

Quantifying Demonstration Quality for Robot Learning and Generalization

Maram Sakr , Zexi Jesse Li , H. F. Machiel Van der Loos, Dana Kulić , and Elizabeth A. Croft 

Abstract—Learning from Demonstration (LfD) seeks to democratize robotics by enabling non-expert end-users to teach robots. However, most LfD techniques assume users provide optimal demonstrations, which may not be accurate. Demonstration quality plays a crucial role in robot learning and generalization. Hence, it is important to quantify the quality of the provided demonstrations before using them for robot learning. In this letter, we propose quantifying the generalizability of the demonstrations based on how well they perform in the learned task. The proposed approach is validated in a user study ($N = 27$). Participants with different robotics expertise levels were recruited to teach a PR2 robot a generic task (pressing a button) under different task constraints. They taught the robot in two sessions on two different days to capture their teaching behaviour across sessions. The task performance was utilized to classify the provided demonstrations into high-quality and low-quality sets. The results show a significant correlation between task performance and generalization performance across all participants. We also found that users clustered into two groups: Users who provided high-quality demonstrations from the first session (the *fast-adapters*), and users who provided low-quality demonstrations in the first session and then improved with practice (the *slow-adapters*). This approach for assessing demonstrations allows us to determine whether users require more training in order to provide high-quality demonstrations.

Index Terms—Learning from demonstration, physical human-robot interaction, design and human factors.

I. INTRODUCTION

AS ROBOTS enter the human environment to assist people in their daily lives, the ability of everyday users, who do not have any robotics or programming background, to work with these robots will soon be essential. A common approach to allowing ordinary people to program a robot is Learning from

Demonstration (LfD). LfD is a paradigm that allows robots to perform tasks after observing a teacher performing them without explicit programming [1], [2].

One of the key elements for deploying LfD in the real world is ensuring fast learning and generalization. This means that the robot has the ability to learn the taught task quickly and apply the learned knowledge to new circumstances. It is not practical to reprogram the robot for every small change in the task and/or environment. To design a generalizable learning system, we need to consider both the human teacher and the robot learner as active stakeholders who have a role in improving efficiency [3]. Most of the literature to date has focused on the robot side, that is, on advancing techniques to boost robot learning efficiency [4]. Much less attention has been paid to the human teacher's role in this process.

As noted in [1], a robot's performance depends heavily on the quality of the demonstrations provided. Hence, it is crucial to define and measure demonstrations quality when teaching a robot. There are different factors that affect demonstrations quality, here we focus on the provider of the demonstrations. The human teacher's adaptation to the task may be varying and is related to their previous expertise. Kobak and Mehring [5] highlighted that after practice with motor tasks sharing structural similarities, new tasks of the same type can be learned faster. Thus, users' stated robotics expertise may not be sufficient to judge their provided demonstrations quality or their adaptation level to a specific task [6], [7]. This necessitates the need for a quantitative approach to quantifying demonstrations quality. In addition, it is important to identify the adaptation level of the users to the taught task to know how much time is needed for users' practice with the robot and when the provided demonstrations can be used for robot learning [8].

The term "poor quality demonstrations" has been used loosely in the literature to refer to different issues in the data itself. These issues include undesired motions [4], failed demonstrations [9], and ambiguous demonstrations [10]. However, the explicit definition and measurements for the *demonstrations quality* is still an open question [4].

In this paper, we define the quality of the provided demonstrations based on how well the learning algorithm reproduces the task, regardless of the demonstrator's stated expertise level. This will be done by assessing the learning algorithm on the same task set as demonstrations, then on a generalized version of the demonstrated task. Furthermore, the quality of the provided demonstrations in repeated trials will be used to define the adaptation level of the users to the taught task.

II. BACKGROUND AND RELATED WORK

A number of researchers have attempted to assess the quality of the provided demonstrations. Ureche and Billard [11]

Manuscript received 24 February 2022; accepted 24 June 2022. Date of publication 18 July 2022; date of current version 26 July 2022. This letter was recommended for publication by Associate Editor S. Chernova and Editor J. Kober upon evaluation of the reviewers' comments. This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). (Corresponding author: Maram Sakr.)

Maram Sakr is with the Mechanical Engineering, University of British Columbia, Vancouver, BC V6T 1Z4, Canada, and also with the Faculty of Engineering, Monash University, Clayton, VIC 3800, Australia (e-mail: maram.sakr@ubc.ca).

Zexi Jesse Li and H. F. Machiel Van der Loos are with the Mechanical Engineering, University of British Columbia, Vancouver, BC V6T 1Z4, Canada (e-mail: lijesse@student.ubc.ca; vdl@mech.ubc.ca).

Dana Kulić and Elizabeth A. Croft are with the Faculty of Engineering, Monash University, Clayton, VIC 3800, Australia (e-mail: dana.kulic@monash.edu; elizabeth.croft@monash.edu).

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the UBC's Behavioural Research Ethics Board under Application No. H20-03740.

Digital Object Identifier 10.1109/LRA.2022.3191950

proposed three metrics for assessing human demonstrations of bi-manual tasks. These metrics are the ability to maneuver the tool, the consistency in teaching and the degree of coordination between the two arms. The main limitation of these metrics is that the consistency in teaching does not directly reflect the quality. The users may provide consistent demonstrations with the same errors due to their lack of understanding of *what* and *how* the robot learns during task demonstration [8], [12].

Kaiser *et al.* [13] listed different sources of sub-optimality in the human demonstrations including i) unnecessary actions that do not contribute to the final goal; ii) incorrect actions that negatively affect the usefulness of the demonstrations; and, iii) unmotivated actions by the human teacher that are measured by sensors that are not available to the robot. To avoid unmotivated actions, teaching methods that record the demonstration data from the robot's body can be used [14]. To avoid unnecessary actions, a metric was proposed to detect such actions, leaving it up to the experimenter to include demonstrations with unnecessary actions or ignore them. In [13], the authors noted that it was difficult to distinguish incorrect actions from correct ones.

Fischer *et al.* [15] compared different teaching interfaces in robot learning from demonstration. They found three main errors commonly committed by the users: i) applying too much pressure to the gripper's fingers, ii) moving into a singularity, and iii) moving into a self-collision. However, they did not study the effect of these errors on robot learning and generalization. Jaquier *et al.* [16] proposed a framework that allows robots to learn and reproduce the joint space trajectories with their particular manipulability indices.

Recently, Sena and Howard [10] identified three main issues with poor quality demonstrations. These issues are undemonstrated states, ambiguous demonstrations and unsuccessful demonstrations. Undemonstrated states refer to the states in which the robot can not perform the task because they have insufficient demonstrations. Ambiguous demonstrations refer to demonstrations that offer little or no new information to the learning model. Unsuccessful demonstrations refer to the demonstrations that do not achieve the task goal [9]. Sena and Howard propose two metrics for measuring demonstration sufficiency: i) teaching efficacy, that is, how much the robot can generalize over the entire task space; and ii) teaching efficiency, that is, teaching efficacy normalized by the number of provided demonstrations. These two metrics are used to solve one of the main issues when collecting demonstrations from human teachers, namely data sparsity [1]. However, these metrics do not address the *quality* of the provided demonstrations that affect robot learning and generalization [16].

In this paper, we propose relating the quality of the provided demonstrations to task learning and generalization they achieve. Learning is defined with respect to the task space covered by the demonstrations, and generalization is defined with respect to the task space spanned by the demonstration points. We hypothesize that if demonstrations lead to high performance in task learning at specific points, they will also achieve high performance at nearby points. This is valid if the learning algorithm is compatible with the data type and size, and the test conditions in the generalization are from the same distribution as the conditions for the demonstrations [17].

III. PROPOSED APPROACH

Several factors contribute to the generalization ability of any LfD system. For instance, the number of the provided

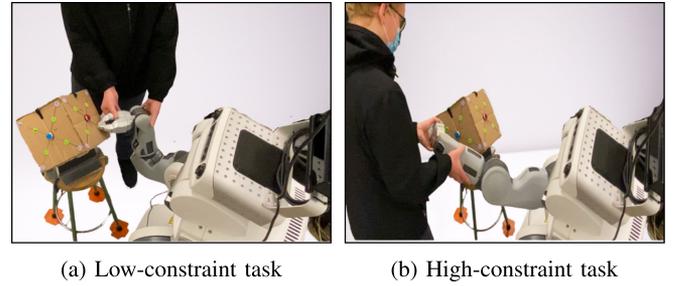


Fig. 1. Human demonstrators kinesthetically teach a robot to press a button on a box on (a) a low-constraint face close to the robot and (b) a high-constraint face such that the robot must be maneuvered around the box in a tight space.

demonstrations, their distribution over the task space, the demonstrator's expertise level, and the complexity of the task, among others. In this paper, we focus on the quality of the provided demonstrations over two levels of task complexity while fixing all other factors (i.e., the number of provided demonstrations and the task space).

Since task performance and generalization are the most important goals in any LfD system [4], we propose using the task performance and generalization performance to judge the quality of the input demonstrations. The provided demonstrations are fitted into a learning model. This model will be evaluated on the same task set covered by the demonstrations and on an adjacent task set to evaluate its generalizability. If the learning model performs the task on the same demonstration conditions with a success rate higher than a pre-defined threshold (δ), the provided demonstrations will be considered high-quality ones. Otherwise, the provided demonstrations will be considered low-quality ones, as shown in Fig. 2. The success criteria are task-dependent metrics that need to be defined beforehand. After classifying the provided demonstrations, the learning model will be tested on a new task set to evaluate its generalizability. This will be done to evaluate our hypothesis that the demonstrations that have high performance on the demonstrated task set will also have high generalization performance and vice versa.

A. Robot Task Learning

The learning process starts with the user providing a demonstration set consisting of M demonstrations. Each demonstration contains T_m state data points. Here we used the state measurements $\xi_n = (t_n, \mathbf{x}_n^T, \epsilon_n^T) \in \mathbb{R}^8$, which involves the time t , the end effector position \mathbf{x}_n^T , and the end effector orientation represented by quaternions ϵ_n^T .

A task parameterized Gaussian Mixture Model (TP-GMM) combined with Gaussian Mixture Regression (GMR) is used for task learning. TP-GMMs have been extensively used in the LfD literature [4], [18], and provide good generalization using a limited set of demonstrations. TP-GMM models a task using task parameters defined by a sequence of coordinate frames. In a D -dimensional space, each task parameter/coordinate frame is given by an $A \in \mathbb{R}^{D \times D}$ matrix indicating its orientation and a $\mathbf{b} \in \mathbb{R}^D$ vector indicating its origin, relative to the global frame. A K -component mixture model is fitted to the data in each local frame of reference. Each GMM is described by $(\pi_k, \mu_k^{(j)}, \Sigma_k^{(j)})$, referring to the prior probabilities, means, and covariance matrices for each component k in frame j , respectively. An expectation-maximization (EM) algorithm is used to

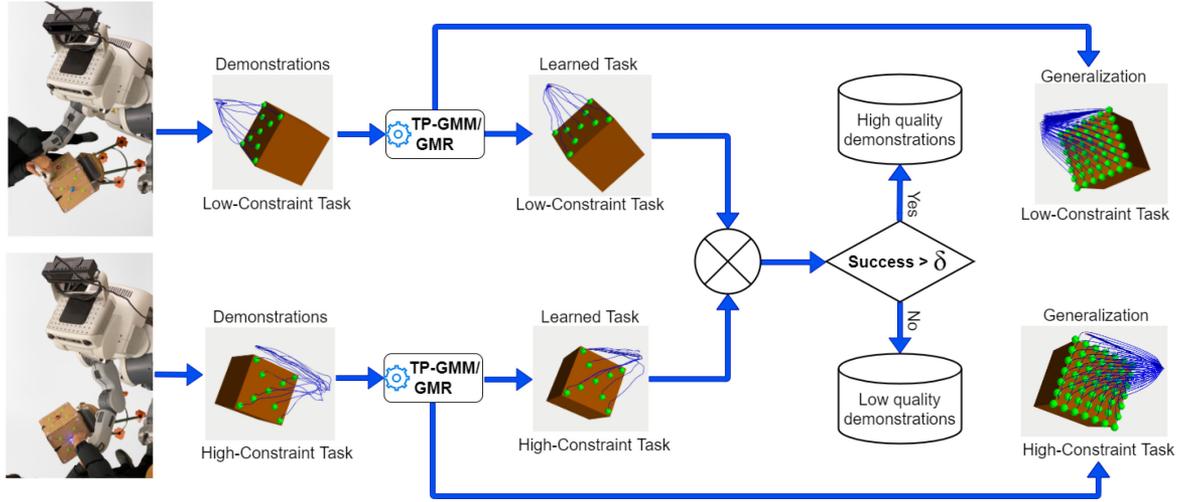


Fig. 2. The proposed approach for assessing demonstrations using learned task performance in a task with two levels of constraints.

estimate these parameters. To use the local models for trajectory generation, they must be projected back into the global frame of reference and then combined into one global model. This is achieved through a linear transformation of the local models with their respective task parameters, followed by a product of Gaussians. A new GMM with components $\{\pi_k, \mu_{k,t}, \Sigma_{k,t}\}_{k=1}^K$ at time t in the global frame $\{O\}$ can be computed as:

$$\mathcal{N}(\mu_{k,t}, \Sigma_{k,t}) \propto \prod_{j=1}^J \mathcal{N}(\mathbf{A}_t^{(j)} \mu_k^{(j)} + \mathbf{b}_t^{(j)}, \mathbf{A}_t^{(j)} \Sigma_k^{(j)} (\mathbf{A}_t^{(j)})^T) \quad (1)$$

Then, GMR can be used to obtain the next trajectory point. This procedure is repeated for each time step in the trajectory. Calinon [19] provides more in-depth detail for this approach.

B. Joint Space Learning

Given the task space trajectory from TP-GMM represented as task position $x(t)$ and task velocity $\dot{x}(t)$, the goal is to find a feasible joint space trajectory as joint position $q(t)$ and joint velocity $\dot{q}(t)$ that reproduce the given trajectory. The differential kinematics equation establishes a linear mapping between joint space velocities and task space velocities, and it can be utilized to solve for joint velocities. However, due to the non-square Jacobian matrix for the 7-DOF manipulator, the basic inverse solution to the kinematic equation is obtained by using the pseudoinverse J^\dagger of the matrix J . Where the pseudoinverse J^\dagger can be computed as $J^T(JJ^T)^{-1}$.

To avoid kinematic singularities, we used a singularity-robust (SR) inverse [20]. For a 7-DoF redundant manipulator, a nonempty null space exists due to the excess of input space relative to the manipulable space. We used a common method of including the null space in a solution with the formulation in [21] as follows:

$$\dot{q} = J^*(q)\dot{x} + (I - J^\dagger(q)J(q))\pi(q) \quad (2)$$

J^* is the SR inverse Jacobian, the pseudoinverse J^\dagger can be computed as $J^T(JJ^T)^{-1}$, the matrix $(I - J^\dagger(q)J(q))$ is the null space projection operator, and $\pi(q)$ is the null space policy.

$\pi(q)$ can be used to control motion in joint space without affecting the task-space motion. Since our goal is to mimic human demonstrations in both Cartesian and joint spaces, we used the closest demonstration as the null space policy $\pi(q)$ for the generated trajectory.

The open-loop solutions of joint variables through numerical integration unavoidably leads to drift and accumulated task space errors. To overcome these drawbacks, the closed-loop inverse kinematics (CLIK) algorithm with error feedback is utilized. The CLIK algorithm can be expressed by the following equation:

$$\dot{q} = J^*(q)(\dot{x}_d + k_p(x_d - x)) + (I - J^\dagger(q)J(q))\pi(q) \quad (3)$$

where k_p is a symmetric positive definite matrix, and the choice of k_p guarantees that the error uniformly converges to zero. \dot{x}_d is the desired velocity in task space, and $(x_d - x)$ is the position error in task space.

IV. EXPERIMENTAL DESIGN

A. Robot Platform

The robot platform used in this work is the PR2 (Willow Garage, Personal Robot 2), a mobile manipulator with two 7-DoF arms. The passive spring counterbalance system in PR2's arms provides gravity compensation, giving users the ability to kinesthetically move the robot's arms within their kinematic range. Each arm has a 1-DoF under-actuated gripper. In this experiment, we only used the right arm in gravity compensation mode with the gripper closed.

B. Task Definition

The exemplar task chosen is pressing a button; this is a general task for pressing a doorbell, elevator call button, pedestrian crossing button, etc. The task was chosen to be sufficiently generic that it does not require domain expertise, but does require to practice with the robot to provide high-quality demonstrations. This task comprises both a constrained reaching task as well as a fine control motion for pressing the button. Furthermore, the task was subject to two different levels of configuration

and task-space constraints: low-constraint and high-constraint. This line of investigation is motivated by Fitt's law [22] where the logarithm of the ratio of the target distance and the target width is used to represent the difficulty level. Here, we used target distance and constraint level as the difficulty aspects.

Fig. 1 shows the experimental setup used for data collection. As shown in the figure, a cardboard box was fixed on one of its vertices such that all buttons are reachable by the robot gripper. Only two faces of the box were used in the data collection. On each face, buttons were placed in the centre (large green button), corners (purple foam markers), and at locations midway between the corners and the centre of each face (green foam markers). A total of nine-goal positions for each face were used. Face-1 represents a low-constraint task as the robot can easily reach all the target points. Face-2 represents a high-constraint task as the participant needs to maneuver the robot arm around the box in a constrained space to reach the goal positions while avoiding self collisions and collisions with the box as shown in Fig. 1-b.

C. User Study

We recruited participants for our user study through advertisements on university campus, and social media. A total of 27 participants (21 male, 6 female) with $\mu_{age} = 22$ with different robotics expertise levels ranging from no robotics experience to 6 years or more of experience were recruited. Prior to conducting the study, we obtained research ethics approval from the university's Behavioural Research Ethics Board (application ID H20-03740). We obtained informed consent from each participant before commencing the experiment.

The experiment was conducted in a two-session regime on two different days. This two-session format is motivated by the fact that practice is essential for motor skill learning [23]. We were particularly interested to explore which aspects improved from the first session to the second and how this will be reflected in robot learning and generalization. Participants were scheduled ahead of time for one hour time slots. Upon arrival, they were given an informed consent form to sign. The experimenter briefly stated the long-term goal of the research, and told the participants that their task will be to program new skills on the PR2. The participants were told that the PR2 robot will imitate and learn from their demonstrations and should be able to press a button located anywhere in the box's face after their demonstrations. The experimenter did not give any instructions on how to program the PR2, rather participants were asked to learn by doing, in order to record their first-time (novice) interaction with the robot. The robot was set in gravity compensation mode. Participants were asked to hold the robot's right arm and physically guide it to press a target button (kinesthetic teaching). The right arm always started at the same position with elbow at 90° and the gripper pointed up (untucked position). The experimenter teleoperated the right arm via joystick to bring the arm back to the initial position after each demonstration. The participants started with guiding the robot to press the nine Face-1 buttons (low-constraint task) and then progressed to the nine Face-2 buttons (high-constraint task) with three trials for each face. The robot joint angles were recorded during each demonstration and saved as ROSbag files for offline analysis. A total of 54 demonstrations were collected in a session for each participant. At the end of the first session, the experimenter asked the participants to reflect on what they had learned and practiced in the first session on how to teach the robot and to schedule a

second session on a different day. This procedure was motivated by the findings of Walker *et al.* [24] that a night of sleep after training on motor skill significantly improves the skill level in a later retesting. In the second session, each participant repeated the same procedure as the first session. At the end of the second session, each participant answered a questionnaire¹ about their previous experience in robotics, if any. We explicitly instructed the participants not to include the experience gained throughout the experiment as their prior experience. We also asked them to write more details about their previous experience to make sure that they did not consider this experiment in their answer. In addition, they answered open questions about the aspects of their performance that they felt improved between the two sessions.

D. Performance Metrics

To evaluate the collected demonstrations, a TP-GMM model was created using each trial's demonstrations as described in Section III-A. Then, the model was tested on the same demonstrated task set as well as 49 new target positions (generalization as shown in Fig. 2). Each face of the box in Figure 1 was discretized into a grid of 7×7 that gave a total of 49 new targets. The grid size was chosen based on the dimensions of the PR2 arm's tip and the box dimensions to avoid any overlap between targets. The arm tip's dimensions are ($W = 2.1$ cm, $L = 2.2$ cm, $H = 3.5$ cm), and the box is a cube of equally sized edges of 26 cm. We specified the target point as a sphere of 3 cm in diameter based on the centre button. The inverse kinematics (IK) procedure in Section III-B was used to compute joint-space trajectory for the learned Cartesian trajectory from TP-GMM. From the joint-space trajectory we can check whether the learned trajectory is a feasible and compute the collision rate of the robot with the experimental setup.

The learned trajectory is considered successful if it reached the goal position (within the goal sphere) while avoiding self collisions and collisions with the box. To account for the non-zero size of the end effector tip, we consider the robot to have reached the goal if any point of the tip touches the goal sphere. To calculate the success rate for each trial, we divide the number of successful trajectories over the total number of the tested points (nine in the task performance test and 49 in the generalization test).

E. Fast-Adapter and Slow-Adapter Users Definition

The proposed approach in Fig. 2 was utilized to classify the provided demonstrations in the two sessions into low-quality and high-quality ones. We found a group of users who provided high-quality demonstrations in both sessions, while others provide low-quality demonstrations in the first session and then after practice they provide high-quality ones in the second session. We call the first group of users *fast-adapters* as they quickly adapt to the task and provide high-quality demonstrations consistently. The other group was called *slow-adapters* as they need some time and practice to adapt to the task and provide high-quality demonstrations.

V. HYPOTHESES

We expect that task performance and generalization performance will be highly correlated, as discussed in Section II.

¹[Online]. Available: <https://forms.gle/TJ6ULsWVkv2dwUpY6>

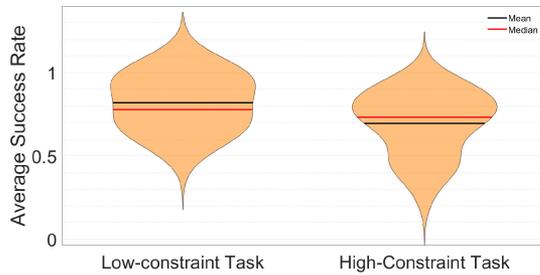


Fig. 3. Average success rate distribution of task performance of Session-1 on both task constraints.

We used the first session task performance for clustering users into fast and slow adapters as described in Section III. We are interested to see if fast adapters will consistently provide high-quality demonstrations in the second session. Also, we are interested to see if slow-adapters improve their performance in their second session. With practice, we expect slow-adapters will converge to a strategy for teaching the robot how to do a task. This will be reflected in the improvement in the success rate from the first session to the second one. For fast-adapters, they already demonstrated good performance in the first session, so we expect that fast-adapters will not have a significant difference in their performance in the two sessions. The support of these hypothesis will show the effectiveness of the proposed approach for assessing demonstrations quality and for defining the adaptation level of the users. Based on these expectations, we formulate the following hypotheses:

- **H1:** There is a significant correlation between task performance and generalization performance.
- **H2:** The task performance of fast-adapters' demonstrations is significantly higher than slow-adapters' demonstrations.
- **H3:** The task performance using second session demonstrations is significantly higher than using the first session ones for slow-adapters.
- **H4:** The generalization performance of fast-adapters' demonstrations is significantly higher than slow-adapters' demonstrations.
- **H5:** The generalization performance using second session demonstrations is significantly higher than using the first session ones for slow-adapters.

VI. RESULTS

In order to evaluate all hypotheses, a TP-GMM model was created using each trial's demonstrations and tested on the same demonstrated task set (task performance) as well as 49 new target positions (generalization performance). To study the correlation between task performance and generalization, we checked the normality of the data before calculating the correlation coefficient. We found that the data does not follow a normal distribution, so we used the Spearman correlation coefficient [25]. We found a statistically significant positive correlation between task and generalization performance ($\rho = 0.76, p < 0.0001$). Following the proposed approach in Fig. 2, we used task performance success rate to cluster the users. We used the violin plot to define the number of clusters and the clustering threshold. We looked into the distribution of the average success rate of Session-1 trials for all participants as shown in Fig. 3. The violin figure shows a clear bi-modal distribution especially with

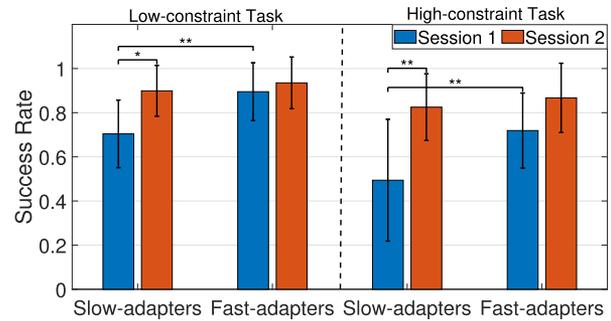


Fig. 4. The average success rate of the learned task for both slow- and fast-adapters in the two sessions for the high and low constraint tasks.

high-constraint tasks. This shows that the data is clusterable into two clusters. Since the mean and median of both task constraints are different, the threshold is defined as a value between them. A threshold of 80% was chosen to categorize the users into: fast-adapters and slow-adapters. We found that 12 participants out of the 27 participants are fast-adapters with a success rate $> 80\%$ for both low- and high- constraint tasks using the first session's demonstrations. Only six participants out of these 12 fast-adapters claimed previous robotics experience in the questionnaire, while four participants out of the 15 slow-adapters claimed previous experience in robotics. An interesting finding is that all participants adapted to the experimented task, they only differ in their adaptation level across the experiment.

A. Task Performance Analysis

Fig. 4 shows the mean and standard deviation for success rate in task performance for both fast-adapters and slow-adapters in two sessions and two task constraints. Overall, Session-2 has a higher success rate with less variance than Session-1 for both fast- and slow-adapters in both task constraints. Furthermore, the improvement between the two sessions in the high-constraint tasks is higher than the improvement in the low-constraint tasks. Finally, there is a large difference between fast-adapters and slow-adapters in Session-1 while in Session-2 this difference decreased.

A 2x2x2 mixed model ANOVA was conducted to investigate the impact of a) Adaptation level, b) Sessions and c) Task constraints on the learned task success rate. The data were checked to be compatible with the relevant statistical assumptions. We found that six sets of data out of the eight we have do not follow a normal distribution. ANOVA tests are noted as being robust to violations of normality [26], and considering that the excess kurtosis for the six groups is $-0.525, 0.854, 7.561, 0.897, 0.538,$ and 0.522 (where a normal distribution would have an excess kurtosis value of zero), the violation is considered minor except for one case and the data are assumed to follow a normal distribution. The assumption of sphericity was not applicable as there were only two levels of all factors. Levene's test showed that all data sets do not violate the homogeneity of variances assumption.

Table I shows the significance level of all factors and their interaction for task performance results. The pairwise comparisons with Bonferroni correction showed that slow-adapters have significant improvement in their success rate from Session-1

TABLE I
2 X 2 X 2 MIXED MODEL ANOVA RESULTS OF TASK PERFORMANCE
SUCCESS RATE

Factor(s)	F	p-value	Levels	Marginal Means
Sessions	32.94	$p < 0.001$	S1 S2	0.70 ± 0.027 0.88 ± 0.018
Adaptation Level	2170.77	$p < 0.001$	Slow Fast	0.73 ± 0.023 0.85 ± 0.025
Task Constraints	15.26	$p < 0.001$	Low High	0.86 ± 0.019 0.73 ± 0.028
Adaptation Level X Sessions	7.36	$p = 0.012$	Slow in S1 Slow in S2 Fast in S1 Fast in S2	0.60 ± 0.037 0.86 ± 0.023 0.81 ± 0.041 0.90 ± 0.026

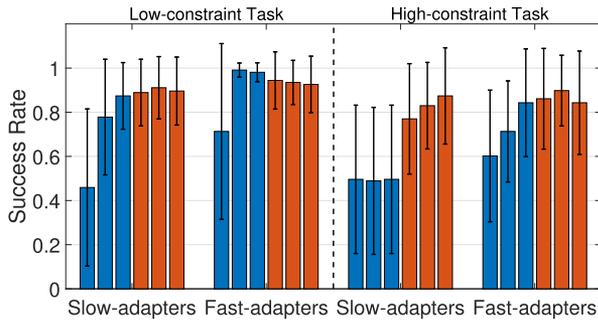


Fig. 5. The success rate of the learned task for both slow- and fast- adapters in all trials. Blue bars represent Session-1 trials and red bars represent Session-2 trials.

to Session-2 $p < 0.001$ while fast-adapters have marginal improvement from Session-1 to Session-2 $p = 0.053$. Another interesting finding is that there is a significant difference between fast- and slow-adapters in Session-1 $p < 0.001$ while there is no significant difference between them in Session-2.

Fig. 5 shows the success rate in the individual trials in both sessions. It is shown that slow-adapters tend to take a longer time to reach high performance than the fast-adapters. Particularly in the high-constraint tasks, slow-adapters have a steady low performance in the first session's trials with 0.49 ± 0.34 , then they gradually improve in the second session's trials till they reach 0.87 ± 0.22 . On the other hand, fast-adapters gradually improve in Session-1 and almost converge to high performance at the end of their first session with 0.90 ± 0.16 . For the low-constraint tasks, slow-adapters converge to their highest performance at the end of Session-1 with 0.87 ± 0.15 while fast-adapters converge to a higher performance of 0.99 ± 0.032 than slow-adapters right in the second trial of Session-1.

B. Task Generalization Performance Analysis

Fig. 6 shows the mean and standard deviation for success rate in task generalization for both fast- and slow-adapters for two sessions and two task constraints. Overall, it shows a similar trend as the task performance figure with Session-2 having a higher success rate with less variance than Session-1 for both fast- and slow-adapters in both task constraints. Also, there is a large difference between slow- and fast-adapters in Session-1 while in Session-2 this difference decreased. Finally, the low-constraint tasks shows a consistent success rate in

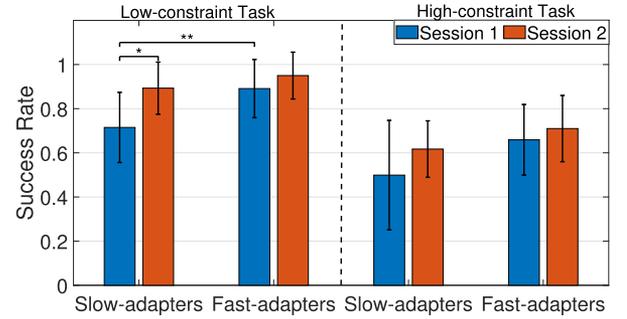


Fig. 6. The average generalization success rate for both slow- and fast-adapters in the two sessions for the high and low constraint tasks.

TABLE II
2 X 2 X 2 MIXED MODEL ANOVA RESULTS OF TASK GENERALIZATION
SUCCESS RATE

Factor(s)	F	p-value	Levels	Marginal Means
Sessions	14.93	$p < 0.001$	S1 S2	0.69 ± 0.025 0.79 ± 0.014
Adaptation Level	2402.87	$p < 0.001$	Slow Fast	0.68 ± 0.020 0.80 ± 0.023
Task Constraints	48.36	$p < 0.001$	Low High	0.86 ± 0.02 0.62 ± 0.026

both task learning and generalization performance while the high-constraint tasks has a degradation in the success rate in generalization performance compared to task learning.

A $2 \times 2 \times 2$ mixed model ANOVA was conducted to investigate the impact of a) Adaptation level, b) Sessions and c) Task constraints on generalization success score. The data were checked to be compatible with the relevant statistical assumptions. We found that only three sets of data out of the eight did not follow a normal distribution. The excess kurtosis for the three groups is -0.823 , 0.083 , and 7.74 ; therefore the violation is considered minor except for one case and the data are assumed to follow a normal distribution. Levene's test showed that all eight data sets did not violate the homogeneity of variances assumption. Table II shows the significance level of all factors and their interaction for task generalization results.

Examining the individual trials of both sessions, we found a similar trend as with learned task performance, with slow-adapters taking longer to reach high performance than the fast-adapters as shown in Fig. 7. Particularly in the high-constraint tasks, slow-adapters have a steady low performance in the first session's trials with 0.51 ± 0.34 and then gradually improve in the second session until they reach 0.70 ± 0.15 . Fast-adapters demonstrate gradual improvement in Session-1 trials and almost converge to high performance at the end of their first session with 0.74 ± 0.21 . For the low-constraint tasks, slow-adapters converge to their highest performance in the second trial of Session-2 with 0.94 ± 0.11 while fast-adapters converge to a higher performance of 0.98 ± 0.06 than slow-adapters in the second trial of Session-1.

C. Mixed Quality Demonstrations Analysis

To show the importance of clustering the demonstrations, we used mixed quality demonstrations for learning and compare

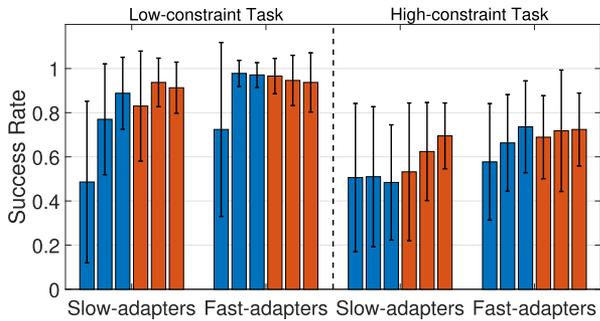


Fig. 7. The success rate of task generalization for both slow- and fast-adapters in all trials. Blue bars represent Session-1 trials and red bars represent Session-2 trials.

the resulting performance to the performance using only high quality demonstrations. We did this test with slow-adapters as they provide a larger range of mixed quality demonstrations. We combined all demonstrations from the three trials of the two sessions, providing 27 demonstrations per session. A TP-GMM model was trained using these mixed quality demonstrations and then task performance and generalization was calculated. This performance is compared to the performance of using only nine high-quality demonstrations. We found that using one trial of high-quality demonstrations results in performance 14% higher than using 27 demonstrations from both sessions of the high-constraint task.

VII. DISCUSSION

The results show a strong correlation between task performance and generalization. Also, a clear difference between fast- and slow-adapters in both task performance and generalization was shown. In addition, practice makes a significant difference in slow-adapters' performance, approaching fast-adapters' performance. This finding agrees with our previous work demonstrating the need for training for slow-adapters to provide high-quality demonstrations for robots [27].

The categorization of users into fast- and slow-adapters showed that the stated robotics expertise may not be reflected in the task performance and generalization. The adaptation level to the task may be affected by previous expertise in domains other than robotics [5]. This suggests that using only the questionnaire may not be sufficient as skills from other domains could be transferable to the robotics domain and contribute to the participant's performance and adaptation [6].

Results show support for **H1** with a significant correlation between task performance and generalization. Task learning results show support for **H2** with a significant main effect of the adaptation level on the success rate. Similarly, support for **H3** was found as there is a significant difference between both sessions for slow-adapters with 43% improvement between sessions, while fast-adapters have only 11% improvement. The importance of practice for slow-adapters was shown with the significant difference between fast- and slow-adapters in Session-1 with an estimated mean difference of 0.21 ± 0.055 , while no significant difference between them was observed in Session-2, the mean difference was 0.04 ± 0.035 . This result agrees with Thorndike [28] who suggests that the individual differences in learning could be determined by assessing whether individuals converged or diverged after practice on a task. If people became

more alike after practice, then individual differences observed at the start of practice were more likely attributable to previously acquired learning skills.

Another interesting finding is that the difference between the two sessions in the high-constraint task is double the difference in the low-constraint task. The reason is the greater room for improvement in the high-constraint task while the low-constraint task is straightforward and does not require the same amount of practice. This suggests taking task complexity into consideration when providing training for slow-adapters before teaching robots by demonstration.

Task generalization results show a similar pattern as the task learning results, providing support for **H4**, and partial support for **H5**. Although a non-significant interaction between sessions and adaptation was detected, the estimated marginal means increased by 25% from Session-1 to Session-2. Furthermore, the difference between fast- and slow-adapters in the first session was 0.19 ± 0.049 , while in Session-2 this difference dropped to 0.08 ± 0.028 . Unlike task learning results, the improvement between sessions in the low-constraint task was 0.12 ± 0.032 , while in the high-constraint task it was 0.09 ± 0.047 . This is because extrapolation in the constrained space is challenging, especially for the farthest points on the box's face.

Looking into individual trials performance in Fig. 5 and Fig. 7, we found that both fast- and slow-adapters achieve low performance in the very first trial with high variance. This suggests that the first demonstrations for a task should not be considered for robot learning regardless of the user's stated expertise level. This reasoning aligns with [8]. Furthermore, fast-adapters quickly reach their highest performance after the first trial, while slow-adapters take a longer time. This suggests that fast-adapters may need three trials or less for training on the task while slow-adapters may need six trials of training to provide high-quality demonstrations. It was also interesting to see that the participants' performance improved across trials without any feedback about their performance. Not only that, but also some participants noticed some well-known issues in the provided demonstrations (e.g., unnecessary motions [13] and joint limits [15]) and then they avoid these issues in the later trials.

One way to explain these findings is through the motor skill learning literature. In [29], Fitts and Posner describe a model for skill acquisition. Their model proposes three stages for acquiring a motor skill: the cognitive stage, associative stage and autonomous stage. In the first stage, a trainee processes the received information about the task and tries out several strategies to perform it. This results in a high cognitive load as well as highly variable performance. After finding a good strategy, the trainee moves to the associative stage at which he/she tries to refine the strategy to improve the task performance. In the last stage, the autonomous stage, the trainee can perform the task more accurately with less cognitive load.

VIII. CONCLUSION

In this letter, we proposed a framework for defining the quality of the provided demonstrations using task learning and generalization performance. This proposed framework was validated in a generic learning task with two levels of constraints. The demonstrations data were collected in two sessions on two different days to determine the improvement pattern in robot teaching skills for both slow- and fast-adapters. The results

show a significant correlation between task performance and generalization performance. In addition, a significant difference between slow- and fast-adapters was shown in both task constraints. Furthermore, slow-adapters show a significant improvement from the first session to the second one and even get close to the performance of the fast-adapters in the second session. Fast-adapters show consistent performance between the two sessions as they defined the strategy for approaching the task faster than slow-adapters. We showed how these results are consistent with the motor skills learning literature.

We believe that the results of this paper open up several directions for future research. The proposed framework can be utilized in active learning in which the user will provide more demonstrations if the provided ones are low-quality and limit robot learning and generalization. Furthermore, a promising future direction would be exploring the properties of the provided demonstrations to better understand what makes high-quality demonstrations and low-quality ones [30].

REFERENCES

- [1] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robot. Auton. Syst.*, vol. 57, no. 5, pp. 469–483, 2009.
- [2] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, "Recent advances in robot learning from demonstration," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 3, pp. 297–330, 2020.
- [3] S. Calinon and A. G. Billard, "What is the teacher's role in robot programming by demonstration?: Toward benchmarks for improved learning," *Interact. Stud.*, vol. 8, no. 3, pp. 441–464, 2007.
- [4] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, and J. Peters, "An algorithmic perspective on imitation learning," vol. 7, no. 1/2, pp. 1–179, 2018.
- [5] D. Kobak and C. Mehring, "Adaptation paths to novel motor tasks are shaped by prior structure learning," *J. Neurosci.*, vol. 32, no. 29, pp. 9898–9908, 2012.
- [6] H. T. Everson *et al.*, "Do metacognitive skills and learning strategies transfer across domains?," ERIC, Tech. Rep., 1997.
- [7] G. P. Latham, "Human resource training and development," *Annu. Rev. Psychol.*, vol. 39, no. 1, pp. 545–582, 1988.
- [8] M. Cakmak and A. L. Thomaz, "Eliciting good teaching from humans for machine learners," *Artif. Intell.*, vol. 217, pp. 198–215, 2014.
- [9] D. H. Grollman and A. Billard, "Donut as i do: Learning from failed demonstrations," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2011, pp. 3804–3809.
- [10] A. Sena and M. Howard, "Quantifying teaching behavior in robot learning from demonstration," *Int. J. Robot. Res.*, vol. 39, no. 1, pp. 54–72, 2020.
- [11] A.-L. P. Ureche and A. Billard, "Metrics for assessing human skill when demonstrating a bimanual task to a robot," in *Proc. 10th Annu. ACM/IEEE Int. Conf. Hum.-Robot Interact. Extended Abstr.*, 2015, pp. 37–38.
- [12] T. Hellström and S. Bensch, "Understandable robots-what, why, and how," *Paladyn, J. Behav. Robot.*, vol. 9, no. 1, pp. 110–123, 2018.
- [13] M. Kaiser, H. Friedrich, and R. Dillmann, "Obtaining good performance from a bad teacher," *Program. Demonstration Learn. Examples Workshop ML*, vol. 95, 1995.
- [14] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, "Robot programming by demonstration," in *Springer Handbook Robot*. Springer, 2008, pp. 1371–1394.
- [15] K. Fischer *et al.*, "A comparison of types of robot control for programming by demonstration," in *Proc. 11th ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2016, pp. 213–220.
- [16] N. Jaquier, L. Rozo, D. G. Caldwell, and S. Calinon, "Geometry-aware manipulability learning, tracking, and transfer," *Int. J. Robot. Res.*, vol. 40, no. 2-3, pp. 624–650, 2020.
- [17] M. Sugiyama, *Introduction to Statistical Machine Learning*. San Mateo, CA, USA: Morgan Kaufmann, 2015.
- [18] A. Pervez and D. Lee, "Learning task-parameterized dynamic movement primitives using mixture of GMMs," *Intell. Ser. Robot.*, vol. 11, no. 1, pp. 61–78, 2018.
- [19] S. Calinon, "A tutorial on task-parameterized movement learning and retrieval," *Intell. Serv. Robot.*, vol. 9, no. 1, pp. 1–29, 2016.
- [20] Y. Nakamura and H. Hanafusa, "Inverse kinematic solutions with singularity robustness for robot manipulator control," *J. Dyn. Syst., Meas., Control.*, vol. 108, no. 3, pp. 163–171, 1986.
- [21] A. Liegeois *et al.*, "Automatic supervisory control of the configuration and behavior of multibody mechanisms," *IEEE Trans. Syst., Man, Cybern.*, vol. 7, no. 12, pp. 868–871, Dec. 1977.
- [22] P. M. Fitts, "The information capacity of the human motor system in controlling the amplitude of movement," *J. Exp. Psychol.*, vol. 47, no. 6, 1954, Art. no. 381.
- [23] R. A. Schmidt and C. A. Wrisberg, *Motor Learning and Performance: A Situation-Based Learning Approach*. Human Kinetics, 2008.
- [24] M. P. Walker, T. Brakefield, J. Seidman, A. Morgan, J. A. Hobson, and R. Stickgold, "Sleep and the time course of motor skill learning," *Learn. Memory*, vol. 10, no. 4, pp. 275–284, 2003.
- [25] A. J. Bishara and J. B. Hittner, "Testing the significance of a correlation with nonnormal data: Comparison of pearson, spearman, transformation, and resampling approaches," *Psychol. Methods*, vol. 17, no. 3, 2012, Art. no. 399.
- [26] E. Schmider, M. Ziegler, E. Danay, L. Beyer, and M. Bühner, "Is it really robust? Reinvestigating the robustness of ANOVA against violations of the normal distribution assumption," *Methodol.: Eur. J. Res. Methods Behav. Social Sci.*, vol. 6, no. 4, p. 147, 2010.
- [27] M. Sakr, M. Freeman, H. Machiel, V. der Loos, and E. Croft, "Training human teacher to improve robot learning from demonstration: A pilot study on kinesthetic teaching," in *Proc. 29th IEEE Int. Conf. Robot Hum. Interactive Commun.*, 2020, pp. 800–806.
- [28] E. L. Thorndike, "The effect of practice in the case of a purely intellectual function," *Amer. J. Psychol.*, vol. 19, no. 3, pp. 374–384, 1908.
- [29] P. M. Fitts and M. I. Posner, *Human Performance*. Belmont, CA, USA: Brooks/Cole Pub. Co., 1967.
- [30] M. Sakr, H. M. Van der Loos, D. Kulić, and E. Croft, "What makes a good demonstration for robot learning generalization?," in *Proc. Companion ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2021, pp. 607–609.