 tray both serve as distractor items, and are not relevant to the tasks the participant must teach the robot in this domain. The tasks for the medicine dispensing domain were as follows.
 - 1) Serve the red pills to person 1, and the yellow pills to person 2.
 - 2) Serve the green pills to person 1.
 - 3) Serve the yellow pills to person 1 and the red pills to person 2.

B. Orderings

We ordered the conditions based on a Latin square design to ensure that the ordering is balanced. Each participant experienced one of the conditions.

Condition 1: Table Setting, Soil Mixture, Box Packing, Medicine Dispensing

TABLE II: This table shows the domain orders for each of the four conditions.

Table Ordering	Soil Ordering	Box Ordering	Medicine Ordering
1	2	3	4
3	1	4	2
4	3	2	1
2	4	1	3

Condition 2: Soil Mixture, Medicine Dispensing, Table Setting, Box Packing

Condition 3: Medicine Dispensing, Box Packing, Soil Mixture, Table Setting

Condition 4: Box Packing, Table Setting, Medicine Dispensing, Soil Mixture

```

goToRedPills( $s_i, \tau, s_r, p_t, \dots, p_{p_3}$ ):
  con: Move( $s_i, \tau, s_r$ ), CFreeE( $\tau$ ),
        CFreeTUMS( $p_t, \tau$ ), . . . ,
        CFreePerson3( $p_{p_3}, \tau$ )
  pre: holdingCup=False, LocRobot= $s_i$ ,
        LocTUMS= $p_r, \dots$ , LocPerson3= $p_{p_3}$ 
  eff: LocRobot= $s_r$ , holdingCup=False
goToGreenPills( $s_i, \tau, s_g, p_t, \dots, p_{p_3}$ ):
  con: Move( $s_i, \tau, s_g$ ), CFreeE( $\tau$ ),
        CFreeTUMS( $p_t, \tau$ ), . . . ,
        CFreePerson3( $p_{p_3}, \tau$ )
  pre: holdingCup=False, LocRobot= $s_i$ ,
        LocTUMS= $p_r, \dots$ , LocPerson3= $p_{p_3}$ 
  eff: LocRobot= $s_g$ , holdingCup=False
goToYellowPills( $s_i, \tau, s_y, p_t, \dots, p_{p_3}$ ):
  con: Move( $s_i, \tau, s_y$ ), CFreeE( $\tau$ ),
        CFreeTUMS( $p_t, \tau$ ), . . . ,
        CFreePerson3( $p_{p_3}, \tau$ )
  pre: holdingCup=False, LocRobot= $s_i$ ,
        LocTUMS= $p_r, \dots$ , LocPerson3= $p_{p_3}$ 
  eff: LocRobot= $s_y$ , holdingCup=False
grasp[cup]( $s_i, \tau, s_j, p_c$ ):
  con: Grasp( $s_i, \tau, s_j$ )
  pre: holdingCup=False, LocRobot= $s_i$ ,
        Loc[cup]= $p_c$ 
  eff: holdingCup=True, LocRobot= $s_j$ 
goToPerson1( $s_i, \tau, s_{p_1}, p_t, \dots, p_{p_3}$ ):
  con: Move( $s_i, \tau, s_{p_1}$ ), Upright( $\tau$ ),
        CFreeE( $\tau$ ), CFreeTUMS( $p_t, \tau$ ), . . . ,
        CFreePerson3( $p_{p_3}, \tau$ )
  pre: holdingCup=True, LocRobot= $s_i$ ,
        LocTUMS= $p_r, \dots$ , LocPerson3= $p_{p_3}$ 
  eff: LocRobot= $s_{p_1}$ , holdingCup=True
goToPerson2( $s_i, \tau, s_{p_2}, p_t, \dots, p_{p_3}$ ):
  con: Move( $s_i, \tau, s_{p_2}$ ), Upright( $\tau$ ),
        CFreeE( $\tau$ ), CFreeTUMS( $p_t, \tau$ ), . . . ,
        CFreePerson3( $p_{p_3}, \tau$ )
  pre: holdingCup=True, LocRobot= $s_i$ ,
        LocTUMS= $p_r, \dots$ , LocPerson3= $p_{p_3}$ 
  eff: LocRobot= $s_{p_2}$ , holdingCup=True
release[person]( $s_i, \tau, s_j, p_p$ ):
  con: Release( $s_i, \tau, s_j$ )
  pre: holdingCup=True, LocRobot= $s_i$ ,
        Loc[person]= $p_b$ 
  eff: holdingCup=False, LocRobot= $s_j$ 

```

Fig. 8: Sample formulation in the medicine domain.

APPENDIX G TAMP SUB-TASK FORMULATION

Task and Motion Planning

Task and motion planning (TAMP) problems integrate multi-modal motion planning with representational strategies originating in long horizon task planning [20, 54]. TAMP is a framework that allows for an agent to break down and complete long-horizon tasks while taking into account multi-faceted constraints in a similar manner to Multi-Modal Motion Planning. As such, each sub-task is described by its preconditions, `pre`, that need to be set in order for a sub-task to take place, constraints `con` that is the set of constraint functions that must hold for all continuous actions during the duration of the sub-task, and effects (or transition), `eff`, the changes made to the robot, objects and environments after a sub-task has been executed.

Sample Formulation

Next, we provide a sample formulation for the Medicine Dispensing domain, shown in Figure 8. Each task takes in an initial state s_i , a trajectory provided by a demonstration τ , the final state the robot should reach at the end of the task s_x (x varies depending on the task), and a collection of poses of all objects in the environment.

Variable τ represents a single trajectory. There are few different types of constraints included in this formulation. $\text{Move}(s_i, \tau, s_j) \forall s_i, s_j \in \mathcal{S}$ is a constraint over trajectories where the robot moves from s_i to s_j . The constraints for $\text{CloseGripper}(s_i, \tau, s_j)$ and $\text{OpenGripper}(s_i, \tau, s_j)$ are similar, they represent trajectories that allow the robot to grasp and release the cup respectively. There are also collision-free constraints, such as `CFreeE` and `CFreeTUMS` that tell the robot not to collide with the environment or any of the objects in the environment. The go to person sub-tasks also include an `Upright` constraint telling the robot to keep the cup upright when holding it.

In the go to pills sub-tasks the robot is not holding anything in its gripper, i.e., `holdingCup=False`. As explained above, this means that its range of motion is not nearly as limited as the robot does not have to worry about any object that it may be manipulating. The go to person sub-tasks are transfer mode subtasks [20], as `holdingCup=True`, and so the robot’s range of motion is limited as it has to worry about not spilling the pills while moving the cup; hence `Upright` is included in `con`. `grasp` performs a kinematic switch when grasping a cup; this switch can be seen in the sub-task formulation as `holdingCup=False` in `pre` being switched to `holdingCup=True` in `eff`. The reverse can be seen in `release`.

APPENDIX H QUOTES

Getting better over time

We first wanted to study how domain experience impacts the quality of demonstrations. In our post-interview, we found that

the majority of participants agreed that they were able to teach the robot better as the experiment progressed. Two sample quotes supporting this qualitative finding are as follows.

“I learned from both my experience teaching the robot, and watching the videos and seeing you teach the robot how to do things.”

“After watching the video of how you accomplished the overall work, then I understood what I should do.”

Zero-shot transfer

We found that participants can perform zero-shot transfer of knowledge regarding sub-task abstraction, teaching efficiency, and sub-task redundancy to novel domains. One participant in particular commented on this in the transition between the box packing and table setting domains. They describe their gained understanding about abstracting sub-tasks to account for item multiplicity.

“I followed again the example of the [table setting domain] video so I knew that the two cans of SPAM and two cans of Jell-O were basically equivalent so I only trained it to go to the space and then separately trained it to close the gripper to pick up an item separately...”

Hierarchical Task Planning

In our interviews, we found that participants spoke about low-level motion control in the first couple domains. One participant, when asked whether they thought they taught the robot better as the experiment progressed, stated the following.

“Earlier, there was a lot of joints and the ways the robot could move that I didn’t know of. Or many things that I felt it couldn’t do but it could have that I got to learn while watching the videos. So yeah, as the experiment progressed, I felt pretty comfortable and easier to control the robot.”

In the latter domains, however, participants shifted to discussing higher-level strategies and task planning. One participant, when initially asked about their experience after block touching (the first domain), describes low-level strategies for manipulating the robot.

“I just noticed that it had some joints that could be moved to make it to teach it to do what I wanted so I just moved those and like pointed to the different blocks and then saved as I went.”

The same participant, after their final domain in the Box packing domain, describes their high-level task planning.

“So for this tasks I simply act out the robot how to grab either Jell-O, bananas, or SPAM and then I taught it how to go to the box and then I taught it how to grab and drop so that you would grab either banana, Jell-O, and SPAM and then go to the box and then drop it there.”

The contrast between thinking about the low-level movement in the first demo domain and the higher-level abstractions in

the last domain encountered is especially apparent with these answers.

Workload

We also investigate the effect of demonstration abstraction on participants’ perceived workload. Our findings indicated an inverse correlation between sub-task count and perceived workload. One participant noted fatigue and a high perceived workload at the end of the study.

“Overall, it was very time-consuming, repeating the motions over and over again. I found myself getting really tired near the end, like the last [domain].”

APPENDIX I CODER INSTRUCTIONS


Following are the coder instructions from which we obtain our abstraction and redundancy scores, cited as [22] in our work.

Coder Instructions

Context

Learning from Demonstration enables end-users, who are not robotics experts, to shape robot behavior. In our work, we are interested in how the quality of demonstrations changes over time, after trial and error teaching the robot in different domains.

Domains



Block Touching: A blue, green, and red block are placed in front of the robot. Participants are asked to teach the robot to touch the blocks in a particular order using the robot gripper.

Box Packing: Two plastic bananas, two Jell-O boxes, and two Spam cans, along with a cardboard box are laid out in front of the robot. Participants are asked to teach the robot to pack (pick up and place) a combination of these food items into the cardboard box.

Table Setting: Two forks, two knives, and two plates are placed in front of the robot. Participants are asked to teach the robot to set the table by picking up the utensils and placing them in designated locations around the two plate settings.

Soil Mixing: A bucket of manure, a bucket of sand, and a bucket of lime are placed in front of the robot, along with a mixing bowl into which scoops of each of these materials are to be poured. Participants are asked to teach the robot to create different soil mixtures for different plants. In this domain, the scoop is placed in the robot's gripper by the experimenter.

Medicine Dispensing: Four kinds of medicine (red pill cup, green pill cup, yellow pill cup, and TUMS pill cup) along with three trays labeled persons 1, 2, and 3 are placed in front of the robot. Participants are asked to dispense the proper medication to each person by picking and placing medicine cups into the appropriate person's tray.

Definitions

Domain: The participant will work with the robot in 5 domains, as listed above.

Task: In each domain, the participant must teach the robot 3 tasks (like assemble exactly two scoops of sand and one scoop of lime).

Subtask: The participant teaches the robot sub-tasks (like go to sand) that can be used to accomplish tasks.

Demonstration: The participant physically moved the robot, *demonstrating* to record how to accomplish a subtask.

What data do we have?

In each domain, participants would provide a series of demonstrations to teach subtasks that can be reused to accomplish tasks. They then used these subtasks to accomplish 3 tasks in each domain. We have notes that describe the subtasks provided, and we also have the list of subtasks names assigned for each task. For instance, in the soil mixing domain, a participant could record the following subtasks:

- Go to sand bin
- Go to manure bin
- Go to lime bin
- Scoop
- Go to mixing bucket
- Drop

Then, they would assign these subtasks to the following tasks:

Soil Mixing Domain	Task 1	Task 2	Task 3
Task description	Assemble exactly two scoops of sand and one scoop of lime.	Assemble exactly one scoop of manure and one scoop of lime.	Assemble exactly one scoop of sand and one scoop of manure
Subtasks assigned to task	<ul style="list-style-type: none"> * Go to sand bin * Scoop * Go to mixture bin * Dump * Go to sand bin * Scoop * Go to mixture bin * Dump * Go to lime bin * scoop * Go to mixture bin * Dump 	<ul style="list-style-type: none"> * Go to manure bin * Scoop * Go to mixture bin * Dump * Go to lime bin * Scoop * Go to mixture bin * Dump 	<ul style="list-style-type: none"> * Go to sand bin * Scoop * Go to mixture bin * Dump * Go to manure bin * Scoop * Go to mixture bin * Dump

Types of scoring

Given the notes of each subtask recorded by participants, how do we evaluate the quality of the demonstrations used for each task? We look at two metrics: abstraction score and redundancy score.

Abstraction Score

The abstraction score describes how well the subtasks were abstracted, to maximize the usability of the subtasks. Recall the goal in the soil mixing domain of making soil mixtures for plants. If, for instance, the goal was 2 scoops of lime and 1 scoop of sand.

Insufficiently Abstracted	Optimal Abstraction	Overly Abstracted
<ul style="list-style-type: none">– Go to lime, scoop, go to bucket, drop, go to lime, scoop, go to bucket, drop, go to sand, scoop, go to bucket, drop	<ul style="list-style-type: none">– Go to sand– Go to lime– Scoop– Go to bucket– Drop scoop	<ul style="list-style-type: none">– Move left– Move right– Move up– Move down
The problem here is that we can only reuse this subtask in the exact same scenario (this subtask cannot be used to scoop 1 scoop of lime and 2 scoops of sand).	This optimally abstracted set of tasks will generalize to any number of scoops of lime and sand required.	The problem here is that it is unclear how to combine these tasks to accomplish the goal, since they are too abstracted.

To calculate abstraction score, one point is awarded for each subtask that can be used towards the task when using the sub-tasks to accomplish the task. For example, the insufficiently abstracted example above would receive 1 point as their abstraction score as they only recorded one subtask that could be applied to the task. If they had instead recorded two subtasks:

- Go to lime, scoop, go to bucket, drop, go to lime, scoop, go to bucket, drop
- Go to sand, scoop, go to bucket, drop

Then they would receive an abstraction score of 2.

The optimal abstraction score example above would be used as follows to accomplish the task:

- Go to lime
- Scoop
- Go to bucket
- Drop scoop
- Go to lime
- Scoop
- Go to bucket
- Drop scoop
- Go to sand
- Scoop
- Go to bucket
- Drop scoop

Resulting in an abstraction score of 12.

Redundancy Score

The redundancy score describes how (un)necessary tasks are. A sub-task is deemed redundant if its goal can be met by another sub-task or a combination of sub-tasks previously taught. Sub-task redundancy is defined with respect to a given set of subtasks being taught. For instance, for the same context, where the goal is 2 scoops of lime and 1 scoop of sand, the following table describes two levels of redundancy.

Redundant	Optimal
Go to sand. Scoop sand. Go to lime. Scoop lime. Go to bucket. Drop lime. Drop Sand.	Go to sand. Go to lime. Scoop. Go to bucket. Drop scoop.
The problem here is that we can only reuse drop sand for drop lime (simply drop), so drop lime is unnecessary, and redundant. Similarly, we can use scoop sand for scoop lime (simply scoop), so scoop lime is unnecessary and redundant.	This is the proper subtask breakdown, where the demonstrator recognizes that drop can work for sand, lime, and manure, and scoop similarly can work for sand, lime and manure.

To calculate redundancy score, one point is given for each subtask that can be done using other sub-tasks recorded in the domain. For instance, the redundant example above receives a redundancy score of 2 (one for the redundant drop subtask and one for the redundant scoop subtask). However, the optimal example above receives a redundant score of 0 as no redundant subtasks were taught.

Note that the redundancy score is computed based upon the subtasks taught by that participant, not the optimal subtask list.

Key for each domain

Soil Mixing Domain	Task 1	Task 2	Task 3
Task description	Assemble exactly two scoops of sand and one scoop of lime.	Assemble exactly one scoop of manure and one scoop of lime.	Assemble exactly one scoop of sand and one scoop of manure
Optimal Subtask List	<ul style="list-style-type: none"> * Go to sand bin * Go to manure bin * Go to lime bin * Scoop * Go to mixture bin * Dump 		
Optimal Subtask Assignment	<ul style="list-style-type: none"> * Go to sand bin * Scoop * Go to mixture bin * Dump * Go to sand bin * Scoop * Go to mixture bin * Dump * Go to lime bin * scoop * Go to mixture bin * Dump 	<ul style="list-style-type: none"> * Go to manure bin * Scoop * Go to mixture bin * Dump * Go to lime bin * Scoop * Go to mixture bin * Dump 	<ul style="list-style-type: none"> * Go to sand bin * Scoop * Go to mixture bin * Dump * Go to manure bin * Scoop * Go to mixture bin * Dump
Max Abstraction Score	12	8	8

Block Touching (Demo) Domain	Task 1	Task 2	Task 3
Task description	Touch the red block then touch the blue block	Touch the blue block then touch the green block	Touch the green block then touch the blue block then touch the red block
Optimal Subtask List	<ul style="list-style-type: none"> * Go to and touch the green block * Go to and touch the blue block * Go to and touch the red block 		
Optimal Subtask Assignment	<ul style="list-style-type: none"> * Go to and touch the red block * Go to and touch the blue block 	<ul style="list-style-type: none"> * Go to and touch the blue block * Go to and touch the green block 	<ul style="list-style-type: none"> * Go to and touch the green block * Go to and touch the blue block * Go to and touch the red block
Max Abstraction Score	2	2	3

Box Packing Domain	Task 1	Task 2	Task 3
Task description	Pack exactly two SPAM cans and exactly one jello box in the box	Pack exactly one SPAM can in the box	Pack exactly two jello boxes in the box
Optimal Subtask List	<ul style="list-style-type: none"> * Go to SPAM * Go to jello box * Close gripper * Go to box * Open gripper 		
Optimal Subtask Assignment	<ul style="list-style-type: none"> * Go to SPAM * Close gripper * Go to box * Open gripper * Go to SPAM * Close gripper * Go to box * Open gripper * Go to jello box * Close gripper * Go to box * Open gripper 	<ul style="list-style-type: none"> * Go to SPAM * Close gripper * Go to box * Open gripper 	<ul style="list-style-type: none"> * Go to jello box * Close gripper * Go to box * Open gripper * Go to jello box * Close gripper * Go to box * Open gripper
Max Abstraction Score	12	4	8

Table Setting Domain	Task 1	Task 2	Task 3
Task description	Place the fork and knife for plate setting 2 and knife for plate setting 1	Place fork for plate setting 1	Place knife and fork for plate setting 1
Optimal Subtask List	<ul style="list-style-type: none"> * Go to fork container * Go to knife container * Close gripper * Go to plate setting 1 * Go to plate setting 2 * Go to plate's fork loc * Go to plate's knife loc * Open gripper 		
Optimal Subtask Assignment	<ul style="list-style-type: none"> * Go to fork container * Close gripper * Go to plate setting 2 * Go to plate's fork loc * Open gripper * Go to knife container * Close gripper * Go to plate setting 2 * Go to plate's knife loc * Open gripper * Go to knife container * Close gripper * Go to plate setting 1 * Go to plate's knife loc * Open gripper 	<ul style="list-style-type: none"> * Go to fork container * Close gripper * Go to plate setting 1 * Go to plate's fork loc * Open gripper 	<ul style="list-style-type: none"> * Go to fork container * Close gripper * Go to plate setting 1 * Go to plate's fork loc * Open gripper * Go to knife container * Close gripper * Go to plate setting 1 * Go to plate's knife loc * Open gripper
Max Abstraction Score	15	5	10

Medicine Dispensing Domain	Task 1	Task 2	Task 3
Task description	Serve red pills to person 1, and yellow pills to person 2	Serve green pills to person 1	Serve yellow pills to person 1 and red pills to person 2
Optimal Subtask List	<ul style="list-style-type: none"> * Go to yellow pill cup * Go to red pill cup * Go to green pill cup * Close gripper * Go to person 1 location * Go to person 2 location * Open gripper 		
Optimal Subtask Assignment	<ul style="list-style-type: none"> * Go to red pill cup * Close gripper * Go to person 1 location * Open gripper * Go to yellow pill cup * Close gripper * Go to person 2 location * Open gripper 	<ul style="list-style-type: none"> * Go to green pills cup * Close gripper * Go to person 1 location * Open gripper 	<ul style="list-style-type: none"> * Go to yellow pill cup * Close gripper * Go to person 1 location * Open gripper * Go to red pill cup * Close gripper * Go to person 2 location * Open gripper
Max Abstraction Score	8	4	8

How is the data organized?

Available is a list describing each of the subtasks recorded for that domain. These are notes taken by the experimenter. Available is also the recipe constructed by the participant for each task (task 1 – 3) in that domain. This uses the names the participant called each of their subtasks (the naming scheme may differ between people).

For each person, the coder must read the tasks recorded, understand how they differ from the optimal key provided, and look at how the participant used these subtasks for each task in the domain. Then, they score the abstraction score and redundancy of that domain, by summing the abstraction score across tasks of the domain, and summing redundancy score across tasks of the domain.

The written notes are to provide context into the observed behavior of the recording of subtasks. However, it is possible that participants recorded subtasks that they never used, i.e. in the recipe for the task (if they for instance change their strategy half way through). Therefore, when calculating score please verify that the subtask was used in the recipe saved.

More Information

This section details background in Task and Motion Planning (TAMP) to enable us to answer the question: *How to tell if a subtask is valid/ can be applied to the task?*

Pre-requisites

- 1) an understanding of TAMP (<https://arxiv.org/abs/2010.01083>)
- 2) familiarity with the domain that you are rating. This familiarity can be gained by actually solving the domain and understanding the possible abstractions within the domain to solve the given tasks.

Concepts needed to score –

There are two concepts that a rater of TAMP abstractions needs to understand: Mode Changes and Bottleneck states.

Mode changes: When we interact with objects, we change the mode of contact with them. Consider picking up a fork. When we have no fork in the gripper the gripper is bound by fewer constraints of movement. When the fork is in the gripper the gripper is bound by additional constraints such as the fork not colliding with anything else in the world. This change in constraints of movement lead to obvious points of sub-task creation. Pick up fork needs to be a different sub-task from move fork to the plate.

Bottleneck States: Bottle neck states are states that an agent needs to reach to solve a larger family of sub-tasks. A classic example of a bottleneck state is a door. To access the gas stove or the kitchen cabinet or to cook, you first need to cross the bottleneck state of the kitchen door. Across these bottlenecks the agent has access to a greater number of sub-tasks. In pick and place domains a classic bottle neck is the “pre-grasp” position. Consider the location just above the fork, this allows the robot to pick up the fork in different orientations. Hence, this location just above the fork is considered a pre-grasp position. The robot does not have to constrain its gripper pose when moving to the pre-grasp position, but when initiating the “grasp” sub-task the robot’s gripper needs to be strictly oriented to pick

the object correctly. A similar situation occurs when the robot needs to place knives and forks next to a plate. The move to plate is a good sub-task to teach as all utensils would require the robot to move to the plate. After which the robot can move to the knife's location and place the knife or move to the fork's position to place the fork. Notice that for a different plate we would need to teach a novel reach plate task as the plates are at different locations, however the knives and forks are at the same relative position when compared to the plate's position, making the move to fork or knife position sub-tasks repeatable when taught once. Without teaching this bottleneck abstraction of move to plate, the user would have to teach place knife next to plate 1 and place knife next to plate 2 as different tasks, and repeat these sub-tasks for the forks, the napkins, the spoons, etc, for each plate. Now they just have to teach how to reach the bottleneck state of reaching a plate, and then the robot can just repeat the placement of the knives and forks from previous relative sub-tasks.

How to award points –

Abstraction score - We give points for each sub-task the user taught that is necessary in solving the task. The sub-tasks are necessary if without them the overall task could not have been solved. For example, to place a fork from the fork location next to plate 1 we can break down subtasks in many ways. A user might just pick the fork and place it next to plate 1 in one task. This is a necessary sub-task to solve this problem, hence the user gets one point. Or the user might decide to move to the fork location (pre-grasp), pick up the fork (grasp), move to the plate 1 (bottleneck), move to fork location (pre-place), place fork (place). These provide 5 sub-tasks all necessary and repeatable hence the user gets 5 points. The users might not distinguish between a pre-grasp and a grasp or a pre-place and place and they might only get 3 points (for pick, move to bottleneck and place).

Redundancy score - Sometimes users teach sub-tasks but do not use them, or use them but they were unnecessary to begin with. For example, a lot of users teach a sub-task to move the robot back to a home position before the robot does the next sub-task. This is not necessary. The users do it because of a poor understanding of how the robot works and this is penalized when measuring the redundancy score.

What about over-abstracting? Are we providing too many points for users to create too many sub-tasks? Does moving 1 cm to the left a sub-task – no, there is no mode change or reaching of a bottleneck state because of this sub-task, hence it is not necessary in solving the task. However, you can teach two ways of getting to the fork, is that important? What if I make a different sub-task to pick every fork in the fork location, now I have tons of sub-tasks. Well picking up one fork from a fork basket is no different from picking another so just wasted your time, and this is a redundant sub-task!

How the scoring strategy applies to each domain –

Now for each domain we will provide our scoring key and list of bottleneck and mode changes. These should be sufficient to score both the abstraction scores and redundancy scores.

Soil domain -

The mode changes here happen when contacting the objects being picked with the scoop. The bottleneck states are the pre-grasp positions over the positions of the objects and the mixing bowl. The optimal abstraction moves over objects (bottleneck state / pre-grasp), scoops them (grasp), moves to the mixing bowl (bottleneck state / pre-place), dumps contents (place).

Table setting domain -

The mode changes happen when forks or knives are picked up (grasp) or placed (released). The bottle neck states are the pre-pick and pre-place over the knife and fork locations and the move to plate, move to fork pick location (which is a fork dispenser like a drawer partition with forks alone) and a move to knife pick location (again a knife dispenser).

Medicine delivery domain -

The mode changes here are pick and place of the medicine boxes. The bottle neck states are to get to the position of each of the colored medicine pill boxes and person's location for delivery of the medicines as pre-grasp and pre-place positions.

Box packing domain -

The mode changes are in the pick and place of the items and are taught as close gripper and open gripper. The bottle-neck abstractions are to reach the location of the spam containers, and jello-box containers storage locations, and reach the location of the box being packed. Again, the spam and jello boxes are present at the same location (like a shelf) and do not need the robot to be taught different tasks to reach them.

Block touching -

Here the only mode change is when the blocks are touched. There is no need of a bottleneck state as there are no constraints over the pose of the arm when touching the blocks.

ACKNOWLEDGMENTS

This work was supported by a grant from the National Science Foundation (IIS-2112633) and a gift from Konica Minolta, Inc.

REFERENCES

- [1] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1, 2004.
- [2] Bernardo Aceituno and Alberto Rodriguez. A hierarchical framework for long horizon planning of object-contact trajectories. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 189–196. IEEE, 2022.
- [3] Baris Akgun, Maya Cakmak, Karl Jiang, and Andrea L Thomaz. Keyframe-based learning from demonstration. *International Journal of Social Robotics*, 4(4):343–355, 2012.
- [4] Rachid Alami, Thierry Siméon, and Jean-Paul Laumond. A geometrical approach to planning manipulation tasks. the case of discrete placements and grasps. In *International Symposium on Robotics Research*, 1991.
- [5] Rachid Alami, Jean-Paul Laumond, and Thierry Siméon. Two manipulation planning algorithms. In *Proceedings of the Workshop on Algorithmic Foundations of Robotics*, 1995.
- [6] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- [7] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International journal of social robotics*, 1(1):71–81, 2009.
- [8] Aude Billard, Sylvain Calinon, Ruediger Dillmann, and Stefan Schaal. Robot programming by demonstration. In *Springer handbook of robotics*, pages 1371–1394. Springer, 2008.
- [9] John Brooke. System usability scale (sus): a quick-and-dirty method of system evaluation user information. *Reading, UK: Digital equipment co ltd*, 43:1–7, 1986.
- [10] Daniel S. Brown, Wonjoon Goo, and Scott Niekum. Better-than-demonstrator imitation learning via automatically-ranked demonstrations. In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura, editors, *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, pages 330–359. PMLR, 30 Oct–01 Nov 2020. URL <https://proceedings.mlr.press/v100/brown20a.html>.
- [11] M. Cakmak and L. Takayama. Teaching people how to teach robots: The effect of instructional materials and dialog design. *2014 9th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 431–438, 2014.
- [12] Alexandre Campeau-Lecours, Véronique Maheu, Sébastien Lepage, Hugo Lamontagne, Simon Latour, Laurie Paquet, and Neil Hardie. Jaco assistive robotic device: Empowering people with disabilities through innovative algorithms. In *Rehabilitation Engineering and Assistive Technology Society of North America (RESNA) Annual Conference*, 2016.
- [13] Rich Caruana. *Multitask learning*. Springer, 1998.
- [14] 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.
- [15] Sonia Chernova and Andrea L Thomaz. *Robot Learning from Human Demonstration*. Springer, 2014.
- [16] Luis C. Cobo, Charles L. Isbell, and Andrea L. Thomaz. Automatic task decomposition and state abstraction from demonstration. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1, AAMAS '12*, page 483–490, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 0981738117.
- [17] Tesca Fitzgerald and Andrea L Thomaz. Skill demonstration transfer for learning from demonstration. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts*, pages 187–188, 2015.
- [18] Tesca Fitzgerald, Ashok Goel, and Andrea Thomaz. Abstraction in data-sparse task transfer. *Artificial Intelligence*, 300:103551, 2021.
- [19] Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2017.
- [20] Caelan Reed Garrett, Rohan Chitnis, Rachel Holladay, Beomjoon Kim, Tom Silver, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Integrated task and motion planning. *Annual review of control, robotics, and autonomous systems*, 4:265–293, 2021.
- [21] Lewis R Goldberg. The development of markers for the big-five factor structure. *Psychological assessment*, 4(1): 26, 1992.
- [22] Nakul Gopalan, Nina Moorman, Manisha Natarajan, and Matthew Gombolay. Negative result for learning from demonstration: Challenges for end-users teaching robots with task and motion planning abstractions.
- [23] Siddhant Haldar, Vaibhav Mathur, Denis Yarats, and Lerrel Pinto. Watch and match: Supercharging imitation with regularized optimal transport. *CoRL*, 2022.
- [24] Sandra G Hart and Lowell E Staveland. Development of nasa-tlx (task load index): Results of empirical and theoretical research. In *Advances in psychology*, volume 52, pages 139–183. Elsevier, 1988.
- [25] Valentin N Hartmann, Ozgur S Oguz, Danny Driess, Marc Toussaint, and Achim Menges. Robust task and motion planning for long-horizon architectural construction planning. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages

- 6886–6893. IEEE, 2020.
- [26] Kris Hauser and Jean-Claude Latombe. Multi-modal motion planning in non-expansive spaces. *The International Journal of Robotics Research*, 29(7), 2010.
- [27] Kris Hauser and Victor Ng-Thow-Hing. Randomized multi-modal motion planning for a humanoid robot manipulation task. *The International Journal of Robotics Research*, 30(6), 2011.
- [28] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29, 2016.
- [29] Auke Jan Ijspeert, Jun Nakanishi, Heiko Hoffmann, Peter Pastor, and Stefan Schaal. Dynamical movement primitives: learning attractor models for motor behaviors. *Neural computation*, 25(2):328–373, 2013.
- [30] Leslie Pack Kaelbling and Tomás Lozano-Pérez. Hierarchical planning in the now. In *Workshops at the Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.
- [31] Wonchul Kim, Chungkeun Lee, and H. Jin Kim. Learning and generalization of dynamic movement primitives by hierarchical deep reinforcement learning from demonstration. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3117–3123, 2018. doi: 10.1109/IROS.2018.8594476.
- [32] Moritz Knaust and Dorothea Koert. Guided robot skill learning: A user-study on learning probabilistic movement primitives with non-experts. In *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*, pages 514–521. IEEE, 2021.
- [33] W. B. Knox and P. Stone. Interactively shaping agents via human reinforcement: the tamer framework. In *K-CAP '09*, 2009.
- [34] George Konidaris and Andrew Barto. Skill discovery in continuous reinforcement learning domains using skill chaining. *Advances in neural information processing systems*, 22:1015–1023, 2009.
- [35] George Konidaris, Scott Kuindersma, Roderic Grupen, and Andrew Barto. Robot learning from demonstration by constructing skill trees. *The International Journal of Robotics Research*, 31(3):360–375, 2012. doi: 10.1177/0278364911428653.
- [36] Sanjay Krishnan, Animesh Garg, Richard Liaw, Lauren Miller, Florian T Pokorny, and Ken Goldberg. Hirl: Hierarchical inverse reinforcement learning for long-horizon tasks with delayed rewards. *arXiv preprint arXiv:1604.06508*, 2016.
- [37] Clemente Laurettil, Francesca Cordella, Eugenio Guglielmelli, and Loredana Zollo. Learning by demonstration for planning activities of daily living in rehabilitation and assistive robotics. *IEEE Robotics and Automation Letters*, 2(3):1375–1382, 2017.
- [38] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- [39] Boyao Li, Jiayi Li, Tao Lu, Yinghao Cai, and Shuo Wang. Hierarchical learning from demonstrations for long-horizon tasks. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4545–4551. IEEE, 2021.
- [40] Anahita Mohseni-Kabir, Charles Rich, Sonia Chernova, Candace L Sidner, and Daniel Miller. Interactive hierarchical task learning from a single demonstration. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 205–212, 2015.
- [41] Vivek Myers, Erdem Biyik, Nima Anari, and Dorsa Sadigh. Learning multimodal rewards from rankings. In *Conference on Robot Learning*, 2021.
- [42] Jakob Nielsen. *Usability engineering*. Morgan Kaufmann, 1994.
- [43] Kai Ploeger, Michael Lutter, and Jan Peters. High acceleration reinforcement learning for real-world juggling with binary rewards. In *Conference on Robot Learning*, pages 642–653. PMLR, 2021.
- [44] Harish Ravichandar, Athanasios S Polydoros, Sonia Chernova, and Aude Billard. Recent advances in robot learning from demonstration. *Annual review of control, robotics, and autonomous systems*, 3:297–330, 2020.
- [45] Frank Ritter and Lael Schooler. The learning curve. In *International encyclopedia of the social and behavioral sciences.*, 2001.
- [46] Stefan Schaal. Is imitation learning the route to humanoid robots? *Trends in cognitive sciences*, 3(6):233–242, 1999.
- [47] Mariah L Schrum, Erin Hedlund, and Matthew C Gombolay. Improving robot-centric learning from demonstration via personalized embeddings. *arXiv preprint arXiv:2110.03134*, 2021.
- [48] Mariah L. Schrum, Muyleng Ghuy, Erin Hedlund-Botti, Manisha Natarajan, Michael J. Johnson, and Matthew C. Gombolay. Concerning trends in likert scale usage in human-robot interaction: Towards improving best practices. *J. Hum.-Robot Interact.*, nov 2022. doi: 10.1145/3572784. URL <https://doi.org/10.1145/3572784>.
- [49] Mariah L Schrum, Erin Hedlund-Botti, and Matthew Gombolay. Reciprocal mind meld: Improving learning from demonstration via personalized, reciprocal teaching. In *6th Annual Conference on Robot Learning*, 2022.
- [50] Mariah L Schrum, Erin Hedlund-Botti, Nina Moorman, and Matthew C Gombolay. Mind meld: Personalized meta-learning for robot-centric imitation learning. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 157–165. IEEE, 2022.
- [51] Dag Sverre Syrdal, Kerstin Dautenhahn, Kheng Lee Koay, and Michael L Walters. The negative attitudes towards robots scale and reactions to robot behaviour in a live human-robot interaction study. *Adaptive and emergent behaviour and complex systems*, 2009.
- [52] Daniel Ullman and Bertram F Malle. What does it mean to trust a robot? steps toward a multidimensional measure of trust. In *Companion of the 2018 acm/ieee international*

conference on human-robot interaction, pages 263–264, 2018.

- [53] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017.
- [54] Kai Zhang, Eric Lucet, Julien Alexandre Dit Sandretto, Selma Kchir, and David Filliat. Task and motion planning methods: applications and limitations. In *19th International Conference on Informatics in Control, Automation and Robotics ICINCO 2022*, pages 476–483. SCITEPRESS-Science and Technology Publications, 2022.
- [55] Yu Zhang and Qiang Yang. An overview of multi-task learning. *National Science Review*, 5(1):30–43, 2018.
- [56] Zuyuan Zhu and Huosheng Hu. Robot learning from demonstration in robotic assembly: A survey. *Robotics*, 7(2):17, 2018.
- [57] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.