



PDF Download
3757279.3785606.pdf
16 March 2026
Total Citations: 0
Total Downloads: 0

 Latest updates: <https://dl.acm.org/doi/10.1145/3757279.3785606>

RESEARCH-ARTICLE

Investigating the Impact of Robot Degree of Redundancy on Learning from Demonstration

MUHAMMAD BILAL, University of Melbourne, Melbourne, VIC, Australia

D. ANTONY CHACON, University of Melbourne, Melbourne, VIC, Australia

NIR LIPOVETZKY, University of Melbourne, Melbourne, VIC, Australia

DENNY OETOMO, University of Melbourne, Melbourne, VIC, Australia

WAFI JOHAL, University of Melbourne, Melbourne, VIC, Australia

Open Access Support provided by:

University of Melbourne

Published: 16 March 2026

Citation in BibTeX format

HRI '26: 21st ACM/IEEE International
Conference on Human-Robot Interaction
March 16 - 19, 2026
Scotland, Edinburgh, UK

Conference Sponsors:
SIGAI
SIGCHI

Investigating the Impact of Robot Degree of Redundancy on Learning from Demonstration

Muhammad Bilal

School of Computing and Information Systems, The University of Melbourne
Melbourne, Australia
m.bilal@unimelb.edu.au

D. Antony Chacon

School of Computing and Information Systems, The University of Melbourne
Melbourne, Australia
antony.chacon@unimelb.edu.au

Nir Lipovetzky

School of Computing and Information Systems, The University of Melbourne
Melbourne, Australia
nir.lipovetzky@unimelb.edu.au

Denny Oetomo

Department of Mechanical Engineering, The University of Melbourne
Melbourne, Australia
doetomo@unimelb.edu.au

Wafa Johal

School of Computing and Information Systems, The University of Melbourne
Melbourne, Australia
wafa.johal@unimelb.edu.au

Abstract

Learning from Demonstration allows robots to acquire skills from human demonstrations, making them more accessible to a wider range of users. Among different approaches, kinesthetic teaching allows humans to manipulate the robot joints directly, making it effective method for demonstrating constrained tasks. However, robots with kinematic redundancy enable multiple joint configurations to achieve a desired task, which could influence human teaching performance. One one hand, it could make it easier, allowing more freedom to demonstrate the task, but on the other, it also increases the number of joints that needs to be manipulated, potentially affecting cognitive and physical load of the demonstrator. Therefore, it is crucial to investigate how the number of degrees of redundancy (DoR) impact human performance during kinesthetic demonstrations, and then how these demonstrations influence robot performance. We simulated high and low DoR by locking one of the robot joint on a 7-DoF Panda robotic arm. We conducted a within-subject user study ($N = 24$) with two conditions: unconstrained condition (high DoR) and constrained condition (low DoR). We used a motion capture system to capture participants physical interaction with the robot when demonstrating two tasks: button pressing and cuboid block insertion. The results show that the robot's DoR significantly affects mental workload, demonstration time, number of failed attempts, and physical interaction with the robot. Likewise, joint constraints significantly influence robot performance, measured by task completion using the learned model. These findings highlight the importance of considering robot DoR during demonstrating constrained tasks, allowing novice users to provide effective demonstrations.

CCS Concepts

• **Computing methodologies** → **Learning from demonstrations**; • **Human-centered computing** → *User studies*.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

HRI '26, Edinburgh, Scotland, UK

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2128-1/2026/03

<https://doi.org/10.1145/3757279.3785606>

Keywords

robot degree of redundancy, kinesthetic teaching, human performance, robot performance

ACM Reference Format:

Muhammad Bilal, D. Antony Chacon, Nir Lipovetzky, Denny Oetomo, and Wafa Johal. 2026. Investigating the Impact of Robot Degree of Redundancy on Learning from Demonstration. In *Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction (HRI '26)*, March 16–19, 2026, Edinburgh, Scotland, UK. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3757279.3785606>

This work was partially supported by the Australian Research Council (Grants No.: DE210100858; CE260100108; DP260101082; FT250100459).

1 Introduction

Learning from Demonstration (LfD) allows novice users to teach robots everyday tasks through demonstrations, eliminating the need of coding each specific task [5]. The efficacy of an LfD process can be enhanced by improving both its key elements: (1) the effectiveness of the robotic learning algorithms [11, 28, 33] and (2) the information content in the human-provided demonstrations [6, 34, 35]. In the literature, researchers have devoted considerable attention to the algorithmic approaches, with comparatively less emphasis on the human-centric aspect for enhancing LfD [25, 28, 31].

Among the available modalities for providing demonstrations [32], kinesthetic teaching offers an intuitive approach in which novice users physically guide the robot through the desired motion. This modality removes the need for additional sensing hardware, external input devices, or specialized interfaces, thereby lowering the barrier to entry for non-expert operators [5, 32]. In contrast, teleoperated teaching typically restricts user input to the end-effector motion, relying on inverse kinematics to compute the corresponding joint configurations [40, 41]. Kinesthetic teaching, however, directly captures user intent not only in Cartesian space (as in teleoperation) but also in the joint space, enabling the expression of preferences at a more granular level. As a result, it supports more comprehensive demonstrations of constrained tasks—particularly those involving joint limits, conceptual task constraints, or physical joint restrictions—because users can manipulate all robot joints to satisfy task requirements [22, 26, 27].

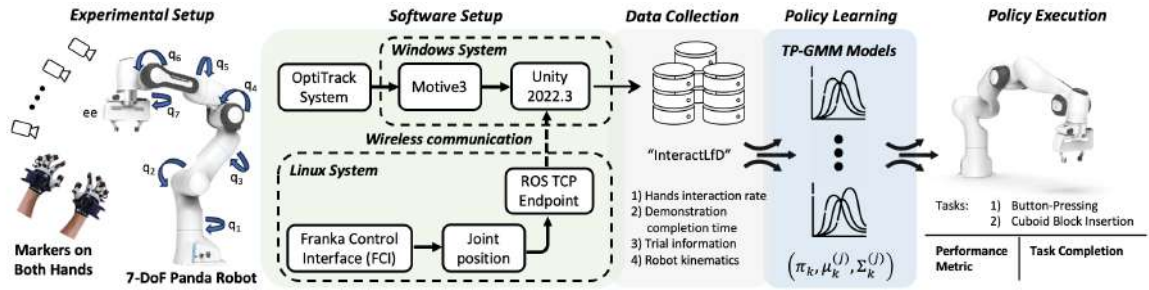


Figure 1: Study overview showing the robotic platform, motion-capture system, and experimental setup.

Kinematic redundancy—where a robot has more joint degrees of freedom than required for a given task—provides flexibility during kinesthetic demonstrations, since multiple joint configurations can achieve the same end-effector pose with different levels of efficiency and precision [1]. However, this flexibility also introduces complexity: novice users must navigate a larger set of feasible configurations, which can make it more difficult to select effective poses and provide successful demonstrations. At the same time, constraining redundancy (e.g., by limiting the number of joints to be manipulated) may reduce cognitive effort by narrowing the solution space, but may also reduce adaptability [42]. This trade-off highlights the importance of examining how redundancy shapes user strategies, demonstration quality, and downstream robot performance in LfD.

This study presents an exploratory investigation that has not been done before in the context of robot learning from demonstrations. Beyond a kinematic comparison of redundancy levels, it provides new insights into how users’ physical interaction patterns and demonstration quality influence the robot’s learning performance. To evaluate the impact of robot redundancy, we implemented a joint constraint, by locking one of the robot joints, to reduce redundancy by one joint degree. It is worth to highlight that this study only concerns the reduction of DoFs of the robotic mechanism (kinematic chain) and not the reduction of DoFs along “unwanted” task directions. Our experimental setup uses motion capture to record hand contacts with individual robot joints, enabling joint-level analysis that is, to our best knowledge, not available in current public datasets. The overall study pipeline is shown in Figure 1.

The key contributions of this study are as follows: (1) demonstrating tasks under constrained conditions imposes a higher mental workload on users; (2) task demonstration time increases as the DoR is reduced; (3) constrained condition leads to more failed attempts; (4) reduced redundancy results in increased physical interaction with the robot joints; and (5) demonstrations collected under reduced redundancy negatively affect subsequent robot performance. Furthermore, we collected a comprehensive dataset, *InteractLFD*, in this study for reproducibility and future research (available in the supplementary material).

2 Background and Related Work

One of the key goals of LfD is to allow novice users to efficiently demonstrate desired tasks to the robotic systems [32]. Among different approaches of LfD [23], human-centric based approaches suppose that users are skilled and capable of providing optimal demonstrations [32]. However, this assumption overlooks the

heterogeneous nature of human demonstrators, often resulting in suboptimal demonstrations [4, 11].

The quality of these demonstrations varies significantly among human demonstrators, influenced by their understanding of the task and their ability to deliver effective demonstrations [9]. In addition, the modalities of demonstration also play a crucial role [2, 10, 32], such as kinesthetic teaching shows promising results for constrained tasks [7, 22]. Eventually, LfD model learned using such demonstrations influences the robot performance [5, 6]. The quality of demonstration is shaped by several key factors as follows: first, the user’s mental workload must be considered, as higher mental workload can reduce their ability to focus, make precise adjustments, and ensure engagement during demonstrations [15]. Second, minimizing demonstration time is critical, as lengthy sessions can lead to user fatigue, reducing both the practicality and appeal of LfD for everyday tasks [2]. Third, previous research has shown that frequent failed attempts, demonstrations that do not accomplish the task, can undermine user confidence, further affecting demonstration quality [17, 38]. Finally, users physical interaction with the robotic system needs to be investigated, focusing on their dynamic interactions with the robot joints during demonstrations [3].

Learning from demonstrations under constraints presents key challenges that reflect real-world scenarios. For instance, robots with high DoF often exploit redundancy to perform given tasks efficiently, extending its operation within physical, environmental, or task-specific limitations [12, 24]. To address these challenges, the concept of conceptual constraints was introduced to complement low-level data with high-level abstract information, facilitating successful task execution [26]. Subsequently, Mueller et al. [27] automated the generation of conceptual constraints and incorporated them into human-aware motion planning algorithms. One such approach is to use a combination of positive and negative demonstrations to learn the constraints of a task [21]. By combining these two types of demonstrations, the algorithm can learn the constraints of the task more effectively. Another approach is to use a two-step learning framework to infer continuous constraint functions from demonstrations [29]. While it is important to develop algorithmic methods to incorporate the environmental or task constraints, it remains crucial to also investigate how robot-side constraints (i.e., different degrees of redundancy) influence human actions and, in turn, robot learning outcomes.

During demonstrations, Fischer et al. [14] identified three key challenges faced by novice users: singular robot configurations, self-collisions, and excessive force applied to the end-effector.

Additionally, Hirschmanner et al. [18] examined various user factors, including trust and workload, in the LfD process. Recently, Aliasghari et al. [3] observed that successful human teachers tend to use both hands and frequently adjust their hand placement on the robot’s arm during demonstrations, suggesting that dynamic physical interaction is crucial for effective demonstrations. Our work extends these efforts by examining the impact of robot-side constraints—specifically, the robot’s DoR—on human–robot interaction, an aspect that remains underexplored. We investigate how a robot’s DoR affects cognitive workload, demonstration time, the number of failed attempts, and users’ hand placement and adjustment strategies during task demonstrations. In addition, we examine how the robot’s DoR influences robot task performance. Based on these considerations, we formulate the following hypotheses: **H1**—Robot DoR (joint constraint) affects the time required to complete a task. **H2**—Robot DoR (joint constraint) affects the number of failed attempts. **H3**—Robot DoR (joint constraint) affects user interaction with the robot joints.

3 Selection of Tasks

We considered two manipulation tasks, button-pressing and cuboid block insertion, to collect demonstrations and evaluate the given hypotheses. These tasks have been used in the literature [10, 35] for their generic nature, ensuring they did not require domain-specific expertise.

Button Task: As illustrated in Figure 2(a) and Figure 2(b), two buttons, B1 and B2, were installed at the centers of adjacent sides of a cardboard box, which was positioned such that the robot’s end-effector could reach both buttons. Participants were required to teach kinesthetically the robot to press B1 and subsequently B2 using its end-effector tip. This task required at least five joint DoF to sequentially press both buttons. The task difficulty were defined by two factors: (1) avoiding collisions with the cardboard box as well as self-collisions of the robot, and (2) ensuring the robot’s end-effector remained perpendicular to the plane of each button during pressing.

Insertion Task: In this task, participants were required to teach the robot to pick up a randomly positioned cuboid block (measuring $5 \times 5 \times 20$ cm) and partially insert it into a designated hole with a diameter of 10 cm, as shown in Figure 2(c) and Figure 2(d). Due to the block orientation, vertically standing, completing this task required at least six joint DoF. The difficulty of this task were defined by two aspects: (1) avoiding collisions with the hole’s edges during insertion and self-collisions of the robot, and (2) the requirement to lift the cuboid block vertically, ensuring the z -axis of the robot’s end-effector remained aligned with the block’s z -axis (upright position).

4 Policy Learning

For the learning algorithm, we used the classic task-parametrized Gaussian Mixture Models (TP-GMM) to train LfD models from the data collected [8]. Due to its ability to generalize well from a limited number of demonstrations, TP-GMMs are widely used in LfD research [28, 30], where task modeling is performed across multiple coordinate frames. The robot task learning begins after the user provided a set of M demonstrations, where each demonstration m comprises T_m state data points. Each state n is represented as

$\xi = (t_n, x_n^T, \omega_n^T)$, where t_n , x_n^T , and ω_n^T represent the time, end-effector position, and orientation in quaternion form, respectively.

In a D -dimensional space, each frame j is defined by an orientation matrix $A \in \mathbb{R}^{D \times D}$ and an origin vector $b \in \mathbb{R}^D$, relative to the global frame. A K -component GMM is fitted to the data in each local frame, with parameters $(\pi_k, \mu_k^{(j)}, \Sigma_k^{(j)})$ representing the prior probabilities, means, and covariance matrices for each component k , respectively. These parameters are estimated using the Expectation-Maximization (EM) algorithm. For trajectory generation, the local models are linearly transformed into the global frame and combined using a product of Gaussian mixtures. Gaussian Mixture Regression is then applied at each time step to retrieve the corresponding trajectory point. For a detailed explanation, refer to Calinon’s work [8].

5 Experimental Design

5.1 Experimental Setup

Robotic System: In this study, we used the 7-DoF Panda robot, a collaborative robotic manipulator, where the robot joints are represented as q_1, q_2, \dots, q_7 as shown in Figure 1. This robot arm is equipped with torque sensors in all 7 joints, making it safe and suitable for human-robot interaction applications. We used the default gravity compensation controller to counteract the robot’s inherent weight, allowing participants to kinesthetically manipulate the robot’s joints without additional force/torque.

Motion Capture System: We employed the OptiTrack motion capture (MoCap) system to track physical interactions of participants in real time. This MoCap system has sub-millimeter accuracy. It allowed us to track the position and orientation of the participants’ hands and infer which of the robot’s joints were manipulated during the demonstrations. The setup comprised nine Prime-13W cameras: five mounted on a ceiling-attached rig and four positioned on tripods around the experimental area. These cameras were strategically spaced and focused to cover the workspace between the participants and the robot. To ensure accurate tracking, we used five sets of markers arranged in unique geometries to monitor the robot’s origin (1x), participants’ hands (2x), the box with buttons (1x), and the insertion bar (1x), as depicted in Figure 2. The combination of trackers and the camera setup was intended to reduce tracking errors. Although minor data losses could occur, we estimate that they would be minimal and did not influence the study’s outcomes. To further limit tracking noise, reflective objects were removed from the experimental area.

Joint Constraint: Each task was performed under two conditions: unconstrained (U) and constrained (C). In the unconstrained condition, all of the robot’s joints were free to be manipulated without any mobility restrictions. In contrast, under the constrained condition, the 6^{th} joint was locked at a default value of 1.571 radians. The goal of locking a joint was to reduce redundancy while keeping the feasibility of 3D motions—hence, reducing the number of available solutions for the given tasks.

To determine which joint to constrain, we first conducted a pilot study ($N = 6$) and complemented it with kinematic analysis. The pilot results showed that the proximal joints, particularly the shoulder joints (1–3), play a major role in defining the robot’s reachable workspace. Constraining any of these joints would drastically reduce reachability and render tasks such as pick-and-place infeasible.




Figure 2: Subfigures (a–b) show a participant kinesthetically demonstrating the button task. Subfigures (c–d) depict the insertion task, with (c) showing block pickup and (d) partial insertion.

For this reason, we focused on the distal joints (5^{th} , 6^{th} , and 7^{th}) as more suitable candidates for reducing redundancy while preserving task feasibility. In typical 6- or 7-DoF manipulators, the wrist joints form a spherical configuration where the last three axes intersect to realize roll–pitch–yaw orientations. A rotation of approximately 90° around the roll joint can interchange the functional roles of pitch and yaw, enabling similar orientations to be achieved with two joints instead of three. Consequently, constraining the mid-wrist joint can reduce redundancy without compromising workspace coverage. Finally, our pilot study also indicated that the 5^{th} and 7^{th} wrist joints should remain unconstrained during the block-insertion task to facilitate grasping blocks with random orientations. Based on the insights from pilot study and kinematic analysis, we selected the 6^{th} joint to be locked during the constrained condition.

5.2 Participants Recruitment

We conducted a prior power ($f = 0.3$, $power = 0.85$, $groups = 2$, $measures = 3$) analysis with G^* Power to calculate the sample size [13], and then recruited a total of 24 participants (6F/18M, age = 26.75 ± 4.22) through advertisements on social media and notice boards across our university campus. Along with demographics questions, we asked two additional questions: *Have you had experience with research/industrial robots before?* and *Do you feel uncomfortable or scared when you are around robots?* The first question served to filter participants based on prior experience with robots while the second one assessed participants’ level of comfort in the presence of a robot. A positive answer to either question would make the person ineligible to participate. The University’s IRB approved the research ethics application (ID 2024-24636-51920-3).

5.3 Experiment Protocol

Upon the participant’s arrival, the experimenter explained the study’s purpose, introduced the MoCap system and robotic manipulator, and described the tasks. Each participant signed a consent form before starting the experiment. Markers were attached to the back of participants’ hands to track their kinesthetic interaction with the robot’s joints. With the gravity compensation controller activated, participants kinesthetically interacted with the robot for few minutes to familiarize themselves with its joints and explore its workspace. The experimenter instructed participants not to move around the experimental setup or manipulate the robot’s

joints using only their fingertips. Instead, they were advised to interact using their hands, with palms in contact with the robot, to ensure accurate tracking by the OptiTrack system. The participants were instructed to avoid transient or unintentional contacts to ensure that detected manipulations reflected deliberate interactions.

Participants were required to teach two tasks sequentially, button-pressing task followed by insertion task, under both conditions (*constrained* and *unconstrained*). Based on the order of conditions, participants were divided into two groups. We used a counterbalanced design where participants were exposed to both conditions in different sequences. The first group, *constrained to unconstrained (C2U)*, comprised 12 participants who began demonstrating tasks with the constrained condition, followed by demonstrating tasks with unconstrained condition. The second group, *unconstrained to constrained (U2C)*, also included 12 participants but completed the tasks in the reverse constraint order, starting with the unconstrained condition and transitioning to the constrained condition. This allowed us to control for order effects. Observation from pilot study data revealed only a negligible effect of task type and its interaction with the constraint mode; therefore, we kept a consistent task order (button task followed by insertion task) across both groups (C2U and U2C). This consistency ensured comparability between conditions, although it did not allow us to explore potential learning or interaction effects.

Participants were instructed to provide five successful demonstrations across both tasks and conditions, resulting in four sessions, one for each combination of task and condition. Following each session, the participant completed a questionnaire to reflect their observations on the task and constraint condition. The experiment took on average one and half hour per participant.

5.4 Performance Metrics

We assessed the human performance using four key measures: (1) mental workload, measured using the NASA-TLX questionnaire; (2) demonstration time, recorded as the total duration of each demonstration; (3) number of failed attempts; and (4) joints manipulation frequency distribution, quantifying the user’s physical interaction with the robot. This last measure is calculated by dividing the number of hand touches on the robot’s joints by the demonstration duration, indicating how frequently a user engaged with each specific joint during the demonstration.

Table 1: Statistical results summary for demonstration time (H1) with standardized coefficients. The terms *CC* and *U2C* represent Constrained-Condition and Group Unconstrained-to-Constrained.

		Button Task		Insertion Task	
		Estimate [95% CI]	Standardized Coeff.	Estimate [95% CI]	Standardized Coeff.
Demo. Time (sec)	<i>CC</i>	26.34 ($p < 0.001$) [16.94, 35.73]	0.78 [0.50, 1.06]	33.22 ($p < 0.001$) [25.45, 40.98]	1.17 [0.89, 1.44]
	<i>U2C</i>	2.85 ($p = 0.73$) [-13.27, 18.99]	0.08 [-0.40, 0.57]	-0.24 ($p = 0.98$) [-11.51, 11.01]	-0.0085 [-0.41, 0.39]
	<i>CC : U2C</i>	1.46 ($p = 0.83$) [-11.74, 14.71]	0.04 [-0.35, 0.44]	-3.71 ($p = 0.51$) [-14.69, 7.24]	-0.13 [-0.52, 0.26]

In addition to the human performance, we evaluated the robot performance based on task completion, defined as the robot’s success in carrying out the task using the learned model. We used each participant data under each *tasks* \times *conditions*, and then tested on real robot once per model. For the button-pressing task, each button is modeled as a sphere with a 3 cm radius based on the button center. A trial is considered successful if the robot’s end-effector tip reaches the target position (i.e., within the goal sphere) without self-collision or collision with the box. For the insertion task, the success is defined as the robot accurately picking up a cuboid block and performing a partial insertion without colliding with itself or the experimental setup.

5.5 Data Collection

We utilized OptiTrack’s Motive 3 software on a Windows machine to track participants’ hand positions and orientations in real time. We employed Unity 2022.3 to retrieve this tracking data via the OptiTrack-Unity plugin and compute interactions with the robot’s joints. This was achieved by accessing the robot’s joint positions via a ROS TCP Endpoint, enabling seamless real-time data transfer between Unity on Windows and the robot controller on a Linux system.

Motion Capturing System: To track the origin position and orientation of the 7-DoF Panda robot, we placed a marker set 17 cm behind the robot’s base center at a height of 1 cm (see Figure 2). This setup allowed precise tracking of hand positions and orientations relative to the robot’s origin. Using Motive 3 software, we transmitted this data to Unity at an average speed of 400 KB/s, ensuring an average frame rate above 30 fps.

Unity Integration: A digital model of the 7-DoF Panda robot was incorporated into Unity to visualize hands-robot interactions. Unity natively computed both hands positions and orientations with respect to the robot’s origin, defined in a left-handed coordinate




Figure 3: Performance metrics of demonstration time and reproduction results using learned model.

system. Hands positions were recorded in x, y , and z coordinates, while orientations were captured as quaternions (x, y, z, w) . In order to determine which robot joint was touched at a given time, we used Unity’s collision detection feature. The hands tracking in the MoCap system was calibrated at the start of each session by instructing participants to place their hands at specific positions on the table where the robot was located. Additionally, the positions and orientations of task-related components, such as the two buttons on the box and the insertion bar, were recorded using additional markers. Data collection of the joint manipulation occurred at a rate of 30 fps to ensure smooth and detailed tracking.

6 Results

This section presents results for both tasks, showing how joint constraint affect user demonstration’s ability and robot performance. We used Cumulative Link Mixed Models (CLMM) and Linear Mixed Models (LMM) to analyze the impact on demonstration time, failed attempts, and physical interaction with robot joints. The CLMM approach was used for ordinal response variables, such as the number of attempts, while the LMM approach was used for continuous response variables, such as the time to complete the demonstrations. These statistical approaches also allowed us to assess the interaction effects among joint constraint and groups, providing insights into potential order or learning effects.

6.1 Demonstration Time (H1)

Button Task: Participants in the *U2C* and *C2U* groups provided a total of 125 and 122 successful demonstrations, respectively. Using LMM with the formula: $Time - Constraint * Group + (1 | Participant)$, we analyzed 247 observations from 24 participants. The results showed that the constrained condition significantly increased the demonstration time by an estimated 26.34 seconds ($p < 0.001$; see Table 1).

Insertion Task: *U2C* and *C2U* participants completed 121 and 120 demonstrations, respectively, yielding 241 observations. The constrained condition led to a significant increase of 33.22 seconds in demonstration time ($p < 0.001$).

Figure 3(a) shows the average demonstration times across both tasks and conditions, with significantly longer demonstrations observed in the constrained condition. Since demonstrations were temporally aligned using DTW, we anticipate longer input durations resulted in proportionally longer robot completion times.

Table 2: Statistical results summary (Estimate [95% CI]) for failed attempts (H2).

	Constrained-Condition	Group U2C	Constrained : Group U2C
Button Task	2.49 ($p < 0.001$) [1.81, 3.17]	0.92 ($p = 0.013$) [0.20, 1.64]	-1.07 ($p = 0.018$) [-1.96, -0.18]
Insertion Task	0.40 ($p = 0.21$) [-0.22, 1.03]	0.56 ($p = 0.11$) [-0.12, 1.24]	-0.72 ($p = 0.14$) [-1.61, 0.17]

6.2 Number of Failed Attempts (H2)

Button Task: In the *U2C* group, participants had 18 (21.68%) failed attempts under the unconstrained condition and 50 (45.47%) under constrained. In contrast, the *C2U* group had 48 (44.41%) failures under constrained and only 02 (3.12%) under unconstrained. CLMM results showed a significant increase of 2.49 (counts) failed attempts under constrained condition ($p < 0.001$). Additionally, the *U2C* group had 0.92 more failures under unconstrained condition than the *C2U* group ($p = 0.013$), as shown in Table 2. The *U2C* group’s failure rate under *constrained condition* was significantly higher, with the group effect negatively moderated (see Table 2), indicating more failures compared to *C2U* in the same condition.

Insertion Task: In the *U2C* group, participants had 15 (19.73%) failures under unconstrained and 12 (16.67%) under constrained. The *C2U* group showed 15 (20%) failures under constrained and 06 (9.09%) under unconstrained. CLMM results showed no significant difference in failure rates between task conditions ($p = 0.21$). The order of joint constraint (*C2U* vs. *U2C*) had no significant impact under constrained, indicating no interaction effect.

6.3 User Interaction with Robot Joints (H3)

Button Task: Joint manipulation frequency across both groups is shown in Figures 4(a) and 4(b), with joint q_5 having the highest and joint q_1 the lowest interaction rates. As detailed in Table 3, under constrained condition, participants interacted more with joints $q_1 - q_4$, particularly q_3 ($est. = 6.1, p < 0.001$), while interactions with q_6, q_7 , and the end-effector decreased, notably q_7 ($est. = -7.6, p < 0.001$). Order of constrain condition significantly affected joint interaction. In the *U2C* group, participants interacted less with q_3 ($est. = -1.9, p = 0.03$) and q_4 ($est. = -2.4, p = 0.003$) than in the *C2U* group under constrained condition. Interaction patterns, particularly for q_7 , also reversed depending on order of condition. Full data is summarized in Table 3.

Insertion Task: Under unconstrained condition, the end-effector had the highest interaction frequency across both groups, while joint q_5 saw the most interaction under constrained condition. Under constrained, participants engaged more with joints $q_1 - q_5$ and significantly less with the end-effector (see Table 3). In the *U2C* group, interaction with joints $q_1 - q_3$ was reduced under constrained condition. Without considering order, we observed a general increase in interaction with these joints.

Figure 5 shows the color-scale representation of joint manipulation frequency across tasks and conditions.

6.4 Mental Workload

To evaluate the perceived difficulty level of each task with both conditions, participants responded to two items using a 7-point Likert scale. The first item was: “Rate how easy you think this task was”. After completing the same task under the alternate constraint condition (i.e., following *U2C* or *C2U*), the second item was: “Rate the difficulty level of this task compared to the previous one”. Additionally, we employed the NASA-TLX (Task Load Index) [15] to assess various aspects of participants’ experiences. The findings, as shown in Figure 6, consistently indicate that participants perceived the tasks as more difficult under the constrained condition, highlighting the impact of joint constraint on perceived difficulty and workload.

6.5 Policy Learning

Successful demonstrations were used to train the TP-GMM models. As a first step, we analyzed each set of demonstrations (i.e., each participant demonstrations data under each *task × condition*) by evaluating lengths of each demonstration, completion time, and ensuring task completion. We removed those demonstrations that failed (self-collision, joint limits, excessive force), incomplete, or unusually long (2x avg. duration). Next, we applied a moving average filter to smooth the trajectories. After pre-processing, we employed the dynamic time warping technique to align the trajectories temporally [39]. TP-GMM models were then trained on these

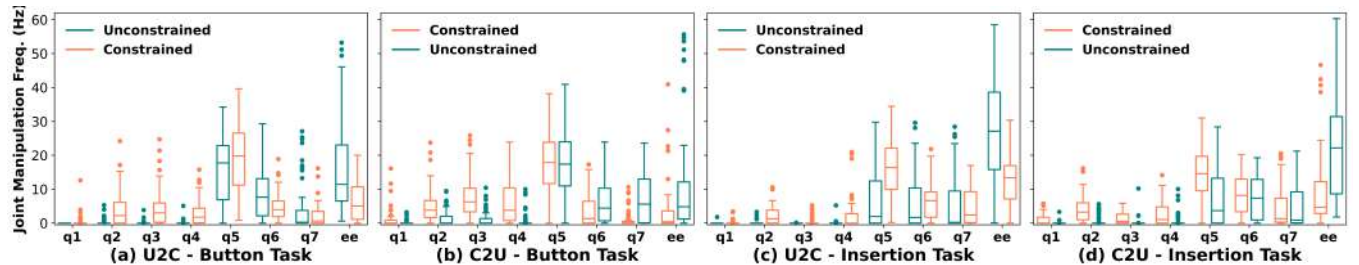


Figure 4: Subfigures (a–b) show joint manipulation frequency for both *U2C* and *C2U* groups during the button-pressing task. Subfigures (c–d) present the corresponding results for the insertion task.




Figure 5: Average joint manipulation frequency; teal and coral indicate unconstrained and constrained condition, respectively.




Figure 6: NASA-TLX scores across six dimensions: Mental (M1), Physical (M2), Temporal (M3), Performance (M4), Effort (M5), and Frustration (M6). Subfigures: (a–b) Button task; (c–d) Insertion task for U2C and C2U groups.

Table 3: Statistical results summary for user interaction (joint manipulation frequency) with the robot’s joints (H3). q_n and ee represent n^{th} robot’s joint (7-DoF) and end-effector, respectively.

Button Task						
	Constrained-Condition		Group U2C		Constrained : Group U2C	
	Estimate [95% CI]	Std. Coeff.	Estimate [95% CI]	Std. Coeff.	Estimate [95% CI]	Std. Coeff.
q_1	1.0 ($p < .001$) [0.48, 1.53]	0.6 [0.3, 0.9]	-0.1 ($p = 0.8$) [-0.9, 0.7]	-0.06 [-0.5, 0.4]	-0.6 ($p = 0.1$) [-1.3, 0.14]	-0.35 [-0.8, 0.08]
q_2	3.8 ($p < .001$) [2.61, 4.9]	0.9 [0.6, 1.2]	-1.1 ($p = 0.3$) [-2.9, 0.8]	-0.24 [-0.7, 0.2]	-0.2 ($p = 0.8$) [-1.9, 1.5]	-0.05 [-0.4, 0.4]
q_3	6.1 ($p < .001$) [4.8, 7.4]	1.2 [0.9, 1.5]	-1.1 ($p = 0.3$) [-3.3, 1.1]	-0.2 [-0.7, 0.2]	-1.9 ($p = 0.03$) [-3.8, -0.2]	-0.4 [-0.8, -0.04]
q_4	5.3 ($p < .001$) [4.2, 6.4]	1.2 [0.9, 1.5]	-0.6 ($p = 0.5$) [-2.6, 1.3]	-0.15 [-0.6, 0.3]	-2.4 ($p = 0.003$) [-3.9, -0.8]	-0.55 [-0.9, -0.2]
q_5	1.4 ($p = 0.3$) [-1.6, 3.9]	0.15 [-0.13, 0.4]	-0.7 ($p = 0.8$) [-6.1, 4.8]	-0.07 [-0.7, 0.5]	2.3 ($p = 0.2$) [-1.3, 5.8]	0.24 [-0.14, 0.6]
q_6	-2.1 ($p = 0.02$) [-4, -0.4]	-0.4 [-0.6, -0.1]	2.2 ($p = 0.2$) [-0.9, 5.2]	0.4 [-0.15, 0.9]	-1.5 ($p = 0.2$) [-3.9, 0.9]	-0.25 [-0.7, 0.2]
q_7	-6.2 ($p < .001$) [-8, -5]	-1.0 [-1.3, -0.8]	-3.4 ($p = 0.07$) [-6.9, 0.15]	-0.55 [-1.1, 0.03]	4.5 ($p < .001$) [2.4, 6.6]	0.74 [0.4, 1.1]
ee	-7.6 ($p < .001$) [-10, -5]	-0.6 [-0.8, -0.4]	4.6 ($p = 0.3$) [-3.3, 12.5]	0.4 [-0.27, 1.0]	-2.1 ($p = 0.3$) [-5.7, 1.6]	-0.17 [-0.5, 0.1]
Insertion Task						
q_1	0.8 ($p < .001$) [0.7, 1.1]	0.9 [0.6, 1.2]	-0.04 ($p = 0.8$) [-0.5, 0.4]	-0.05 [-0.5, 0.4]	-0.6 ($p = 0.001$) [-1.0, -0.3]	-0.7 [-1.1, -0.3]
q_2	3.8 ($p < .001$) [2.9, 4.6]	1.2 [0.9, 1.5]	-0.4 ($p = 0.5$) [-1.7, 0.9]	-0.13 [-0.6, 0.3]	-1.6 ($p = .007$) [-2.7, -0.4]	-0.5 [-0.9, -0.1]
q_3	1.2 ($p < .001$) [0.8, 1.7]	0.9 [0.6, 1.2]	-0.24 ($p = 0.4$) [-0.8, 0.3]	-0.2 [-0.6, 0.2]	-0.8 ($p = 0.01$) [-1.4, -0.1]	-0.5 [-0.9, -0.1]
q_4	1.7 ($p < .001$) [0.9, 2.9]	0.6 [0.3, 0.9]	-0.5 ($p = 0.5$) [-2.3, 1.2]	-0.2 [-0.7, 0.3]	0.6 ($p = 0.4$) [-0.8, 2.0]	0.2 [-0.2, 0.6]
q_5	7.4 ($p < .001$) [5.3, 9.5]	0.8 [0.5, 1.0]	-0.8 ($p = 0.9$) [-6.3, 4.7]	-0.08 [-0.8, 0.5]	2.2 ($p = 0.2$) [-0.7, 5.2]	0.2 [-0.08, 0.5]
q_6	0.8 ($p = 0.4$) [-1.2, 2.8]	0.13 [-0.2, 0.45]	-1.2 ($p = 0.5$) [-4.3, 1.9]	-0.2 [-0.7, 0.3]	-0.4 ($p = 0.7$) [-3.3, 2.4]	-0.07 [-0.5, 0.4]
q_7	-0.2 ($p = 0.8$) [-2.3, 1.9]	-0.03 [-0.4, 0.3]	1.2 ($p = 0.5$) [-1.9, 4.2]	0.2 [-0.3, 0.7]	-1.1 ($p = 0.4$) [-4.1, 1.8]	-0.2 [-0.6, 0.3]
ee	-13.7 ($p < .001$) [-17, -10]	-0.9 [-1.1, -0.7]	5.1 ($p = 0.2$) [-2.8, 13.1]	0.3 [-0.2, 0.9]	-1.4 ($p = 0.6$) [-6.3, 3.5]	-0.09 [-0.4, 0.2]

filtered set of demonstrations for the given task parameters (i.e., matrix A and vector b) and tested on the real robot to evaluate robot performance. The number of Gaussian components (K) was selected using Bayesian Information Criteria, with diagonal covariance regularization ($1e-5$) and EM parameters (min 5, max 100 steps, tolerance $1e-5$). The Franka ROS2 default joint trajectory controller was used.

Figure 3(b) illustrates that the presence of a joint constraint has a notable impact on the robot’s performance. Each bar in the plot represents the number of successful models or rollouts out of a total of 12, corresponding to demonstrations provided by 12 individual users in each group. The results indicate a notable difference between the two experimental conditions. For both manipulation

tasks, the number of successful attempts was consistently higher when the model was trained on demonstrations collected under the unconstrained condition, whereas the success rate decreased when demonstrations were collected under the constrained condition.

7 Discussion

This work examined how robot redundancy shapes novice users’ experiences during kinesthetic teaching. Rather than treating redundancy as a purely kinematic property, our findings highlight its role as an interaction design variable that directly influences how users explore, understand, and demonstrate robot motion. Across tasks,

higher degrees of redundancy supported more efficient and less demanding demonstrations, while also reshaping users' physical engagement with the robot and the consistency of their demonstrations.

Taken together, the results suggest that providing users with greater kinematic freedom can lower barriers to effective demonstration by reducing physical and cognitive effort. However, this same flexibility introduces new challenges related to coordination, collision avoidance, and execution fidelity. In the following section, we discuss how redundancy influenced user physical interaction, demonstration diversity, and failure patterns, before reflecting on the broader implications for LfD.

User Physical Interaction: User physical interaction provided insight into how redundancy affected participants' reasoning about the demonstration through their hands. Under the constrained condition, participants primarily manipulated the lower joints (q_2 – q_5), with q_5 receiving the most interaction. Participants who started with the constrained condition initially expected the unconstrained condition to require similar full-joint manipulation; however, they learned that focusing on joints closer to the end-effector enabled more efficient task completion, leading to a gradual shift in hand contact points toward these joints (see Figure 4). The change in user interaction patterns over time has also been reported by Aliasghari et al. [3]. We also observed significant differences in hand contact sequences between failed attempts and successful demonstrations.

Demonstrations Diversity: In our previous work [6], we observed that demonstration diversity influenced the robot task performance. Under the constrained condition, greater variability in user demonstrations made learning more challenging for the robot, requiring additional demonstrations to achieve satisfactory performance. In contrast, demonstrations collected under the unconstrained condition were more consistent, resulting in improved robot task performance. Although it may appear counterintuitive that reduced redundancy leads to greater demonstration diversity, limiting redundancy increases task difficulty for novice users by reducing valid solutions and obscuring the desired configuration. These results are echoing the importance of previous LfD research [36]. In addition, our work suggests a possible design pathway for diverse demonstrations [37].

Learning Through Practice: Learning effects significantly influenced the number of failed attempts in the button-pressing task: participants who began with the easier, unconstrained condition demonstrated improved performance when transitioning to the more challenging constrained condition. Most failures stemmed from joint-limit violations, excessive force application, and self-collisions. In contrast, no notable differences were observed in the insertion task, likely because participants became progressively familiar with the robot's joint configuration and task requirements over the course of the study, consistent with prior findings [2, 3].

Challenges of Higher Robot Redundancy: Our user study revealed two key challenges associated with higher robot redundancy during kinesthetic teaching. First, *higher redundancy increased the likelihood of violating inherent robot constraints*—such as self-collision—because additional links introduced more opportunities for *interference during demonstration* (e.g., we observed frequent collisions between links 5 and 2 during the insertion task). Second, novice users found it more difficult to manage singular

configurations due to the larger configuration space. These observations highlight the need for assistive feedback systems to support novice users during kinesthetic demonstrations of redundant robots. We also observed implications for execution, as reproducing demonstrations in redundant systems required closer alignment in both task and joint spaces. These findings suggest that, although redundancy offers flexibility, it also introduces coordination and modeling challenges that affect both teaching and execution in LfD. These results are supported from the literature [19, 20].

Limitations

All participants were university students, with a predominantly male sample, reflecting our focus on recruiting LfD novices rather than achieving gender balance. Future work will include industry professionals, such as caregivers and healthcare workers, to broaden the applicability of our findings. In addition, while order effects were observed in the button task but not the insertion task (H2), studies with larger participant samples are needed to better assess learning effects.

We evaluated two manipulation tasks—button pressing and cuboid block insertion—chosen for their comparability with prior work [10, 35]. Future work will explore additional constraint types [16, 27], such as maintaining an upright cup during transport, to evaluate and compare their effects against the current results.

Finally, our study was conducted using a commercially available 7-DoF Franka Emika Panda robot, whose shoulder–elbow–wrist morphology is representative of many collaborative manipulators. However, our findings are based on a single redundancy level and may not fully generalize to robots with fewer or substantially more degrees of freedom. While similar effects may arise in lower-DoF systems for sufficiently constrained tasks, or be amplified in highly redundant robots, future work is needed to systematically examine how redundancy-related effects scale across different robot morphologies.

8 Conclusion

This work examined how robot kinematic redundancy shapes both novice users' kinesthetic teaching strategies and the resulting robot learning performance. Across tasks, constraining redundancy increased users' mental workload, prolonged demonstration time, and led to more frequent failures. In contrast, providing higher redundancy encouraged users to interact more locally around the end-effector, resulting in more efficient demonstrations and consistently stronger robot performance. Together, these findings demonstrate that kinematic redundancy is not merely a mechanical property, but a key interaction factor that influences how users explore the robot's configuration space and how effectively robots learn from human demonstrations.

Key Takeaways:

- Providing novice users with higher robot redundancy (more DoF than strictly required) helps them identify suitable joint configurations within a larger feasible solution space.
- The sequence of hand contact points contains valuable information that can be used to anticipate and prevent potential failures.

Supplementary Materials

The dataset *InteractLFD* is available in the supplementary materials to facilitate replication and further research.

References

- [1] Alessandro Adami, Aris Synodinos, Matteo Iovino, Ruggero Carli, and Pietro Falco. 2025. Learning Stack-of-Tasks Management for Redundant Robots. arXiv:2508.10780. Retrieved from <https://arxiv.org/abs/2508.10780>.
- [2] Pourya Aliasghari, Moojan Ghafurian, Chrystopher L. Nehaniv, and Kerstin Dautenhahn. 2022. Kinesthetic Teaching of a Robot over Multiple Sessions: Impacts on Speed and Success. In *International Conference on Social Robotics (ICSR 2022)*. Springer, 160–170.
- [3] Pourya Aliasghari, Moojan Ghafurian, Chrystopher L. Nehaniv, and Kerstin Dautenhahn. 2024. How Non-experts Kinesthetically Teach a Robot over Multiple Sessions: Diversity in Teaching Styles and Effects on Performance. *International Journal of Social Robotics* 16 (2024), 2079–2105.
- [4] Saleema Amershi, Maya Cakmak, W. Bradley Knox, and Todd Kulesza. 2014. Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine* 35, 4 (2014), 105–120.
- [5] Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57, 5 (2009), 469–483.
- [6] Muhammad Bilal, Nir Lipovetzky, Denny Oetomo, and Wafa Johal. 2024. Beyond Success: Quantifying Demonstration Quality in Learning from Demonstration. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 5120–5127.
- [7] Aude Billard, Sylvain Calinon, Ruediger Dillmann, and Stefan Schaal. 2008. Robot Programming by Demonstration. *Springer Handbook of Robotics* (2008), 1371–1394.
- [8] Sylvain Calinon. 2016. A tutorial on task-parameterized movement learning and retrieval. *Intelligent Service Robotics* 9, 1 (2016), 1–29.
- [9] Sylvain Calinon, Florent Guenter, and Aude Billard. 2007. On Learning, Representing, and Generalizing a Task in a Humanoid Robot. *IEEE Transactions on Systems, Man and Cybernetics, Part B (Cybernetics)* 37, 2 (April 2007), 286–298.
- [10] Jiahao Chen, D. Antony Chacon, Muhammad Bilal, Qiushi Zhou, and Wafa Johal. 2024. Mr.LfD: A Mixed Reality Interface for Robot Learning from Demonstration. In *Proceedings of the 36th Australasian Conference on Human-Computer Interaction (OzCHI '24)*. ACM, 275–285.
- [11] Sonia Chernova and Andrea L. Thomaz. 2014. *Robot Learning from Human Teachers*. Springer Cham.
- [12] Glen Chou, Dmitry Berenson, and Necmiye Ozay. 2021. Uncertainty-Aware Constraint Learning for Adaptive Safe Motion Planning from Demonstrations. In *Proceedings of the 2020 Conference on Robot Learning*. PMLR 155, 1612–1639.
- [13] Franz Faul, Edgar Erdfelder, Axel Buchner, and Albert-Georg Lang. 2009. Statistical power analyses using G* Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods* 41, 4 (2009), 1149–1160.
- [14] Kerstin Fischer, Franziska Kirstein, Lars Christian Jensen, Norbert Kruger, Kamil Kuklinski, Maria Vanessa aus der Wieschen, and Thiusius Rajeeth Savarimuthu. 2016. A comparison of types of robot control for programming by Demonstration. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 213–220.
- [15] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Advances in Psychology* 52 (1988), 139–183.
- [16] Bradley Hayes and Brian Scassellati. 2014. Discovering task constraints through observation and active learning. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 4442–4449.
- [17] Erin Hedlund, Michael Johnson, and Matthew Gombolay. 2021. The effects of a robot's performance on human teachers for learning from demonstration tasks. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI '21)*. ACM, 207–215.
- [18] Matthias Hirschmanner and Markus Vincze. 2022. Robot Learning from Humans in Everyday Life Scenarios. In *Trust in Robots*. TU Wien Academic Press, 179–199.
- [19] Tinghe Hong, Weibing Li, and Kai Huang. 2024. A reinforcement learning enhanced pseudo-inverse approach to self-collision avoidance of redundant robots. *Frontiers in Neurorobotics* 18 (2024), 1375309.
- [20] Kui Hu, Jiwen Zhang, and Dan Wu. 2024. Learning from demonstration for 7-DOF anthropomorphic manipulators without offset via analytical inverse kinematics. *Neurocomputing* 598 (2024), 128036.
- [21] Aleksandra Kalinowska, Ahalya Prabhakar, Kathleen Fitzsimons, and Todd Murphey. 2021. Ergodic imitation: Learning from what to do and what not to do. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 3648–3654.
- [22] Andrey Kurenkov, Baris Akgun, and Andrea L. Thomaz. 2015. An evaluation of GUI and kinesthetic teaching methods for constrained-keyframe skills. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 3608–3613.
- [23] Michael Laskey, Caleb Chuck, Jonathan Lee, Jeffrey Mahler, Sanjay Krishnan, Kevin Jamieson, Anca Dragan, and Ken Goldberg. 2017. Comparing human-centric and robot-centric sampling for robot deep learning from demonstrations. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 358–366.
- [24] Xing Li and Oliver Brock. 2022. Learning from demonstration based on environmental constraints. *IEEE Robotics and Automation Letters* 7, 4 (2022), 10938–10945.
- [25] Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martin-Martin. 2022. What Matters in Learning from Offline Human Demonstrations for Robot Manipulation. In *Proceedings of the 5th Conference on Robot Learning*. PMLR, 164, 1678–1690.
- [26] Carl Mueller, Jeff Venicx, and Bradley Hayes. 2018. Robust Robot Learning from Demonstration and Skill Repair Using Conceptual Constraints. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 6029–6036.
- [27] Carl L. Mueller and Bradley Hayes. 2020. Safe and Robust Robot Learning from Demonstration through Conceptual Constraints. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20)*. ACM, 588–590.
- [28] Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J. Andrew Bagnell, Pieter Abbeel, and Jan Peters. 2018. An Algorithmic Perspective on Imitation Learning. *Foundations and Trends in Robotics* 7, 1-2 (2018), 1–179.
- [29] Baiyu Peng and Aude Billard. 2024. Positive-Unlabeled Constraint Learning for Inferring Nonlinear Continuous Constraints Functions From Expert Demonstrations. *IEEE Robotics and Automation Letters* 10, 2 (2024).
- [30] Affan Pervez and Dongheui Lee. 2018. Learning task-parameterized dynamic movement primitives using mixture of GMMs. *Intelligent Service Robotics* 11, 1 (2018), 61–78.
- [31] Ornnalin Phajit, Claude Sammut, and Wafa Johal. 2023. User Interface Interventions for Improving Robot Learning from Demonstration. In *Proceedings of the 11th International Conference on Human-Agent Interaction (HAI '23)*. ACM, 152–161.
- [32] Harish Ravichandar, Athanasios S. Polydoros, Sonia Chernova, and Aude Billard. 2020. Recent Advances in Robot Learning from Demonstration. *Annual Review of Control, Robotics, and Autonomous Systems* 3, 1 (2020), 297–330.
- [33] Stéphane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. 2011. A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*. PMLR, 15, 627–635.
- [34] Maram Sakr, H. F. Machiel Van der Loos, Dana Kulic, and Elizabeth Croft. 2021. What Makes a Good Demonstration for Robot Learning Generalization?. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI '21)*. ACM, 607–609.
- [35] Maram Sakr, Zexi Jesse Li, H. F. Machiel Van der Loos, Dana Kulic, and Elizabeth A. Croft. 2022. Quantifying Demonstration Quality for Robot Learning and Generalization. *IEEE Robotics and Automation Letters* 7, 4 (2022), 9659–9666.
- [36] Maram Sakr, Juyan Zhang, H. F. Machiel Van der Loos, Dana Kulic, and Elizabeth Croft. 2025. Consistency Matters: Defining Demonstration Data Quality Metrics in Robot Learning from Demonstration. *ACM Transactions on Human-Robot Interaction* 15, 2, Article 34 (2025), 31 pages.
- [37] Maram Sakr, Zhikai Zhang, Benjamin Li, Haomiao Zhang, H. F. Machiel Van der Loos, Dana Kulic, and Elizabeth Croft. 2023. How Can Everyday Users Efficiently Teach Robots by Demonstrations? *ACM Transactions on Human-Robot Interaction* 14, 4, Article 74 (2023), 22 pages.
- [38] Maha Salem, Gabriella Lakatos, Farshid Amirabdollahian, and Kerstin Dautenhahn. 2015. Would You Trust a (Faulty) Robot?: Effects of Error, Task Type and Personality on Human-Robot Cooperation and Trust. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '15)*. ACM, 141–148.
- [39] Pavel Senin. 2008. Dynamic Time Warping Algorithm Review. *University of Hawaii at Manoa, Honolulu, USA*. (2008), 23 pages.
- [40] Weiyong Si, Zhehao Jin, Zhenyu Lu, Ning Wang, and Chenguang Yang. 2024. A Stable Guidance Method for Teleoperation-based Robot Learning from Demonstration. In *2024 IEEE 20th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2376–2381.
- [41] Weiyong Si, Ning Wang, and Chenguang Yang. 2021. A review on manipulation skill acquisition through teleoperation-based learning from demonstration. *Cognitive Computation and Systems* 3, 1 (2021), 1–16.
- [42] Sebastian Wrede, Christian Emmerich, Ricarda Grünberg, Arne Nordmann, Agnes Swadzba, and Jochen Steil. 2013. A user study on kinesthetic teaching of redundant robots in task and configuration space. *Journal of Human-Robot Interaction* 2, 1 (2013), 56–81.

Received 2025-09-30; Accepted 2025-12-01