

Wrapping Haptic Displays Around Robot Arms to Communicate Learning

Antonio Alvarez Valdivia[✉], *Graduate Student Member, IEEE*, Soheil Habibian[✉], *Graduate Student Member, IEEE*, Carly A. Mendenhall[✉], Francesco Fuentes[✉], Ritish Shailly[✉], *Graduate Student Member, IEEE*, Dylan P. Losey[✉], and Laura H. Blumenschein[✉]

Abstract—Humans can leverage physical interaction to teach robot arms. As the human kinesthetically guides the robot through demonstrations, the robot learns the desired task. While prior works focus on how the robot learns, it is equally important for the human teacher to understand what their robot is learning. Visual displays can communicate this information; however, we hypothesize that visual feedback alone misses out on the *physical connection* between the human and robot. In this paper we introduce a novel class of *soft haptic displays* that wrap around the robot arm, adding signals without affecting that interaction. We first design a pneumatic actuation array that remains flexible in mounting. We then develop single and multi-dimensional versions of this wrapped haptic display, and explore human perception of the rendered signals during psychophysics tests and robot learning. We ultimately find that people *accurately distinguish* single-dimensional feedback with a Weber fraction of 11.4%, and identify multi-dimensional feedback with 94.5% accuracy. When physically teaching robot arms, humans leverage the single- and multi-dimensional feedback to provide better demonstrations than with visual feedback: our wrapped haptic display decreases teaching time while increasing demonstration quality. This improvement depends on the location and distribution of the wrapped haptic display.

Index Terms—Haptic display, learning from demonstration, tactile devices.

I. INTRODUCTION

IMAGINE teaching a rigid robot arm to clean objects off a table (see Fig. 1). One intuitive way for you to teach this robot is through *physical interaction*: you push, pull, and guide

Manuscript received 14 July 2022; revised 19 December 2022; accepted 20 January 2023. Date of publication 30 January 2023; date of current version 21 March 2023. This work was supported in part by NSF under Grants 2129201 and 2129155 and in part by the NSF Graduate Research Fellowship Program. This paper was recommended for publication by Associate Editor Matteo Bianchi. (*Corresponding author: Antonio Alvarez Valdivia*.)

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by Virginia Tech IRB and Purdue University IRB, Application Nos. 20-755 and 2021-1283, respectively.

Antonio Alvarez Valdivia, Carly A. Mendenhall, Francesco Fuentes, and Laura H. Blumenschein are with the Mechanical Engineering, Purdue University, West Lafayette, IN 47901 USA (e-mail: alvar168@purdue.edu; cmenden@purdue.edu; francescofuentes@gmail.com; lhblumen@purdue.edu).

Soheil Habibian, Ritish Shailly, and Dylan P. Losey are with the Collaborative Robotics Lab (Collab), Department of Mechanical Engineering, Virginia Tech, Blacksburg, VA 24061 USA (e-mail: habibian@vt.edu; rshailly@vt.edu; losey@vt.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TOH.2023.3240400>, provided by the authors.

Digital Object Identifier 10.1109/TOH.2023.3240400

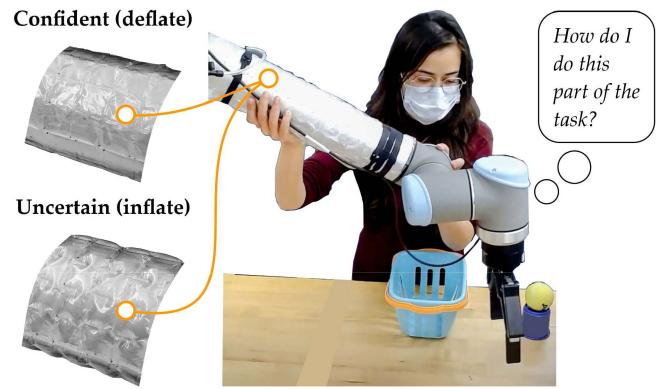


Fig. 1. Human physically teaching a robot arm. We wrap a soft pneumatic display around the arm and render haptic signals by controlling the pressure of the display. The robot learner leverages this haptic display in real-time to communicate the parts of the task that it is confident about, as well as the parts where it is uncertain and needs additional guidance.

the arm along each part of the task. Of course, the robot may not learn everything from a single demonstration, and so you show multiple examples of closing shelves, removing trash, and sorting objects. As you kinesthetically teach the robot you are faced with two questions: i) has the robot learned enough to clear the table by itself and ii) if not, what features of the task is the robot still uncertain about?

While existing work enables robots to learn from physical human interaction [1], [2], [3], [4], having the robot effectively provide *real-time feedback* to human teachers remains an open problem. Ideally, this feedback should not be cumbersome or distracting (i.e., the human must be able to focus on guiding the robot) and should be easily interpretable (i.e., the human must be able to clearly distinguish between signals). These requirements present a tradeoff for haptic feedback as human fingertips provide the densest mechanoreceptors, but placing rigid devices at the hand will impact task performance. Recent research has created communication channels by instead wrapping *haptic devices* around the human's arm [5], [6], [7], but locating feedback at unrelated locations on the human's body can create a disconnect with the task.

Our insight is that — instead of asking the human teacher to wear a feedback device or watch a computer monitor —

We can take advantage of the preexisting physical contact between the human and robot through slim form-factor soft haptic displays that wrap around the robot arm.

Accordingly, in this paper we develop and analyze wrapped haptic displays for communicating robot learning based on soft robotic principles. We distribute these soft displays along rigid robot arms so that the human can physically interact with the robot to demonstrate a task while simultaneously perceiving the robot's feedback. We actively control the *pressures* of the pneumatic display to render where in the task and what features of the task the robot is *uncertain* about: the display inflates for regions and features of the task where the robot is unsure about its actions (and needs additional human teaching), and deflates where the robot is confident about the task (and does not need any additional human guidance). Our hypothesis is that — because the soft wrapped display creates a channel for communication on any surface without impacting the task — humans will be able to more intuitively and robustly use this feedback with a greater level of focus compared to other feedback modalities. We experimentally demonstrate that this pressure-based feedback enables humans i) to determine whether the robot has learned enough to be deployed and ii) to identify parts of the task where kinesthetic teaching is still required. Additionally, we demonstrate the importance of the location and distribution of the feedback on the robot arm for creating this improvement. An interface that provides such intuitive and robust feedback about a robot's latent state in real time could potentially be introduced to factory settings that feature learning systems. Such displays would allow everyday workers with no programming or robot-related training to interact with learning systems.

Parts of this work were previously published in [8], which presented the experimental results for our one degree-of-freedom (DoF) haptic display. This current paper builds on that initial research by demonstrating the design, analysis, and application of *multi-DoF* spatial signals localized or distributed along the robot arm, as well as a follow-up analysis of the 1-DoF device. Overall, we make the following contributions:

Developing Wrapped Haptic Display: We design and build a compliant pneumatic haptic device that wraps around and conforms to the robot, providing haptic stimuli that are localized to the robot arm and distributed along its geometry. This device is manufactured using soft, flexible pouches that render haptic signals through pressure.

Measuring User Ability to Perceive Wrapped Displays: We perform a psychophysics study to find the range of pressures that humans can distinguish. We report the just noticeable difference (JND) for pressures rendered by the soft display.

Applying Wrapped Displays to Communicate Learning: We ask participants to kinesthetically teach a robot arm while the robot provides real-time feedback about its learning. We map the robot's uncertainty to the pressure of our wrapped display. Compared to a graphical user interface, wrapped haptic display feedback leads to faster and more informative human teaching, and is subjectively preferred.

Extension on Wrapped Displays to Multiple Degrees of Freedom: We generalize the wrapped display design to create multi-degree of freedom displays. These displays can be configured to fit different robotic manipulator geometries and to change the interconnections between pouches.

Measuring Effect of Display Distribution on User Perception: We perform a psychophysics study to understand how the spatial distribution of the wrapped haptic display signals affects the accuracy and speed of signal identification. We demonstrate a tradeoff between speed of identification and accuracy as signals are spread further apart.

Measuring Effect of Display Distribution of Multi-Degree of Freedom Displays for Communicating Learning: We repeated the kinesthetic teaching task with three degree of freedom displays, confirming that users still improve demonstrations over baseline as signal complexity increases. When comparing different options to distribute feedback in 3-DoF displays, users performed better with and subjectively preferred wrapped display layouts where all feedback was displayed in the small area where contact was already occurring instead of distributed in larger areas along the robot arm.

II. RELATED WORK

In this paper we introduce a wrapped haptic display for communicating robot learning in real-time during physical human-robot interaction. We build on previous research for kinesthetic teaching, haptic interfaces, and soft displays.

Kinesthetic Teaching: Humans can show robot arms how to perform new tasks by physically demonstrating those tasks on the robot [1], [2], [3], [4]. As the human backdrives it, the robot records the states that it visits and the human's demonstrated actions at those states. The robot then learns to imitate the human's actions and perform the task by itself [9]. One important output of the learning process is the robot's *uncertainty* about the task. The uncertainty can be measured as the robot's *overall* confidence in *what* to do at different states [10], [11], or also measure the robot's confidence on *how* to perform the task [12], [13], [14], [15]. In this paper we explore how robots should *communicate* their learning uncertainty back to the human teacher. Keeping the human up-to-date with what the robot has learned builds trust and improves teaching [16]. Outside of physical human-robot interaction, prior research has developed multiple non-haptic modalities to communicate robot learning and intent: these include robot motion [17], graphical user interfaces [18], projections into the environment [19], and augmented reality headsets [20]. Within a teleoperation domain, our recent work suggests that *haptic interfaces* are particularly effective at communicating low-dimensional representations of robot learning [6]. Here we will leverage these results to develop a real-time feedback interface *specifically for* kinesthetic teaching.

Haptics for Information Transfer: When using haptics to communicating features of robot learning, the type of information being transferred is important to consider. While haptic devices have a general goal of stimulating the human sense of touch, haptics has also been applied to communicate *robot intent* or similar social features. For instance — when studying how humans and robots should interact in shared spaces — prior works have used haptics to explicitly convey the robot's intended motion or actions [5], [21], [22], [23]. Recent work has shown that, given appropriate context, complex human-to-human social

touch signals, like stroking [24], [25], hugging [26], dancing [27], and emotional communication [28], [29], [30], can be replicated and understood in a wearable format. Some other work has shown the use of haptic interfaces for high information tasks, like assisting navigation through rendering patterns with a certain meaning [31] or using haptic signals to reduce the distraction from visual displays in human-robot collaborative task scheduling [32]. Lastly, work has shown communicating alerts with different urgency levels in car driving [33], [34] and communicating contact events in teleoperation and AR/VR through hand-held haptic devices [35], [36]. These past works suggest that a wide range of social and collaborative information can be transferred using haptics with appropriate design of the interface and signals.

Soft Haptic Devices: Soft haptic devices offer an attractive option for human-robot communication due to their compliance and adaptability, through the *flexibility* of the interface or the *compliance* of the actuators. A range of compliant actuation types have been used for haptic devices: pneumatic actuation [37], [38], [39], shape memory alloys [25], dielectric elastomers [40], and fluidic elastomers [41]. Soft wearable fingertip devices have targeted a range of stimuli in the skin [42], such as vibrations [43], [44], [45], indentation [46], [47], [48], skin-stretch [49], [50], or combinations of those [51], [52], [53], [54]. Soft haptic approaches scale easily to increased areas of stimulation; haptic surfaces using arrays of actuators and sensors show scaling to fit varied areas. These developments have typically used rigid elements, such as NFC electronics [55], thin-film strain sensors [56], and piezo films [57], embedded in cloth and silicone layers to create bi-directional interfaces. These rigid elements can limit the flexibility of the device, and lead to issues with wear over time and comfort. Some tabletop haptic displays have used pneumatically actuated soft composite materials [58] or pneumatic actuation with particle jamming [59] to control the shape and mechanical properties of surfaces, leading to complex signals and comfortable interaction.

Soft haptic interfaces also support a range of device types distinguished by method of interaction: graspable, wearable, or touchable [60]. This method can have a large impact on the usability of the devices. Fingertip worn devices provide high fidelity and interpretable signals [42], [45], [51]. These devices are popular for virtual reality where physical contact with the real world is unlikely; in other applications they can reduce the user's ability to use their hands. This motivates wearable devices for other body areas, such as hand dorsal [61], [62], wrists/forearms [25], [37], [38], or gloves that cover the whole hand [63]. Our recent work has demonstrated the use of inflatable pouches to create wearable haptic interfaces that provide feedback to humans in the form of distributed spatial inflation [39]. Placing haptic signals *directly on the human body* enables the human to move about the space while receiving real-time feedback; but, as feedback is moved away from the fingertip and physically separated from the task, it potentially requires additional mental energy to decode the intended message. A different approach has focused on developing touchable haptic surfaces consisting of arrays of actuators and sensors [55], [56], [57]. These devices use the fingertip mechanoreceptors

without burdening the user's hands. Soft touchable displays allow installation of haptic interfaces in common touch areas, like car steering wheels [33]. While not a haptic display, recent work showed pneumatic actuators wrapped around robot arms to visualize the weight load carried by the robot [64]. Based on this past work, we target a touchable device placed at the point of human-robot interaction, and use soft pneumatic actuation to maximize the flexibility and transparency of the display.

III. DEVELOPING A WRAPPED HAPTIC DISPLAY

We first aim to design a soft haptic display that can wrap around a robot arm, conforming to the surface and adding a haptic interface to existing points of contact between the human and robot. This section describes the identification of three critical requirements (low volume, fast inflation, and textured surface). With these requirements, we outline two designs built on the same underlying principle: a 1-DoF display with a large contact area and a N-DoF design with multiple, reduced-width, "ring" sleeves. Finally, we describe the how these wrapped haptics displays were implemented.

A. Requirements

While designing the wrapped haptic display concept we considered three key requirements to improve operation and the haptic sensation: low volume, fast inflation, and textured surface. First, the display should function without using large volumes of air or static materials, keeping the display flexible enough to easily wrap around objects like the robot arm. Limiting the volume of air also aids in fast inflation and deflation, allowing faster changes in the produced signals. Additionally, we aimed to create an inflatable surface with textured tactile sensations. We believe that a textured surface helps users to quickly identify pressure changes in the display by exploring surface features.

B. Soft Haptic Display Concept

To address these requirements, we use thin, heat-sealable, and inextensible materials that are formed into air-tight pouches and heat-sealed with patterns. The heat sealed patterns subdivide the bag, limiting the volume, adding texture, and keeping the surface flexible when inflated. The final design consists of an array of 2.54 cm square-shaped cells patterned into a low-density polyethylene (LDPE) plastic tube, sealed using a linear heat sealer (H-89 Foot-Operated Impulse Sealer, ULINE). A repeated and homogeneous pattern with gaps in the seals between cells allows for smooth and fast inflation. The square-array design is shown in Fig. 2 in two form factors. The dimensions and shape of the display can be varied to fit different applications and surfaces. A unit of a soft wrapped haptic display consists of one or more pouches attached to the same pressure source, forming a degree of freedom (1-DoF). Multiple degrees of freedom can be attached together to form an N degree of freedom (N-DoF) display. Given this general description of the soft haptic display, we will next describe the specific 1-DoF and 3-DoF displays used in testing.

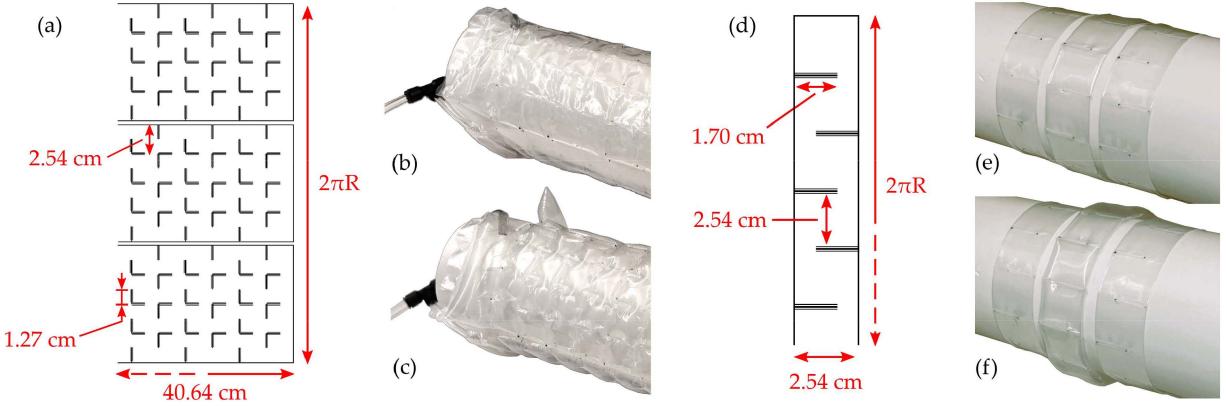


Fig. 2. Overview of the soft wrapped haptic display design. (a) Detailed view of square-cell array implemented in the 1-DoF sleeve display [8]. The thick lines indicate places where the LDPE plastic tube was heat-sealed. The sleeve is composed of 3 pouches taped together to form a sleeve with circumference $2\pi R$. The sleeve display is shown (b) deflated and (c) inflated. (d) Detailed view of the square-cell array implemented in the 1-DoF ring display. Grouping multiple individually-actuated pouches, placed side by side, forms a N-DoF wrapped haptic displays. A 3-DoF ring display is shown in two states: (e) deflated and (f) one of the DoF (center) is actuated, while the others are deflated.

C. Large Surface Display

The 1-DoF soft wrapped haptic display was made from a set of three connected pouches made from a 10.16 cm flat-width LDPE tube (S-5522, ULINE). The LDPE tube matched the length of one section of a UR-10 robotic arm (40.64 cm). The heat-sealed lines are 1.27 cm long, alternated in rows and columns to create the 2.54 cm-squares (Fig. 2(a)). Through-wall straight connectors (5779K675, McMaster-Carr) were attached to one side of each bag strip to allow for individual inflation. The display was made of three bags taped together using viscoelastic adhesive tape (MD-9000, Marker Tape) to construct a sleeve that matched, and entirely wrapped, the cylindrical surface (Fig. 2(a)–(c)). The bags were connected using tee-adapters and inflated using a single pressure line (i.e. a 1-DoF soft wrapped haptic display). The 1-DoF soft wrapped haptic display can be inflated quickly; pressures above 1.5 psi (10.43 kPa) inflate in 0.86 seconds, the pressure can be changed from 1 to 3 psi (6.89 to 20.68 kPa) in 0.72 seconds, and deflate back to 1 psi in 0.18 seconds. The display can operate to a maximum of 3.5 psi (24.13 kPa). Above that pressure the heat-sealed edges begin to tear, producing leaks.

D. Multi-Degree of Freedom Display

We next increased the signal complexity while maintaining the design requirements by building on the 1-DoF design. We grouped multiple, individually-actuated pouches to form a N-DoF wrapped haptic display, as shown in Fig. 2(d)–(f). Each pouch consisted of a 2.54 cm flat-width LDPE tube (S-11155, ULINE), cut to fit the circumference ($2\pi R$) of a segment on a robot arm and form a ring-shaped haptic display. Grommets were placed in the ends of the displays, and elastic bands tied the device around the cylindrical surface. The pattern is modified from the 1-DoF displays to better fit the LDPE tubing. The 2.54 cm square cell grid was achieved by heat sealing 1.7 cm long lines across the length of the tube, alternating sides (Fig. 2(d)). Silicon tubing (0.66 cm OD) was attached to an end of the individual ring display to inflate. For the studies in

Sections VI and VII, three ring displays were placed side by side. Separation between pouches (1.9 cm) is added to assist in making the identification of each DoF easier. Since the N-DoF display segments cover a smaller area, it is easier to mount them in different places of the robot arm. Additionally, since the width of these displays is smaller than the 1-DoF sleeve design, they have smaller volume when inflated and resist higher pressures, producing faster inflation/deflation speeds. These ring-shaped soft wrapped haptic displays can be inflated to pressures above 1.5 psi (10.43 kPa) in 0.55 seconds, and withstand a maximum of 5 psi (34.48 kPa). Switching from 1 to 3 psi (6.89 to 20.68 kPa) occurred in 0.38 seconds, and the display deflate back to 1 psi in 0.12 seconds.

E. Implementation

The haptic displays were mounted on cylindrical surfaces for the studies outlined in the following experiments, either sections of the robot arms or a PVC pipe acting as a stand-in. The mounting arrangements fixed the wrapped display in place, restricting it to 10% contraction. The basic pneumatic control systems used to actuate the wrapped haptic displays consisted of: (1) a pressure regulator that supplied an electronically controlled pressurized-air supply and (2) a pressure release feature for deflating the displays. Two different pressure regulators were used. A pressure regulator with built in sensor and exhaust (QB3, Proportion-Air, McCordsville, Indiana) was used for the studies outlined in Sections IV and VI and was controlled using an Arduino Uno via MATLAB. A different pressure regulator (550-AID, ControlAir, Amherst, New Hampshire) was used for the remaining experiments, controlled using the UR-10's I/O controller (Section V) or an Arduino Uno (Section VII). For this pressure regulator, the inflation pressure was measured using an electronic pressure sensor (015PGAA5, Honeywell Sensing, Gold Valley, Minnesota). If faster switching between inflation and full deflation is needed, on-off solenoid valves can be implemented. It is important to note that each 1-DoF device



Fig. 3. Experimental Setup. The participants were instructed to sit at the desk right in front of the curtain and put on hearing protection headphones.

was connected to an individual pressure supply. For the case of the 3-DoF display, one can configure the device to effectively act as a 1-DoF device by connecting the individual rings to a single pressure, or have 3-DoF control if three pressure regulators are used.

IV. MEASURING HUMAN PERCEPTION OF 1-DOF WRAPPED HAPTIC DISPLAYS

Understanding the human sensory perception of the soft display, especially compared to rigid haptic displays, is essential in applying and controlling the wrapped haptic display. To that end, we first conducted a psychometric user study to measure the ability to distinguish display signals outside of the context of the target scenario. Participants physically interacted with the 1-DoF display and were asked to distinguish between pairs of pressures. We studied the user's ability to differentiate inflation levels in the display to understand what pressure differences produce clear signals. This experiment was previously featured in greater detail in our previous work [8].

A. Experiment Setup

The 1-DoF inflatable haptic display was mounted on a PVC pipe matching the diameter of the UR-10 (Section V), and the pipe was secured flat to the table. A curtain blocked the user's vision, and users wore hearing protection to ensure the perception study focused on tactile sensations (Fig. 3).

The study was conducted as a forced-choice comparison where participants were asked to identify the higher of two pressures. Pressures were shown in pairs (i.e., reference pressure, P_o , vs. test pressure, P), distinguished as "Pressure 1" and "Pressure 2". We selected 2 psi (13.79 kPa) as the reference pressure, and the test pressure values of 1.5, 1.75, 1.875, 2.0, 2.125, 2.25, and 2.5 psi (10.34, 12.07, 12.93, 13.79, 14.65, 15.51, and 17.93 kPa). These pressures are within a safe operating range for the display. We randomized the order in which the pairs were shown to the participant, as well as the order of reference and test pressure in each pair. As a note, in some pairs the reference pressure and test pressure were both 2.0 psi to measure bias in participants' choices when guessing.

The participants sat at the desk, and, before beginning the experiments, we demonstrated the display function, allowing

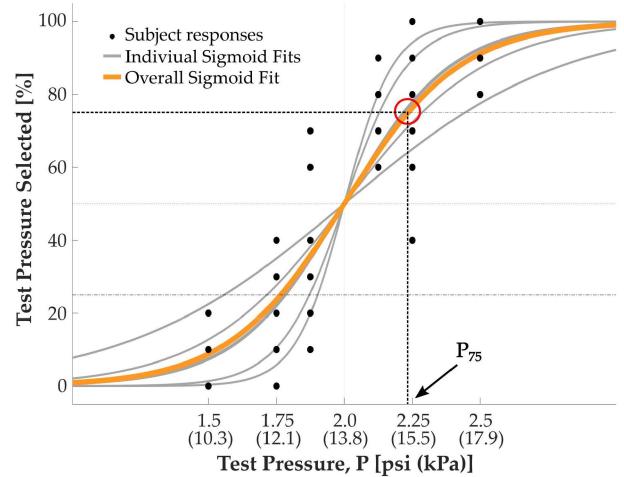


Fig. 4. Raw data and sigmoid function fit for the collection of responses (orange) and each subject (grey). The percentages represent the proportion of times subjects selected the test pressure, P , as higher. The JNDs were calculated using the sigmoid function to solve for the pressure value corresponding to the 75% threshold and subtracting it from the reference pressure. The dots represent percentages associated with individual subject responses. The k steepness factor for the overall sigmoid fit was 4.678, giving a JND of 0.235 psi. The individual steepness factors ranged 2.477-11.15, with JNDs varying between 0.099 and 0.444 psi (0.68-3.06 kPa).

participants to interact with it. Each experimental trial began by inflating the display to "Pressure 1". The participants were told to interact with the display for an unrestricted period of time and then release it. Then, the display was inflated to "Pressure 2," and the participants were asked to interact again. Once they interacted with both pressures, we asked which one felt like a higher inflation pressure. The subjects were not told the correct answers. This procedure was repeated ten times each for the seven test pressures. After completing the interaction portion, the participants were given a post-experiment questionnaire asking about their overall study experience and their previous experiences with haptic technology, robotics, etc. The entire experiment took approximately 35 minutes, with an optional break.

B. Results

A total of 10 participants (4 female, 1 non-binary, 5 male, average age 20.6 years, age range 18 – 23 years) participated in this experiment after giving informed consent. Out of the group, 9 participants were right-handed, and 1 was left-handed. The Purdue Institutional Review Board approved the study protocols. Fig. 4 shows the subjects' responses to the experiment. Each dot shows the percentage of times the test pressure was selected as higher when compared to the reference pressure. The just noticeable difference (JND) was calculated by first fitting a sigmoid function to the data:

$$q = \frac{100}{1 + e^{-k(P - P_0)}} \quad (1)$$

where q is the modeled percentage of times the user chose the test pressure (P) as higher, k is the steepness factor for fitting a sigmoid curve, P is the test pressure, and P_0 is the reference

pressure. Using this fit, the JNDs are calculated by finding the pressure value corresponding to the 75% threshold, P_{75} , and subtracting the reference pressure, P_0 :

$$JND = P_{75} - P_0 = -\frac{1}{k} \ln \left(\frac{100}{75} - 1 \right) \quad (2)$$

The sigmoid function was fit for each of the subjects, as well as for the collection of responses from all subjects.

C. Analysis

The experimental results show that the k steepness factor for the overall sigmoid fit (the orange line in Fig. 4) was 4.678, with 95% confidence bounds between 3.605 and 5.751, giving a JND of 0.235 psi (1.62 kPa). Individual JNDs ranged 0.099-0.444 psi (0.68-3.06 kPa). The mean JND was defined as the mean of the values obtained for all participants, which was found to be 0.228 psi (1.57 kPa) with a standard deviation of 0.109 psi (0.75 kPa). The Weber fraction (WF), calculated as the ratio of the JND and the reference pressure, ranged between 4.9% and 22.2%, with a mean value of 11.4%. Although there was no restriction on how the user could interact with the display, multiple users reported using active interaction to explore the display. This means that participants used reactive force sensing to explore the dynamics of inflation and determine how much pressure was used. Additionally, users reported mainly using their fingertips. Previous studies on fingertip psychophysics tests show similar results. Frediani and Carpi [65] conducted psychophysical tests for a fingertip-mounted pneumatic haptic display, reporting JNDs of 0.12-0.33 psi (0.8-2.3 kPa) for pressures between 0.58 and 2.90 psi (4 and 20 kPa), yielding a WF of 15%. A study evaluating a haptic jamming display found fingertips WF to be 16% ($\sigma = 7.4\%$) and 14.3% ($\sigma = 2.6\%$) for stiffness and size perception, respectively [66]. A different study testing stiffness perception for a rigid vibrotactile, fingertip-mounted haptic device reported WF between 17.7 and 29.9% [67]. The JNDs and WFs obtained in this study show that our wrapped haptic display produced detectable signals and matched previously found psychometric baselines.

As mentioned in Section IV-A, the reference pressure was shown against itself 10 times to measure subject's bias. Subjects overall showed unbiased behavior, choosing "Pressure 1" 45% of the time and "Pressure 2" 55% of the time. However, two subjects had a large preference for choosing "Pressure 2" (80% of the time when guessing). These subjects also scored the highest WF, which may explain their higher bias when guessing compared to the complete participant pool.

The qualitative data from the post-experiment questionnaire shows that, besides the participants already mentioned (who had the highest WF), no other participants struggled to identify the pressures. A majority of the participants (7 out of 10) agreed that they could detect the differences and that they were sure about their answers. It is also worth noting that subjects with the highest performance reported dexterity-related skills, such as playing string musical instruments, piano, knitting, and American Sign Language proficiency.

This study shows that the sensations produced by our wrapped haptic display match the psychometric measures for other haptic devices. Both quantitative and qualitative results show that users were able to distinguish pressure changes over time without a specific task context. Overall, we demonstrated that the soft-wrapped haptic display can perform as well as other haptic devices in displaying tactile signals.

D. Follow-Up Study

As timing became a significant factor during the later studies in Sections V-VII, a follow-up study was conducted, replicating the experimental procedure with the addition of a graphical user interface (GUI). The purpose of the GUI was to enable participants to control the pace of the experiment without the influence of the experimenter and to allow accurate recording of the time spent exploring each pressure. By evaluating time, we can better understand later result on timing and difficulty of interpreting haptic signals.

A total of 12 participants (6 female, 0 non-binary, 6 male, average age 21.9 years, age range 21 – 23 years) participated in the follow-up experiment after providing informed consent. Due to technical difficulties in data collection, 2 participants were removed from the study. 1 additional participant was excluded from analysis as an outlier (performance equivalent to guessing). Of the remaining 9 participants, 7 were right-handed, and 2 were left-handed.

The results show a JND of 0.279 psi (1.923 kPa). Individual JNDs ranged 0.114-0.674 psi (0.788-4.650 kPa), with a mean JND of 0.310 psi (2.136 kPa) and standard deviation of 0.173 psi (1.195 kPa). The WF ranged between 5.7% and 33.7%, with a mean value of 15.5%, consistent with the initial study.

Participants spent an average of 13.84 s on the first pressure ($\sigma = 7.323$ s), and an average of 11.27 s on the second pressure ($\sigma = 5.746$ s), for an average of 25.11 s per pair ($\sigma = 10.855$ s). By one-way ANOVA, total time spent per pair was found to significantly impact correctness ($p = 0.024$). Subjects spent significantly more time assessing the haptic device when answering incorrectly (26.84 s) than when answering correctly (24.56 s). Notably, mean time itself did not have a significant influence on overall accuracy ($p = 0.973$).

V. APPLYING WRAPPED HAPTIC DISPLAYS TO COMMUNICATE 1-DOF ROBOT UNCERTAINTY

So far we have studied the precision with which humans can perceive the 1-DoF wrapped haptic display. Next, we apply this display to convey robot learning from physical interactions. Section V presents a condensed version of the robot experiments in [8], excluding some details on tasks, metrics, and procedure. In this experiment participants kinesthetically taught a UR-10 robot arm to perform cleaning tasks. We applied an existing learning algorithm to measure the robot's uncertainty [11] and then conveyed that uncertainty back to the human in real-time. We highlight two key differences from the experiment in Section IV: here the robot arm is *moving* during interaction (i.e., the wrapped haptic display is not stationary), and the haptic display

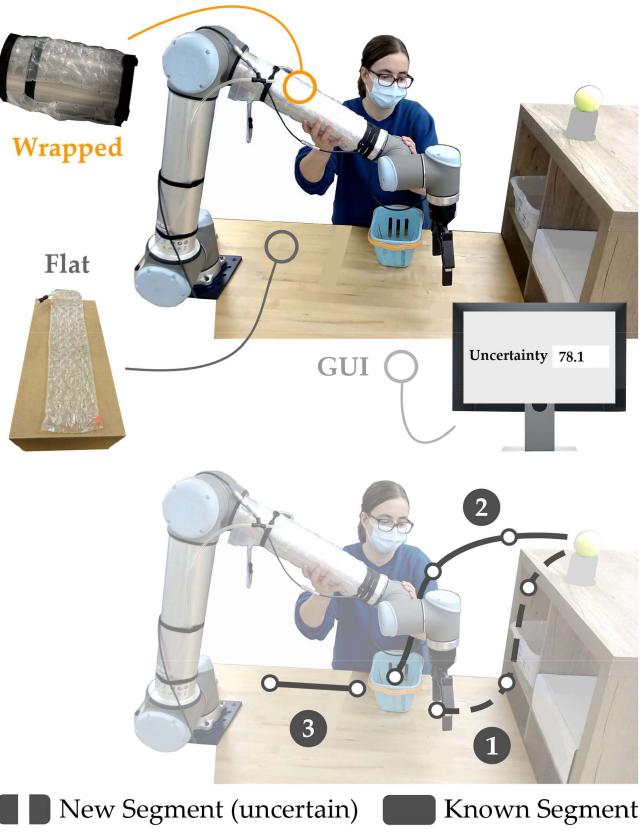


Fig. 5. Participant kinesthetically teaching the robot arm the *Cleaning* task. (Top) We compared our proposed approach (**Wrapped**) to two alternatives. **GUI** displayed the robot’s uncertainty on a screen, while in **Flat** we placed the haptic display on table. (Bottom) We initialized the robot with data from known segments. During their first demonstration the human attempted to identify the region where the robot was uncertain (i.e., the new segment). The human then gave a second demonstration where they only guided the robot through the region(s) where they thought it was uncertain.

now conveys a specific signal that the human must interpret and react to during interaction.

Independent Variables: We compared three different types of feedback (see Fig. 5):

- A graphical user interface (**GUI**) that displayed the robot’s uncertainty on a computer monitor.
- Our soft haptic display placed **Flat** on the table.
- Our proposed approach where we **Wrapped** the haptic display around the robot arm.

All three types of feedback showed the same information but used different modalities. In the **GUI** baseline we displayed uncertainty on a computer screen in front of the user. Here uncertainty was shown as a percentage, where numbers close to 100% indicated that the robot was uncertain at the current state. The **Flat** and **Wrapped** interfaces used the 1-DoF soft haptic display from Section III. Uncertainty was linearly scaled from 1 – 3 psi (6.89 – 20.68 kPa). Here 1 psi (deflated bags) corresponded to 0% uncertainty and 3 psi (inflated bags) corresponded to 100% uncertainty. The **Flat** haptic display was placed in a designated area next to the human, such that participants could periodically touch it during the experiment.

Experimental Setup: Participants completed three different tasks with each of the three feedback conditions (i.e., nine total trials). Tasks involved pushing, grasping, and moving objects around a table and drawers. Fig. 5 shows an example task.

Before conducting any experiments we first initialized the robot’s uncertainty. We collected five expert demonstrations of each task and trained the robot with a behavior cloning approach [11]. This approach outputs the robot’s uncertainty at each state (i.e., uncertainty was a function of the robot’s joint position). We purposely removed segments of the expert’s demonstrations from the training set: specifically, we trained the robot without showing it how to perform either the first segment or the last segment of the task. As a result, when participants interacted with the robot, the robot was uncertain about either the start or the end of the task.

For each trial the participant provided *two demonstrations*. First, the participant kinesthetically guided the robot throughout the entire task while receiving real-time feedback from **GUI**, **Flat**, or **Wrapped**. Based on this feedback, the participant attempted to identify the region of the task where the robot was uncertain (and needed additional teaching). During the second demonstration, the human *only taught the segment* of the task where they believed the robot was *uncertain* (i.e., the region they identified in the first demonstration). If the feedback is effective, participants should only reteach segments where the robot is confused without repeating parts of the task that the robot already knows.

Participants: We recruited 10 participants from the Virginia Tech community to take part in our study (5 female, 0 non-binary, 5 male, average age 22.9, age range 19 – 26 years).

Dependent Measures: To measure how the robot’s feedback affected the human’s teaching, we focused on the *second demonstration* (i.e., the demonstration where users retaught the uncertain part of the task). We recorded the time users spent on this second demonstration (*Teaching Time*) and the percentage of this second demonstration that overlapped with the segment where the robot was actually uncertain (*Correct Segment*). Offline, we retrained the robot using the participant’s second demonstration. We then measured the percentage reduction in uncertainty due to the user’s demonstration (*Improvement*). Finally, we also measured how users *subjectively perceived* each feedback method using a 7-point Likert scale survey.

Hypotheses: We had two hypotheses for this user study:

H1. Participants will most efficiently teach the robot with wrapped haptic displays.

H2. Participants will subjectively prefer our wrapped haptic display over other methods.

Results: We summarize our aggregated results in Fig. 6.

We first ran a repeated measures ANOVA, and found that the robot’s feedback type had a statistically significant effect on *Teaching Time*, *Correct Segment*, and *Improvement*. Post hoc analysis revealed that participants spent less time teaching the robot with **Wrapped** than with either **GUI** or **Flat** ($p < .05$). Participants also better focused their teaching on the region

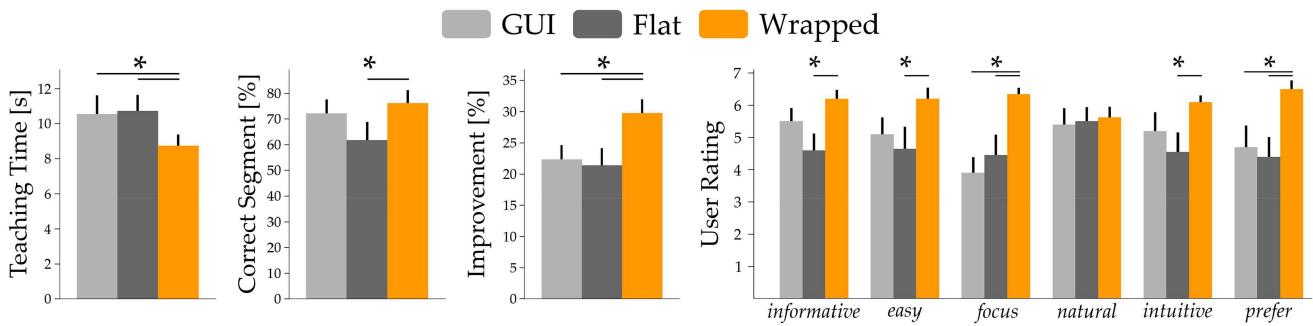


Fig. 6. Objective and subjective results when communicating 1-DoF robot uncertainty in real-time with **GUI**, **Flat**, and **Wrapped** feedback. Participants taught the robot three tasks; we here report the aggregated results across tasks. Error bars show standard error of the mean (SEM), and * indicates statistically significant comparisons ($p < .05$). (Left) Wrapping the haptic display around the robot arm caused participants to spend less time teaching the robot, focused their teaching on regions where the robot was uncertain and improved the robot's understanding of the task after the human's demonstration. (Right) Participants thought that the wrapped display best enabled them to focus on the task, and they preferred this feedback type to the alternatives.

where the robot was actually uncertain: **Wrapped** resulted in a higher *Correct Segment* than **Flat** ($p < .05$). However, here the differences between **Wrapped** and **GUI** were not statistically significant ($p = .287$).

Recall that *Improvement* captures how much more confident the robot is about the task after the participant's demonstration. This metric is especially important: we want to enable humans to teach robots efficiently, and *Improvement* quantifies how much the robot learned from the human's teaching. We found that the robot's confidence improved the most in the **Wrapped** condition as compared to either **GUI** or **Flat** ($p < .05$). Overall, these results support **H1**: when users get real-time feedback from a haptic display wrapped around the robot arm, they provide shorter duration kinesthetic demonstrations that more precisely hone in on the robot's uncertainty and efficiently correct the robot.

We next analyzed our Likert scale survey to understand how users perceived each type of feedback. After confirming that our six scales were reliable (using Cronbach's α), we grouped these scales into combined scores and ran a one-way repeated measures ANOVA on each resulting score. Post hoc analysis showed that participants thought that **Wrapped** was more informative, easier to interact with, less distracting, and more intuitive than either one or both of the alternatives ($p < .05$). Participants also indicated that they preferred **Wrapped** over **GUI** and **Flat**. When explaining this preference, one participant said, "*I definitely prefer Wrapped over other methods. I was able to clearly focus and the other methods were distracting.*". Our subjective results support **H2**, and indicate that users perceived wrapped haptic displays as preferable when compared to alternatives like visual interfaces. We also note that the overall results from videos, user ratings, and teaching time indicate that participants were able to detect **Wrapped** feedback during their kinesthetic demonstrations, i.e., in **Wrapped** users did not need to stop moving, explore the pouches, and then resume their demonstration.

VI. MEASURING HUMAN PERCEPTION OF 3-DOF WRAPPED HAPTIC DISPLAYS

Having explored the human perception and application of the 1-DoF wrapped haptic display in the shape of a sleeve,

we next pursue a study that will help us understand how the spatial distribution of displays affects the perception of both 1-DoF and multiple-DoF soft haptic displays. Both temporally and spatially varying signals can help us add complexity when we need to communicate multiple haptic signals within the space that a human might contact. We also seek to understand how the spatial distribution of signals might affect the effectiveness of the display in terms of accuracy of identification and the time needed to identify signals. To do so, we conducted a user study to measure the ability to distinguish haptic signals in different spatial distributions and outside of the context of the target application scenario. We select pressure levels considering the psychometric baselines (JNDs) obtained in Section IV, and designed a study in which participants physically interacted with 3-DoF displays. The displays were arranged in two ways: (1) a 3-DoF ring display placed in a single location, and (2) three 1-DoF displays made up of three rings each (by interconnecting the individual rings) and placed at three different locations. We called these arrangements **Global** (for the 3-DoF display), meaning all information was available at the single point of contact, so it would be "globally" available, and **Local** (for the three 1-DoF displays), meaning the information for each degree of freedom was only available locally. In each display, the user was asked to identify the signal with the highest pressure out of the three, and we hypothesized that the distribution of the signals (whether three in a single location or spread over a distance) would affect performance. As a note, these same methods are later used in the experiment in Section VII, but there three of the **Global** displays are used instead of one to keep the total area of the display on the robot constant. This also allows different users to contact the robot arm at different locations based on preference while still receiving the same feedback.

A. Experiment Setup and Procedure

The 3-DoF wrapped haptic displays were mounted on passive stand-ins. For the **Local** method, three stand-ins with three ring displays configured as 1-DoF displays each were placed on the table, with a separation in between each. For the **Global** method, a single stand-in with a 3-DoF display was used. Both methods essentially have 3-DoF, but the difference is the spatial

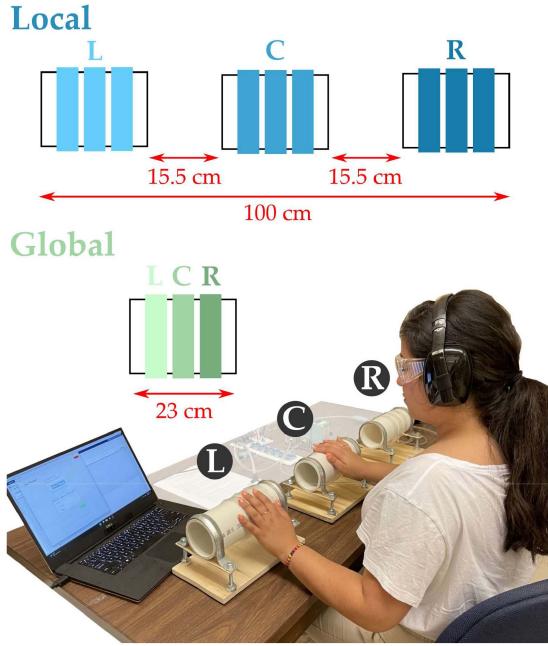


Fig. 7. Experimental Setup. (Top) The **Local** setup consists of three sets of 3-DoF displays configured as 1-DoF each, with a separation in between each. (Middle) The **Global** setup consists of a single 3-DoF display. Both methods essentially have 3-DoF, but the difference is the spatial distribution of each of the DoF. The three DoF were named *Left*, *Center*, and *Right*, for both methods. (Bottom) Participants were instructed to sit in front of the setup; here, we show a participant interacting with the **Local** setup.

distribution of each of the degrees of freedom; for **Global**, all signals are located in a small space, while for **Local** the signals are distributed in a 1 m space. The three degrees of freedom were named *Left*, *Center*, and *Right*, for both methods. The setups are illustrated in Fig. 7. Participants were instructed to wear hearing protection and safety glasses during the study. The task was to identify which of the signals, *Left*, *Center*, or *Right*, was inflated to the higher pressure. Two of the degrees of freedom were inflated to a reference pressure P_o (2 psi) and one to a high pressure P_H (2.75 psi). Subjects were not told that two degrees of freedom had the same pressure, they were just instructed to identify the one inflated to the different pressure. We selected the P_o and P_H values based on the findings of the previous psychophysics study and taking into consideration that there is an increase in the complexity of haptic signals for this new study. As reported in Section IV-C, the average JND found in the previous study was 0.228 psi. However, some of the participants had JNDs almost double of the mean. With that in mind, we determined that a pressure difference between the signals of $\Delta P = 0.75$ psi was large enough so that we could be sure all subjects would perform to an adequate level in this study.

Each DoF (*Left*, *Center*, and *Right*) was rendered to the participant as the P_H a total of 16 times each, for a total of 48 trials. The process was performed for both **Global** and **Local** methods. Half of the participants completed the procedure with **Global** first and the other half completed **Local** first. The study was as follows. Participants were instructed to sit at the desk in front of the arrangements. They interacted with a GUI developed in

LOCAL				GLOBAL				
Signal	R	1.88	2.50	95.63	R	3.75	2.50	93.75
	C	1.88	95.63	2.50	C	3.13	93.75	3.13
	L	97.50	1.25	1.25	L	90.63	5.00	4.38
		L	C	R		L	C	R
Response					Response			

Fig. 8. Confusion matrices showing the mean accuracy for each signal rendered (*Left*, *Center*, *Right*) in both methods (**Local** and **Global**).

MATLAB to navigate through the study. The GUI first guided the participants through a demo to demonstrate the study procedure. The GUI showed a red light that would turn green to indicate when the participant was allowed to touch the displays. For each of trial, the GUI asked the participant to click a “Next” button to continue. Once clicked, the light would turn green once the displays reached their corresponding steady-state pressures. The participants were allowed to touch the displays for an unrestricted period of time and they could explore the displays using any method, including using both hands if desired. During this time, the GUI displayed the question “*Which one has the different pressure?*”, and showed options for *Left*, *Center*, and *Right*. After the participants selected an option and confirmed by clicking an “Enter” button, the GUI showed whether they were correct and, if incorrect, what the right answer was. Note that the GUI was configured to measure the participants’ response time; an internal timer would start when the light turned green and would stop when the participants answered the question. To continue with the next trial, the participants then had to press “Next.” The procedure was repeated until 48 trials were completed for the first method and then for the second method. Participants were offered a break halfway through each method and another break in between methods. After completing the interaction portion of the experiment, participants answered a post-experiment questionnaire. The questionnaire asked about how distinguishable the signals were, if they were often unsure about their answers and if they were increasingly confident about their answers as the study progressed. We also asked about the overall experience during the study (clarity of instructions, sense of safety during the experiment) and about their previous experiences and familiarity with haptic technology, robotics, etc. The study was 45 minutes long.

B. Results

We recruited 10 participants (4 female, 0 non-binary, 6 male, average age 22.1 years, age range 19 – 25 years) from the Purdue community. All participants completed the study after giving informed consent. The Purdue Institutional Review Board approved the study protocols (IRB #2021-1283). In the group, 9 participants were right-handed; one was left-handed.

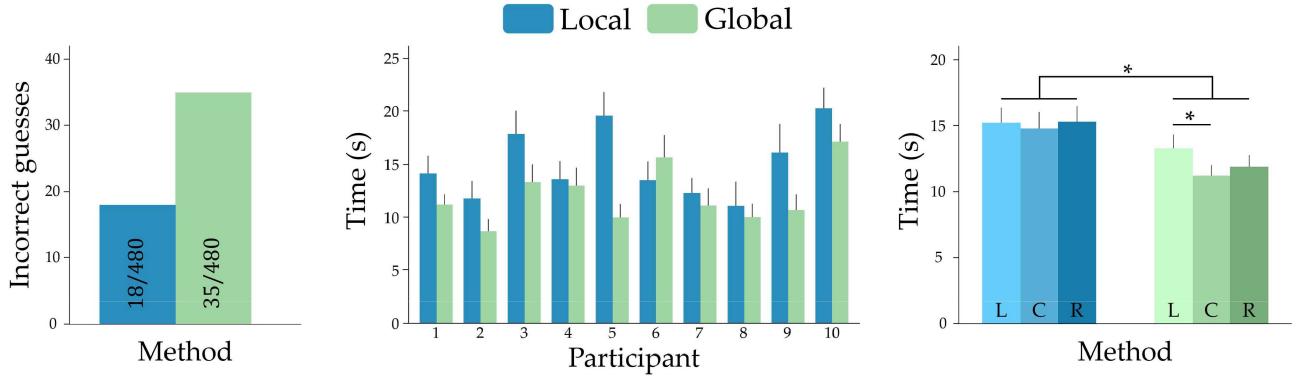


Fig. 9. Experimental results. (Left) Count of incorrect guesses for each of the methods. A Wilcoxon signed-rank test showed that there is a significant association between participants' accuracy and methods ($Z = -2.335, p < .05$). (Center) Mean response time of individual participants for each method. Nine out of ten participants had a higher response time using the **Local** method. (Right) Mean response time for both methods, displayed by signal type (*Left*, *Center*, *Right*). Signal type had a statistically significant effect on response time ($p = .047$) but to a lesser extent than the method type ($p < .001$). For **Global**, participants spent more time responding the question when the *Left* signal was the highest pressure than when it was *Center* ($p < .01$) or *Right* ($p = .078$).

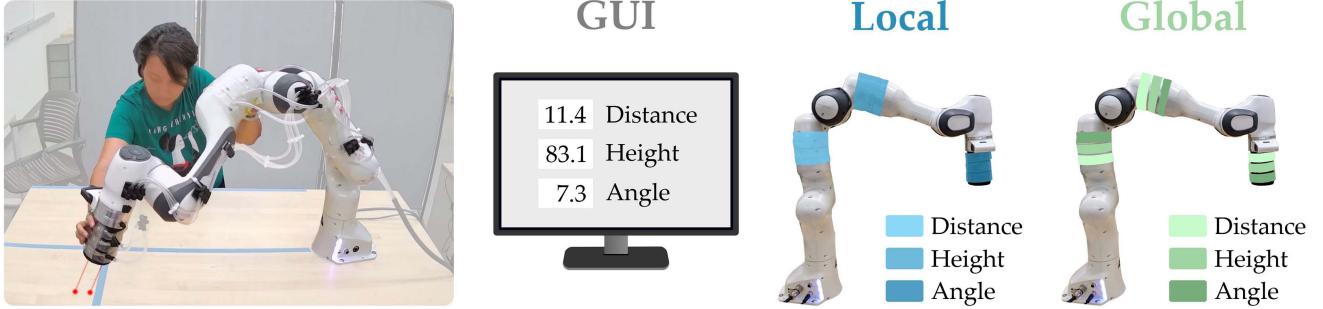


Fig. 10. Experimental setup and independent variables for the user study from Section VII. (Left) Participants physically demonstrated a mock welding task to a Franka Emika robot arm. We mounted two lasers to the robot's end-effector: the robot prompted human teachers to keep the end-effector and lasers at the correct distance, height, and orientation. (Right) The robot indicated which feature(s) it needed help with using three different feedback modalities: **GUI**, **Local**, and **Global**. For **GUI** the robot printed its percentage uncertainty about each feature on a computer monitor placed in front of the workstation. Both **Local** and **Global** leveraged our wrapped haptic displays. In **Local** we attached three 1-DoF displays, and the location of the display indicated the desired feature. By contrast, in **Global** we used three 3-DoF displays such that each row of the displays corresponded to a separate feature.

The confusion matrices in Fig. 8 summarize the accuracy of participants. Overall, participants' accuracy was higher for the **Local** method (average $\bar{x} = 96.25\%$, standard deviation $\sigma = 3.88$) than **Global** ($\bar{x} = 92.71\%$, $\sigma = 7.17$). Participants spent an average of 15.09 s ($\sigma = 7.55$) using the **Local** method, and 12.12 s ($\sigma = 5.877$) for **Global**. Interestingly, looking at the complete pool of participants' responses (whether global or local), we found that participants had a greater response time when they responded incorrectly ($\bar{x} = 16.89$ s, $\sigma = 7.68$ s) than when they answered correctly ($\bar{x} = 13.41$ s, $\sigma = 6.83$). Fig. 9 shows the average time spent by each participant for both **Local** and **Global** methods.

C. Analysis

The two quantitative measures that we used to understand the results are *Accuracy* and *Response Time*. To further analyze the accuracy of participants, a Wilcoxon signed-rank test was conducted to understand the relation between accuracy and the methods used. The results showed that there is a significant association between participants' accuracy and methods

($Z = -2.335, p < .05$). This means that although participants responded faster to the task while using **Global** as shown by the mean response time values, participants were not as accurate at detecting the higher pressure as when they were using **Local**. Fig. 9 shows the count of incorrect guesses for both **Local** and **Global** methods. Another Wilcoxon test was conducted to determine whether the order in which the experiments were conducted (**Local** first, then **Global**, or vice-versa) affected subjects' accuracy. The results showed that there was no significant association ($Z = -0.143, p = .886$), suggesting that subjects did not benefit from learning to improve their accuracy for the second half of the study.

To analyze response time, we used a one-way repeated measures ANOVA. We found that the method type had a statistically significant effect on response time. Post hoc analysis revealed that participants spent less time identifying the target signal with **Global** as compared to **Local** ($p < .001$). This observation matches the mean values for the response time previously mentioned and also the mean response time for each participant, as shown in Fig. 9. Nine out of ten participants spent more time

using **Local** compared to **Global**. We also found that the rendered signal (whether *Left*, *Center*, or *Right*) had a statistically significant effect on answering time ($p = .047$) but with a smaller effect size than the method type. Data shows that while using **Global**, participants spent more time responding when the *Left* signal was the highest pressure than when it was *Center* ($p < .01$) or *Right* ($p = .078$). For **Local**, we did not find any statistically significant distinction between signals and their mean response time. These results can be observed in Fig. 9, where we show the response time of participants for each signal type (*Left*, *Center*, *Right*) when using **Local** and **Global** methods.

This study shows that the spatial distribution of haptic displays is an important factor to consider since it has an effect on both accuracy of detection and response time. Using the psychometric measures found in the previous study, we showed that participants were better able to identify the highest pressure out of a set of three when the signals were spatially distributed (**Local**) than when the signals were condensed in a smaller space (**Global**). However, the response time for the spatially distributed signals was higher; this makes sense because participants moved around a larger space to interact with the places where the haptic signals were located. Participants reported the pressure differences were detectable, they were sure about their answers throughout the experiment, and that they felt safe interacting with the displays. Some participants mentioned that during the **Local** portion of the experiment, they wished they could place the displays together to make the exercise easier; this suggests that users consciously thought that having displays dispersed in different locations was an inconvenience, even though the results show participants were slightly more accurate with this method than with **Global**. To summarize, the **Global** method had the faster response time, but **Local** had the higher accuracy. These observations show the trade-off between response time and accuracy when we increase the complexity of haptic signals in a smaller space or distribute them in a larger space.

VII. USING MULTI-DOF WRAPPED HAPTIC DISPLAYS TO COMMUNICATE 3-DOF ROBOT LEARNING

In Section V we demonstrated that robot arms can leverage a haptic display to communicate with human teachers. However, this haptic device only had 1-DoF: the same pressure was rendered along the entire robot arm. One degree-of-freedom is sufficient when the robot learner wants to convey whether or not it is uncertain — but what if the robot needs to communicate more complicated feedback? For instance, the robot may want to indicate *what* it is confused about or *how* the human teacher could improve their demonstrations.

In our final user study we wrap multiple 3-DoF haptic displays around a Franka Emika robot arm. Participants physically teach the robot to perform a mock welding task, and the robot applies multi-dimensional feedback to indicate what aspects of the task the human teacher must emphasize. Overall, our goal is to compare the two different feedback distributions shown in Section VI and understand how they impact the human's physical demonstrations. Remember that we are wrapping haptic displays along the robot arm. One option is to *localize* different

signals to different parts of the arm, such that the place where the bags inflate helps indicate and remind users what the robot is uncertain about. Our second option is to *distribute* all three signals along the entire arm; here the human perceives the same haptic rendering no matter where they grasp the robot. In this user study we explore how human teachers perceive and leverage multiple displays that use both feedback layouts.

Independent Variables: Participants kinesthetically guided the robot arm through a mock welding task. The robot displayed feedback in real-time to guide the human through the task. We compared three different types of feedback for communicating when the robot was uncertain and what motions it needed the human teacher to emphasize (see Fig. 10):

- A **GUI** baseline where the robot showed its numerical uncertainty on a computer monitor.
- Three 1-DoF wrapped haptic displays with signals localized to different regions of the robot arm (**Local**).
- Three 3-DoF wrapped haptic displays with signals distributed across the entire robot arm (**Global**).

All conditions provided the same information to the participants. Similar to Section V, in **GUI** the robot displayed its uncertainties as a percentage: values close to 100% meant that the robot needed assistance. For **Local** and **Global** we actuated three separate wrapped haptic displays with pressures between 1–3 psi (6.89 – 20.68 kPa). In **Local** each location of the haptic display had a single pressure signal; i.e., bags at the end-effector were one pressure, bags at the base of the arm were another pressure, and bags in the middle of the arm were a third pressure. In **Global** each haptic display location rendered all three of the potentially different pressures using three independent degrees of freedom, and all **Global** displays rendered those same three pressures — participants could feel the same feedback at the base, middle, and end of the robot arm. **GUI**, **Local**, and **Global** each provided a total of 3-DoF feedback. We emphasize that with **Global** participants had to discern which segments of the 3-DoF haptic display were inflated, while with **Local** participants needed to determine at which parts of the robot arm the haptic displays were inflated.

Experimental Setup: Participants physically interacted with a 7-DoF robot arm (Franka Emika) to complete a mock welding task (see Fig. 10). Recall that users interacted with a UR10 robot in Section V — from the user's perspective, the Franka Emika robot is smaller, has one more joint, and is easier to backdrive. We mounted lasers to the robot's end-effector: participants kinesthetically guided the robot across a table while the lasers marked where the robot was “welding.”

The welding task consisted of three features: how close the end-effector was to the edge of the table, the end-effector's height from the table, and the orientation of the end-effector. When the task started participants would guide the robot arm towards the fixed goal position. As they moved, the robot would leverage its feedback to notify the human *which feature* they needed to emphasize. For example, during the first third of the task the robot may prompt the human to keep the lasers close to the table; in the middle of the task the human should move the end-effector to the table edge; and during the final third of the task the human might need to align the robot's orientation. Participants had to

TABLE I
QUESTIONS ON THE LIKERT SCALE SURVEY FROM SECTION VII. WE GROUPED QUESTIONS INTO FIVE SCALES AND EXAMINED THEIR RELIABILITY USING CRONBACH'S α . QUESTIONS EXPLORED WHETHER PARTICIPANTS THOUGHT THE ROBOT'S FEEDBACK WAS EASY TO INTERPRET, IF THEY COULD FOCUS ON TEACHING, HOW DISTINGUISHABLE THE ROBOT'S SIGNALS WERE, WHICH METHODS WERE INTUITIVE, AND THEIR OVERALL PREFERENCES. FOR PREFERENCE WE DID NOT CHECK FOR RELIABILITY SINCE THERE WAS ONLY A SINGLE ITEM

Questionnaire Item	Reliability	$F(2, 22)$	p-value
– It was hard to figure out what the robot was trying to convey to me.	.75	1.699	.206
– I could easily tell what the robot wanted.			
– I could focus on the robot's feedback without having to look up or move my hands.	.74	6.266	< .01*
– I had to physically go out of my way to get the robot's feedback.			
– It was easy to distinguish the different feedback signals.	.64	1.733	.215
– I had to think carefully about what I was seeing / feeling to determine the signal.			
– The way the robot provided feedback seemed intuitive to me.	.86	.081	.923
– I thought the robot's feedback was unintuitive and hard to understand.			
– Overall, I prefer this communication modality.	—	5.189	.191

We then performed a one-way repeated measures anova on the grouped scores: here an * denotes statistical significance.

dynamically determine *what* feature the robot currently needed help with and then *modify* their motion to emphasize that feature. Note that the robot asked for assistance with all three features at different segments of the task — we randomized these segments so that participants could not anticipate the robot's feedback.

Participants and Procedure: We recruited 12 participants (5 female, 0 non-binary, 7 male, average age 28, age range 19 – 35 years) from the Virginia Tech community. All participants provided informed written consent consistent with university guidelines (IRB # 20-755). None of the participants for this study took part in the previous study from Section V. Three of the twelve participants reported that they had physically interacted with robot arms before.

Each user completed the welding task four times. First, we asked users to demonstrate the task without any feedback from the robot. We used this initial demonstration as a baseline to measure their improvement. Next, participants completed the welding task with **GUI**, **Local**, and **Global**. We counterbalanced the order of these feedback conditions: four participants started with **GUI**, four participants started with **Local**, and four participants started with **Global**.

Dependent Measures – Objective: We measured the total time it took for participants to demonstrate the welding task (*Teaching Time*). We also measured the *Improvement* between the human's initial demonstration and their demonstration under each feedback condition. Let $f(\xi) \in \mathbb{R}^k$ be a k -length vector of feature counts along trajectory ξ , e.g., the distance, height, and angle. We define $e(\xi) = \|f(\xi) - f(\xi^*)\|^2$ as the error between the human's demonstration ξ and the ideal trajectory ξ^* . *Improvement* captures how this error changes after receiving robot feedback: $(e(\xi_{initial}) - e(\xi))/e_{max} \cdot 100$, where $e_{max} = \|f(\emptyset) - f(\xi^*)\|^2$ is a normalizer (i.e., the error when the human does not provide any demonstration). *Improvement* captures the percentage change in demonstration quality for each feedback condition: positive *Improvement* reveals that the human is demonstrating the task more accurately.

Dependent Measures – Subjective: Participants responded to a 7-point Likert scale survey after each feedback condition. Our

survey was composed of four multi-item scales and one single-item scale (see Table I). We asked participants how *easy* it was to understand the robot's feedback, whether they could *focus* on the task, how *distinguishable* was the robot's feedback, if the feedback was *intuitive* for this task, and to what extent they *prefer* this condition as a communication modality. Finally, after participants had finished working with all the conditions they responded to a forced-choice comparison: “Which method did you like the most?”

Hypotheses: We had two hypotheses for this user study:

H3. *Distributing multi-DoF haptic feedback along the robot arm (**Global**) will lead to improved demonstrations and lower teaching time.*

H4. *Participants will prefer distributed feedback (**Global**) as compared to localized feedback (**Local**).*

Results – Objective: The results from this user study are summarized in Fig. 12. To get a sense of the users' experience, we also show participant demonstrations in Fig. 11.

Let us start our analysis by looking at the objective results. Using a one-way repeated measures ANOVA, we determined that feedback type had a significant effect on *Teaching Time* ($F(2, 22) = 3.423, p < .05$). Post hoc tests revealed that participants spent less time demonstrating the task with **Global** than with **GUI** ($p < .05$), while the differences between **Global** and **Local** were not significant ($p = .675$). To explain these results we measured the amount of idle time during the demonstration. We found that with **GUI** users needed to stop, look at the monitor, and think about their next action: shifting attention back-and-forth between the monitor and the welding task contributed to the increased *Teaching Time*.

So with **Global**, participants taught the robot more quickly — but did they provide accurate, informative demonstrations? Remember that to measure *Improvement* we first collected a demonstration without feedback, and then compared that initial demonstration to the user's behavior under each condition. The

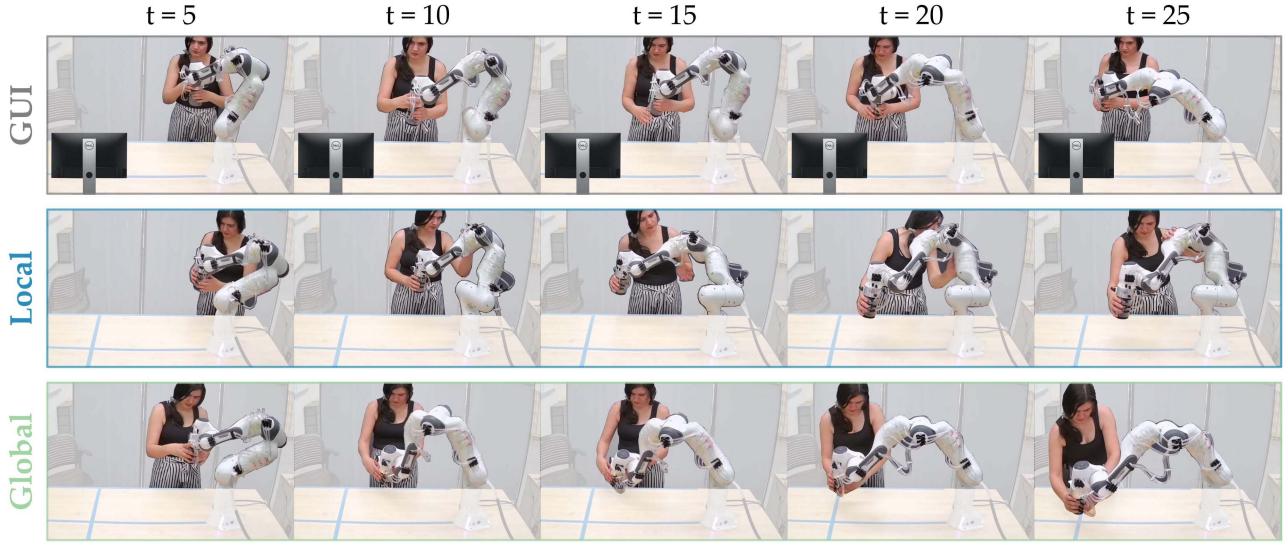


Fig. 11. Participant teaching the welding task with **GUI**, **Local**, or **Global**. We show task progress in 5 s intervals. (Top) With **GUI** users needed to look at the computer monitor to obtain feedback. The monitor is placed on the near side of the table: this participant is looking at the **GUI** at times $t = 5$, $t = 10$, and $t = 25$ seconds. (Middle) With **Local**, participants must move their hands — and change their grasp — to sense the different wrapped displays. This participant keeps one hand on the end-effector, and then moves their other hand between the haptic displays at the middle and base of the robot arm. (Bottom) Finally, with **Global** the participants receive feedback through 3-DoF haptic displays. **Global** helped this user remain focused on the task: notice that they are continually looking at the welding task, and keep both hands on the end-effector (where a 3-DoF haptic display is located).

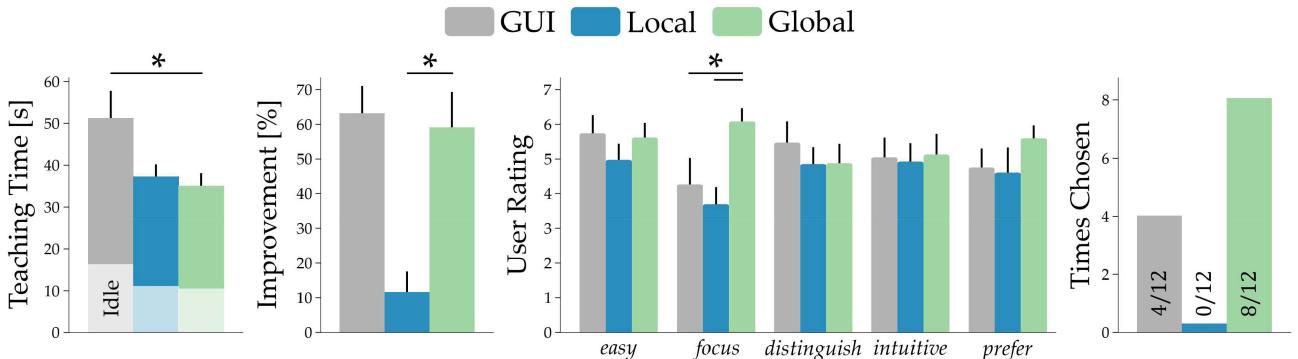


Fig. 12. Objective and subjective results when communicating multi-dimensional robot feedback. We compared using a computer monitor (**GUI**), localizing wrapped haptic feedback to specific parts of the robot (**Local**), and distributing 3-DoF wrapped haptic feedback along the arm (**Global**). Error bars show standard error of the mean and an * indicates statistically significant comparisons. (Left) Participants spent less time teaching the robot with **Global** as compared to **GUI**: shaded regions show the amount of time where participants stopped moving the robot to think about their next actions. The human's demonstrations improved more with **Global** feedback as compared to **Local** feedback. (Middle) Participants perceived the multi-DoF wrapped haptic display as similar to the alternatives, but indicated that **Global** enabled them to focus on teaching the robot. (Right) At the end of the experiment users were asked to choose their favorite method. Of the 12 total participants, 8 selected **Global**, 4 selected **GUI**, and none selected **Local**.

type of robot feedback had a significant effect on *Improvement* ($F(2, 22) = 12.707$, $p < .001$). With both **GUI** and **Global** the participants made similar improvements to their teaching ($p = .769$). However, *Improvement* was significantly lower for **Local** as compared to **Global** ($p < .01$). When participants received **Local** feedback they frequently had to change their grasp and move their hands across the three haptic displays; by contrast, in **GUI** and **Global** the participants could maintain a fixed grasp (Fig. 11). Overall, our objective results support **H3**. **Global** enabled users to teach robots more seamlessly than **GUI** and more accurately than **Local**.

Results – Subjective: Table I and Fig. 12 outline the results of our Likert scale survey and forced-choice comparison. We

first checked the reliability of our four multi-item scales: *easy*, *focus*, and *intuitive* were reliable (Cronbach's $\alpha > 0.7$) but *distinguish* was not. We then grouped each scale into a combined score and performed a one-way repeated measures ANOVA on the result. Note that we did not check for reliability in *prefer* because we only had one item (i.e., one question) on this scale.

We found that participants perceived **GUI**, **Local**, and **Global** to be similar along several axes. For instance, users did not think that any of the feedback types were more distinguishable ($p = .215$) or intuitive ($p = .923$) than the others. However, users reported that they were better able to *focus* on the task with **Global** than with **GUI** ($p < .05$) or with **Local** ($p < .001$).

After the experiment was finished we asked users to select their favorite feedback type: eight of the twelve participants chose **Global**, and the remaining four selected **GUI**. These subjective results support **H4**. Given the results in Section **VI**, we were particularly interested to find that participants preferred **Global** feedback over **Local** feedback — this suggests that the convenience of having all three signal available at each contact point along the arm in the **Global** feedback condition outweighed the slight decrease in accuracy. One participant mentioned that “*I liked Local the least, since it requires repositioning hands to get feedback.*”

VIII. CONCLUSION

In this paper we presented a novel approach for communicating a robot’s internal state during physical interaction. Specifically, we introduced a class of soft, wrapped haptic displays that are mounted on the robot arm at the point of contact between the human and robot; these displays provide real-time feedback as the robot learns from human demonstrations.

We first designed wrapped pneumatic devices using flexible pouches that render one or more pressure signals (Section **III**). We performed psychophysics and robotics experiments with (a) 1-DoF displays and (b) N-DoF displays. With the 1-DoF display, humans could accurately distinguish between different pressures rendered by the wrapped haptic display (Section **IV**). This feedback enabled participants to kinetically teach robot arms more rapidly and effectively as compared to the alternatives (Section **V**).

We next explored N-DoF haptic displays to communicate more detailed feedback. We compared two approaches: localizing separate 1-DoF haptic displays to different regions of the robot arm, or distributing identical 3-DoF displays along the entire arm. From a psychophysics perspective, localized feedback resulted in more accurate communication but at slower speeds, i.e. larger spatially distribution of signals increased accuracy, but required participants to move their hands to perceive each region, and recognize the signal (Section **VI**). We applied both types of haptic displays to a robot learning task. Here we found that distributed 3-DoF signals were preferable to localized 1-DoF signals in terms of teaching time, demonstration improvement, and subjective responses (Section **VII**). We also note that the results of our user studies were consistent across two different industry-standard robot arms, suggesting that approach is not tied to one specific arm type or geometry. Overall, using multi-DoF haptic displays to concentrate signals into a smaller space resulted in more seamless communication and teaching.

Future work will focus on further increasing the complexity of signals that the soft wrapped haptic displays can render. The stacking of pneumatic pouches developed by Do et al. [39] may allow better spatial resolution by eliciting separate force and contact area signals. The integration of sensing technology to the concept of localized pressure distribution may also allow us to break down local pressure measurements (i.e. local user forces) and map them into desired robot motions.

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] B. Akgun, M. Cakmak, K. Jiang, and A. L. Thomaz, “Keyframe-based learning from demonstration,” *Int. J. Social Robot.*, vol. 4, no. 4, pp. 343–355, 2012.
- [3] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal, “Learning and generalization of motor skills by learning from demonstration,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2009, pp. 763–768.
- [4] D. P. Losey, A. Bajcsy, M. K. O’Malley, and A. D. Dragan, “Physical interaction as communication: Learning robot objectives online from human corrections,” *Int. J. Robot. Res.*, vol. 41, no. 1, pp. 20–44, 2022.
- [5] Y. Che, A. M. Okamura, and D. Sadigh, “Efficient and trustworthy social navigation via explicit and implicit robot–human communication,” *IEEE Trans. Robot.*, vol. 36, no. 3, pp. 692–707, Jun. 2020.
- [6] J. F. Mullen, J. Mosier, S. Chakrabarti, A. Chen, T. White, and D. P. Losey, “Communicating inferred goals with passive augmented reality and active haptic feedback,” *IEEE Robot. Automat. Lett.*, vol. 6, no. 4, pp. 8522–8529, Oct. 2021.
- [7] N. Dunkelberger et al., “A multisensory approach to present phonemes as language through a wearable haptic device,” *IEEE Trans. Haptics*, vol. 14, no. 1, pp. 188–199, Jan.–Mar. 2021.
- [8] A. Alvarez Valdivia, R. Shailly, N. Seth, F. Fuentes, D. P. Losey, and L. H. Blumenschein, “Wrapped haptic display for communicating physical robot learning,” in *Proc. IEEE Int. Conf. Soft Robot.*, 2022, pp. 823–830.
- [9] S. Ross, G. Gordon, and D. Bagnell, “A reduction of imitation learning and structured prediction to no-regret online learning,” in *Proc. Int. Conf. Artif. Intell. Statist.*, 2011, pp. 627–635.
- [10] R. Hoque, A. Balakrishna, E. Novoseller, A. Wilcox, D. S. Brown, and K. Goldberg, “ThriftyDAGger: Budget-aware novelty and risk gating for interactive imitation learning,” in *Proc. Conf. Robot Learn.*, 2022, pp. 598–608.
- [11] K. Menda, K. Driggs-Campbell, and M. J. Kochenderfer, “EnsembleDAGger: A Bayesian approach to safe imitation learning,” in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2019, pp. 5041–5048.
- [12] J. Hough and D. Schlangen, “It’s not what you do, it’s how you do it: Grounding uncertainty for a simple robot,” in *Proc. ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2017, pp. 274–282.
- [13] M. Cakmak and A. L. Thomaz, “Designing robot learners that ask good questions,” in *Proc. ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2012, pp. 17–24.
- [14] C. Basu, M. Singhal, and A. D. Dragan, “Learning from richer human guidance: Augmenting comparison-based learning with feature queries,” in *Proc. ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2018, pp. 132–140.
- [15] S. Habibian, A. Jonnavittula, and D. P. Losey, “Here’s what i’ve learned: Asking questions that reveal reward learning,” *ACM Trans. Hum.-Robot Interact.*, vol. 11, no. 4, pp. 1–28, 2022.
- [16] T. Hellström and S. Bensch, “Understandable robots—what, why, and how,” *J. Behav. Robot.*, vol. 9, pp. 110–123, 2018.
- [17] A. D. Dragan, K. C. Lee, and S.S. Srinivasa, “Legibility and predictability of robot motion,” in *Proc. ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2013, pp. 301–308.
- [18] S. H. Huang, D. Held, P. Abbeel, and A. D. Dragan, “Enabling robots to communicate their objectives,” *Auton. Robots*, vol. 43, no. 2, pp. 309–326, 2019.
- [19] R. S. Andersen, O. Madsen, T. B. Moeslund, and H. B. Amor, “Projecting robot intentions into human environments,” in *Proc. IEEE Int. Symp. Robot Hum. Interactive Commun.*, 2016, pp. 294–301.
- [20] M. Walker, H. Hedayati, J. Lee, and D. Szafir, “Communicating robot motion intent with augmented reality,” in *Proc. ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2018, pp. 316–324.
- [21] F. Cini, T. Banfi, G. Ciuti, L. Craighero, and M. Controzzi, “The relevance of signal timing in human–robot collaborative manipulation,” *Sci. Robot.*, vol. 6, no. 58, 2021, Art. no. eabg1308.
- [22] A. Casalino, C. Messeri, M. Pozzi, A. M. Zanchettin, P. Rocco, and D. Prattichizzo, “Operator awareness in human–robot collaboration through wearable vibrotactile feedback,” *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 4289–4296, Oct. 2018.
- [23] S. Grushko, A. Vysocký, P. Oščádal, M. Vocetka, P. Novák, and Z. Bobovský, “Improved mutual understanding for human–robot collaboration: Combining human-aware motion planning with haptic feedback devices for communicating planned trajectory,” *Sensors*, vol. 21, no. 11, 2021, Art. no. 3673.

[24] C. M. Nunez, B. N. Huerta, A. M. Okamura, and H. Culbertson, "Investigating social haptic illusions for tactile stroking (SHIFTS)," in *Proc. IEEE Haptics Symp.*, 2020, pp. 629–636.

[25] S. Muthukumarana, D. S. Elvitigala, J. P. F. Cortes, D. J. Matthes, and S. Nanayakkara, "Touch me gently: Recreating the perception of touch using a shape-memory alloy matrix," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2020, pp. 1–12.

[26] J. K. S. Teh, A. D. Cheok, R. L. Peiris, Y. Choi, V. Thuong, and S. Lai, "Huggy pajama: A mobile parent and child hugging communication system," in *Proc. ACM Int. Conf. Interact. Des. Child.*, 2008, pp. 250–257.

[27] T. Kobayashi, E. Dean-Leon, J. R. Guadarrama-Olvera, F. Bergner, and G. Cheng, "Whole-body multicontact haptic human–humanoid interaction based on leader–follower switching: A robot dance of the box step," *Adv. Intell. Syst.*, vol. 4, no. 2, 2022, Art. no. 2100038.

[28] M. Salvato et al., "Data-driven sparse skin stimulation can convey social touch information to humans," *IEEE Trans. Haptics*, vol. 15, no. 2, pp. 392–404, Apr.–Jun. 2022.

[29] Y. Ju, D. Zheng, D. Hynds, G. Chernyshov, K. Kunze, and K. Minamizawa, "Haptic empathy: Conveying emotional meaning through vibrotactile feedback," in *Proc. Extended Abstr. CHI Conf. Hum. Factors Comput. Syst.*, 2021, pp. 1–7.

[30] C. Rognon, B. Stephens-Fripp, J. Hartcher-O'Brien, B. Rost, and A. Israr, "Linking haptic parameters to the emotional space for mediated social touch," *Front. Comput. Sci.*, vol. 4, pp. 1–7, 2022.

[31] S. Panéels, M. Anastassova, S. Strachan, S. P. Van, S. Sivacoumarane, and C. Bolzmacher, "What's around me? Multi-actuator haptic feedback on the wrist," in *Proc. World Haptics Conf.*, 2013, pp. 407–412.

[32] R. Maderna, M. Pozzi, A. M. Zanchettin, P. Rocco, and D. Prattichizzo, "Flexible scheduling and tactile communication for human–robot collaboration," *Robot. Comput.-Integr. Manuf.*, vol. 73, 2022, Art. no. 102233.

[33] P. Di Campi San et al., "Purring wheel: Thermal and vibrotactile notifications on the steering wheel," in *Proc. ACM Int. Conf. Multimodal Interact.*, 2020, pp. 461–469.

[34] A. Löcken, H. Buhl, W. Heuten, and S. Boll, "Tacticar: Towards supporting drivers during lane change using vibro-tactile patterns," in *Proc. Int. ACM Conf. Automot. User Interfaces Interactive Veh. Appl.*, 2015, pp. 32–37.

[35] I. Choi, E. Ofek, H. Benko, M. Sinclair, and C. Holz, "CLAW: A multi-functional handheld haptic controller for grasping, touching, and triggering in virtual reality," in *Proc. ACM CHI Conf. Hum. Factors Comput. Syst.*, 2018, pp. 1–13.

[36] S. Mintchev, M. Salerno, A. Cherpillod, S. Scaduto, and J. Paik, "A portable three-degrees-of-freedom force feedback origami robot for human–robot interactions," *Nature Mach. Intell.*, vol. 1, no. 12, pp. 584–593, 2019.

[37] M. Raitor, J. M. Walker, A. M. Okamura, and H. Culbertson, "WRAP: Wearable, restricted-aperture pneumatics for haptic guidance," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 427–432.

[38] N. Agharese et al., "HapWRAP: Soft growing wearable haptic device," *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 5466–5472.

[39] B. H. Do, A. M. Okamura, K. Yamane, and L. H. Blumenschein, "Macro-mini actuation of pneumatic pouches for soft wearable haptic displays," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 14499–14505.

[40] H. Zhao et al., "A wearable soft haptic communicator based on dielectric elastomer actuators," *Soft Robot.*, vol. 7, no. 4, pp. 451–461, 2020.

[41] J. Barreiros et al., "Fluidic elastomer actuators for haptic interactions in virtual reality," *IEEE Robot. Automat. Lett.*, vol. 4, no. 2, pp. 277–284, Apr. 2019.

[42] J. Yin, R. Hinche, H. Shea, and C. Majidi, "Wearable soft technologies for haptic sensing and feedback," *Adv. Funct. Mater.*, vol. 31, no. 39, 2021, Art. no. 2007428.

[43] X. Ji et al., "Untethered feel-through haptics using 18- μm thick dielectric elastomer actuators," *Adv. Funct. Mater.*, vol. 31, no. 39, 2021, Art. no. 2006639.

[44] Y.-L. Feng, C. L. Fernando, J. Rod, and K. Minamizawa, "Submerged haptics: A 3-DoF fingertip haptic display using miniature 3D printed airbags," in *Proc. ACM SIGGRAPH Posters*, 2017, pp. 1–2.

[45] R. Hinche, V. Vechev, H. Shea, and O. Hilliges, "Dextres: Wearable haptic feedback for grasping in VR via a thin form-factor electrostatic brake," in *Proc. ACM Symp. User Interface Softw. Technol.*, 2018, pp. 901–912.

[46] D. Leonardis, G. Santamato, M. Gabardi, M. Solazzi, and A. Frisoli, "A parallel-elastic actuation approach for wide bandwidth fingertip haptic devices," *Meccanica*, vol. 57, no. 3, pp. 739–749, 2022.

[47] H. Boys, G. Frediani, M. Ghilardi, S. Poslad, J. C. Busfield, and F. Carpi, "Soft wearable non-vibratory tactile displays," in *Proc. IEEE Int. Conf. Soft Robot.*, 2018, pp. 270–275.

[48] D. Hwang, J. Lee, and K. Kim, "On the design of a miniature haptic ring for cutaneous force feedback using shape memory alloy actuators," *Smart Mater. Struct.*, vol. 26, no. 10, 2017, Art. no. 105002.

[49] K. Minamizawa, D. Prattichizzo, and S. Tachi, "Simplified design of haptic display by extending one-point kinesthetic feedback to multipoint tactile feedback," in *Proc. IEEE Haptics Symp.*, 2010, pp. 257–260.

[50] D. Leonardis, M. Solazzi, I. Bortone, and A. Frisoli, "A wearable fingertip haptic device with 3 DoF asymmetric 3-RSR kinematics," in *Proc. IEEE World Haptics Conf.*, 2015, pp. 388–393.

[51] Z. Zhakypov and A. M. Okamura, "FingerPrint: A 3-D printed soft monolithic 4-degree-of-freedom fingertip haptic device with embedded actuation," in *Proc. IEEE Int. Conf. Soft Robot.*, 2022, pp. 938–944.

[52] E. Leroy, R. Hinche, and H. Shea, "Multimode hydraulically amplified electrostatic actuators for wearable haptics," *Adv. Mater.*, vol. 32, no. 36, 2020, Art. no. 2002564.

[53] J.-H. Youn, H. Mun, and K.-U. Kyung, "A wearable soft tactile actuator with high output force for fingertip interaction," *IEEE Access*, vol. 9, pp. 30206–30215, 2021.

[54] F. H. Giraud, S. Joshi, and J. Paik, "Haptigami: A fingertip haptic interface with vibrotactile and 3-DoF cutaneous force feedback," *IEEE Trans. Haptics*, vol. 15, no. 1, pp. 131–141, Jan.–Mar. 2022.

[55] X. Yu et al., "Skin-integrated wireless haptic interfaces for virtual and augmented reality," *Nature*, vol. 575, no. 7783, pp. 473–479, 2019.

[56] H. A. Sonar, A. P. Gerratt, S. P. Lacour, and J. Paik, "Closed-loop haptic feedback control using a self-sensing soft pneumatic actuator skin," *Soft Robot.*, vol. 7, no. 1, pp. 22–29, 2020.

[57] C. Suh, J. C. Margarit, Y. S. Song, and J. Paik, "Soft pneumatic actuator skin with embedded sensors," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2014, pp. 2783–2788.

[58] L. Yao, R. Niizuma, J. Ou, S. Follmer, C. Della Silva, and H. Ishii, "PneUI: Pneumatically actuated soft composite materials for shape changing interfaces," in *Proc. ACM Symp. User Interface Softw. Technol.*, 2013, pp. 13–22.

[59] A. A. Stanley and A. M. Okamura, "Controllable surface haptics via particle jamming and pneumatics," *IEEE Trans. Haptics*, vol. 8, no. 1, pp. 20–30, Jan.–Mar. 2015.

[60] H. Culbertson, S. B. Schorr, and A. M. Okamura, "Haptics: The present and future of artificial touch sensation," *Annu. Rev. Control. Robot. Auton. Syst.*, vol. 1, no. 1, pp. 385–409, 2018.

[61] J.-B. Chossat, D. K. Chen, Y.-L. Park, and P. B. Shull, "Soft wearable skin-stretch device for haptic feedback using twisted and coiled polymer actuators," *IEEE Trans. Haptics*, vol. 12, no. 4, pp. 521–532, Oct.–Dec. 2019.

[62] Z. Wang et al., "A three-fingered force feedback glove using fiber-reinforced soft bending actuators," *IEEE Trans. Ind. Electron.*, vol. 67, no. 9, pp. 7681–7690, Sep. 2020.

[63] H. In, B. B. Kang, M. Sin, and K.-J. Cho, "Exo-glove: A wearable robot for the hand with a soft tendon routing system," *IEEE Robot. Automat. Mag.*, vol. 22, no. 1, pp. 97–105, Mar. 2015.

[64] B. Larsen, P. Manoonpong, and J. Jørgensen, "WISARD: Weight informing soft artificial robotic dermis," in *Proc. IEEE Int. Conf. Soft Robot.*, 2022, pp. 1–8.

[65] G. Frediani and F. Carpi, "Tactile display of softness on fingertip," *Sci. Rep.*, vol. 10, 2020, Art. no. 20491.

[66] A. M. Genecov, A. A. Stanley, and A. M. Okamura, "Perception of a haptic jamming display: Just noticeable differences in stiffness and geometry," in *Proc. IEEE Haptics Symp.*, 2014, pp. 333–338.

[67] A. T. Maereg, A. Nagar, D. Reid, and E. L. Secco, "Wearable vibrotactile haptic device for stiffness discrimination during virtual interactions," *Front. Robot. AI*, vol. 4, 2017, Art. no. 42.



Antonio Alvarez Valdivia (Graduate Student Member, IEEE) received the B.S. degree in mechanical engineering with Iowa State University, Ames, IA, USA, in 2021. He is currently working toward the Ph.D. degree in mechanical engineering with Purdue University, West Lafayette, IN, USA. He is currently a Research Assistant with the RAAD Laboratory, Department of Mechanical Engineering, Purdue University. His research interests include soft haptics, soft robotics, and mechatronics.



Soheil Habibian (Graduate Student Member, IEEE) received the B.Sc. degree in mechanical engineering from Qazvin Azad University, Qazvin, Iran, in 2015, and the M.Sc. degree in mechanical engineering from Bucknell University, Lewisburg, PA, USA, in 2020. He is currently working toward the Ph.D. degree in mechanical engineering with Virginia Tech, Blacksburg, VA, USA. He is currently a Research Assistant with Collaborative Robotics Laboratory, Department of Mechanical Engineering, Virginia Tech. His research interests include robot learning, human-robot teams, and artificial intelligence.



Carly A. Mendenhall received the B.S. degree in mechanical engineering in 2022 from Purdue University, West Lafayette, IN, USA, where she is currently working toward the Ph.D. degree in mechanical engineering. Since 2021, she has been working jointly with RAAD Laboratory and Tepole Mechanics and Mechanobiology Laboratory, Purdue University, where she is currently a Research Assistant. Her research interests include soft haptics, soft robotics, and computational modeling of soft materials.



Francesco Fuentes received the B.S. degree in mechanical engineering from Marquette University, Milwaukee, WI, USA, in 2020. Since then, he has been working toward the Ph.D. degree in mechanical engineering with RAAD Laboratory, Purdue University. His research interests include soft robotics, bio-inspired robots, and growing robots.



Ritish Shailly (Graduate Student Member, IEEE) received the B.S. degree in mechanical engineering in 2021 from Virginia Polytechnic Institute and State University (Virginia Tech), Blacksburg, VA, USA, where he is currently working toward the M.S. degree in mechanical engineering. From 2021 to 2022, he was with Collaborative Robotics Laboratory, Virginia Tech. His research interests include haptics, robot learning, and human perception.



Dylan P. Losey received the bachelor's degree from Vanderbilt University, Nashville, TN, USA, in 2014, and the Doctoral degree from Rice University, Houston, TX, USA, in 2018. He is currently an Assistant Professor with the Mechanical Engineering Department, Virginia Tech. His research interests include learning and control algorithms for robots that interact with people. He was previously a Postdoctoral Scholar with Stanford University, Stanford, CA, USA. He was the recipient of the 2020 Conference on Robot Learning Best Paper Award and 2017 IEEE/ASME TRANSACTIONS ON MECHATRONICS Best Paper Award.



Laura H. Blumenschein received the B.S. and M.S. degrees in mechanical engineering from Rice University, Houston, TX, USA, in 2015 and 2016, respectively, and the Ph.D. degree in mechanical engineering from Stanford University, Stanford, CA, USA, in 2019, under the supervision of Prof. A. Okamura. She is currently an Assistant Professor of mechanical engineering with Purdue University. Her research interests include soft robotics, actuator design, modeling, haptics, and growing robots.