

Data-driven sparse skin stimulation can convey social touch information to humans

M. Salvato, Sophia R. Williams, Cara M. Nunez, Xin Zhu, Ali Israr, Frances Lau, Keith Klumb, Freddy Abnoui, Allison M. Okamura, Heather Culbertson

Abstract—During social interactions, people use auditory, visual, and haptic cues to convey their thoughts, emotions, and intentions. Due to weight, energy, and other hardware constraints, it is difficult to create devices that completely capture the complexity of human touch. Here we explore whether a sparse representation of human touch is sufficient to convey social touch signals. To test this we collected a dataset of social touch interactions using a soft wearable pressure sensor array, developed an algorithm to map recorded data to an array of actuators, then applied our algorithm to create signals that drive an array of normal indentation actuators placed on the arm. Using this wearable, low-resolution, low-force device, we find that users are able to distinguish the intended social meaning, and compare performance to results based on direct human touch. As online communication becomes more prevalent, such systems to convey haptic signals could allow for improved distant socializing and empathetic remote human-human interaction.

Index Terms—social communication, data-driven actuation, tactile devices, social distancing

1 INTRODUCTION

SOCIAL touch is a natural mode of communication between humans, and its importance is becoming more evident as we increase remote communication through videoconference, email, and text messages. Without the depth and subtlety achievable in person, remote communication causes increased feelings of loneliness and social isolation [1], which can have a significant impact on mental and physical health [2]. The need to social distance [3] exacerbates this issue, and further highlights the importance of more expressive remote communication.

Previously, researchers have sought ways to virtually replicate the feelings of social touch through wearable and holdable haptic devices. Several methods of conveying touch have been explored, such as vibrations [4], thermal stimulation [5], and air jets [6]. However, most social touch haptic devices have focused on replicating a single form of social touch, such as a handshake [7], a hug [8], [9], or a stroke on the arm [6], [10], [11], [12], [13]. In addition, most of these devices leverage manually generated signals, often optimized for simplicity of description (such as frequency and intensity), rather than conveying the richness and subtlety of touch interaction. Humans, on the other hand, have shown moderate success interpreting touch cues from partners even without context [14], [15], and their social touch signals can be complex and varied. Even devices that directly map human interaction from a sensor to an actuator

have relied on tight coupling between the sensing and actuation methods [7], [10], [16]. Rather than focusing on specialized hardware, our method seeks to richly represent a range of social touch messages by using a data-driven method to leverage recorded human social touch data. This allows for complex and varied social touch signals using a variety of recording and actuation devices.

Social haptic devices vary in the type of contact they achieve with the skin. Some devices are limited to applying haptic stimulation, such as vibration, heat, or pressure, to a local region of the skin [4], [5]. On the other end of the spectrum are those which can move continuously along the skin, applying stimulation anywhere over a range of skin [6], [10]. We focus on what we define as “sparse” devices, which apply haptic stimulation multiple, separate, local regions of the skin [11], [13], [17], [18]. We focus on sparse devices because they can be less mechanically complex than continuous devices, but more expressive than applying stimulus to only a single local skin region.

This paper explores whether social touch data can be represented using sparse actuation given appropriate processing, allowing for the creation of systems that can represent a variety of social touch messages mapped directly from recorded human interactions. We develop a consistent and generalizable, data-driven algorithm to map recorded human touch to a sparse representation. We leverage multi-object tracking techniques [19] to find data likely to represent sustained or high-pressure touch interactions, then find optimal regions to render on a haptic system. We display the sparse touch signal with eight fixed points of contact, via an arm-worn array of one-degree-of-freedom voice coil actuators, each 37-50 mm apart. These contact points are at a larger distance than the threshold for discrimination by afferents in the skin, based on direction discrimination tests [20]. We tested participants’ ability to recognize social touch scenarios based on passively felt signals produced by the actuator array.

- M. Salvato (corresponding author: msalvato@cs.stanford.edu; web.stanford.edu/~msalvato) and A. M. Okamura are with the Department of Mechanical Engineering, Stanford University, Stanford, CA 94305.
- S. R. Williams is with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305.
- C. M. Nunez is with the Departments of Bioengineering and Mechanical Engineering, Stanford University, Stanford, CA 94305.
- X. Zhu and H. Culbertson are with the Department of Computer Science, University of Southern California, Los Angeles, CA 90089
- A. Israr, F. Lau, K. Klumb, and F. Abnoui are with the Facebook Inc., 1 Hacker Way, Menlo Park, CA 94025

This work was performed as a 2-stage study. In the first stage (Section 2), we created a novel social touch dataset. This dataset involved human-human interaction with a provided scenario, gave no instruction on which gestures the human should use, and recorded force data directly. In related studies, human social touch interaction was recorded with defined gestures on a mannequin [21] and small robot [22]. These data involved simple gestures, with explicit instruction to the user of how each gesture was to be performed. They were collected with the primary goal of classification [23], rather than generative modeling. Hauser et al. [14] used camera-based recording and electromagnetic trackers to measure touch direction and movement associated with six thoughts to be conveyed for open-ended interaction. However, the study lacked force sensing, and used only single-word prompts. We provide scenarios so participants can touch in an informed way with similar context across pairs.

In the second stage of our study, we applied a mapping algorithm to instances from our social touch dataset (Section 3) and examined if a new set of participants could use the sensation they felt to determine the scenario the data was drawn from (Section 4). We found that participants achieved 45% accuracy on 6 scenarios, in comparison with 57% accuracy found for human-human interaction between close friends and partners [14]. This suggests that social touch signals can be represented via sparse skin stimulation.

The contributions of this work are as follows. The social touch classification accuracy in our human subject test was achieved via the use of a novel social touch dataset that was passed through a data-driven mapping algorithm and rendered on a sparse haptic device. We describe a detailed account of our data collection, and provide the resulting dataset for public use. We also present a novel rendering algorithm and provide the associated code to facilitate future work in haptic rendering. Because our patterns were found to be interpretable, we posit it is possible that sparse touch signals are sufficient for humans to classify social touch.

2 A NATURALISTIC SOCIAL TOUCH DATASET

To analyze human social touch and create artificial touch signals, we collected a naturalistic social touch dataset. Participants were recruited in pairs who self-identified as either close friends or romantic partners who felt comfortable interacting with each other through touch. We recruited participants through Stanford University email lists that included hundreds of students and postdocs enrolled in various programs of study. Participants gave informed consent, and the protocol was approved by the Stanford University Institutional Review Board (protocol #22514). Twenty pairs of partners were recruited, for 40 participants total (17 male, 23 female, ages 22-34 years old). Thirty-five participants self-reported that they were right-handed, three participants that they were left-handed, and two participants that they were ambidextrous. Eleven of the pairs were close friends. Out of these, three pairs were male-female, three were male-male, and five were female-female. Nine of the pairs were romantic partners. Out of these, eight pairs were male-female and one pair was female-female. No male-male romantic partners were tested. All pairs included at least one participant who was affiliated with Stanford.

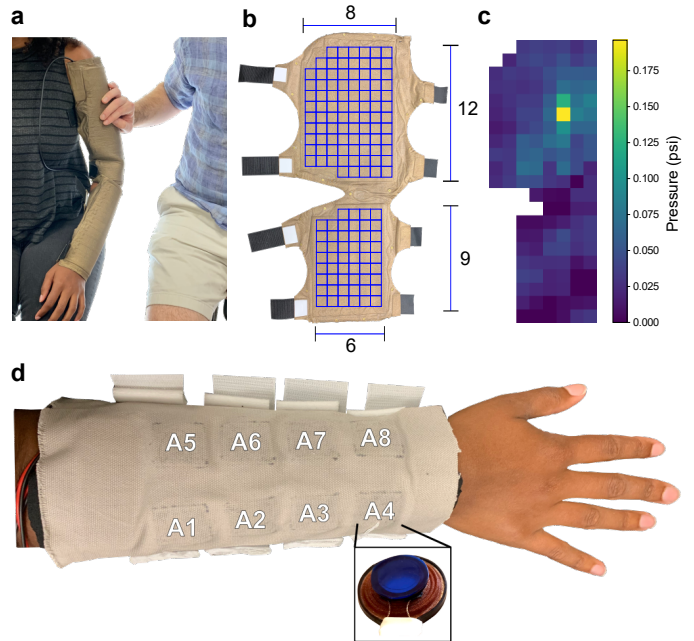


Fig. 1. Data recording and actuator hardware. (a-c) Social touch data recording setup and sensor. (a) Image of the experimental setup. One participant wore a pressure sensor sleeve while the other participant touched the sensor sleeve arm to express social meaning. Complementary social scenarios were played over headphones to participants to cue the social touch. (b) Image of a sensor sleeve. The sensor sleeve is shown laid flat, with squares indicating the individual sensor cells. Each cell is 1 in². This is the larger of two sleeves used during the study, to account for differing arm sizes. The smaller sleeve has upper arm dimension 8x10 and lower arm dimension 6x8. (c) An image of a single data frame from the sensor. The sensor can record pressure as high as 2.96 psi at 20 Hz. (d) Actuator sleeve. An image of a user wearing our actuator sleeve, with an example actuator. The actuator is covered with a thin piece of rigid plastic. While the sensor covers the upper arm and forearm, the actuators are limited to the forearm.



Fig. 2. Image of social touch recording room. The participants sat on the two provided stools. All lights in the room were color changing bulbs, which changed hue based on the prompt.

A dataset entry involves one participant wearing a soft pressure sensor on their arm, and another touching them on the worn sensor to express a response to a described scenario (Fig. 1a). We used detailed scenarios in an attempt to keep the participants maximally invested, and so all pairs imagined similar contexts. We attempted to create a setting where participants were comfortable expressing themselves by creating a room where subjects were isolated from the experimenters, artificial plants and a rug were placed in the room, and floor and lamp lighting were provided (Fig. 2).

TABLE 1

Scenario identifiers and associated touch meaning categories. Our 6 scenarios span 5 meaning categories. For the repeated category (“Announcing a response”) we have a positive and a negative valence example. These terms motivated our scenarios and were presented to the toucher and receiver prior to the scenario.

Scenario Identifier	Touch Meaning Category
Attention seeking	Attention-getting
Gratitude	Appreciation
Happiness	Announcing a response (positive valence)
Calming	Support
Love	Affection
Sadness	Announcing a response (negative valence)

We focused on the interaction from one person’s hand to another’s arm, based on prior work of socially acceptable places to give and receive social touch [24] and informal pilot testing. We placed no restriction on the duration or nature of their movements beyond interaction on the arm to prevent the subjects from feeling constrained in their actions. With this population we expect that participants were able to map the provided scenario to similar ones they previously experienced with that partner. This data collection scenario is unique compared to other datasets [21], [22], which involve a human interacting with a non-living object. While [17] similarly uses recorded touch on a human-worn sensor, they were performed by the experimenter in a controlled setting and limited to simple gestures. They refer to this as a controlled “best case” recording scenario.

We based our set of scenarios on previous work studying direct social touch between individuals. [26] found that social touch can be divided by meaning into 12 categories, such as support touches, appreciation touches, and affection touches. To have a variety of scenarios, we leveraged prior work in emotion communication through touch [15], [27], and selected 6 scenarios spanning 5 of the categories presented by [26] (Table 1). For the repeated category we have a positive valence and a negative valence example. Our final set of scenarios is: attention seeking, gratitude, happiness, calming, love, and sadness.

For each pair of participants, one participant, the “receiver”, wore a soft pressure sensor wrapped on their arm. For each scenario, the lighting level and color changed slightly to match the scenario (Fig. 2). The receiver and the other participant, the “toucher”, each listened to an audio prompt. The receiver listened to a prompt that prepared them to be touched, and the toucher listened to a different prompt describing the scenario and which ended by telling them to touch the receiver. The toucher touched the receiver on the sensor array using one hand. The toucher was told to touch for as long as they wished, and to return their hands to their lap when finished (Supplementary Mov. SM1). This procedure was repeated for each scenario presented in pseudorandom order in three sets, for a total of 18 interactions. The process was then repeated with the toucher and receiver roles switched. While we refer to scenarios by single word identifiers, our results apply to the *scenarios* presented to the participants. Text and audio recordings of the prompts are available at Supplementary Information: Social Touch Dataset and Supplementary Dataset SD1 respectively. As an example, below is the *love scenario* presented to the toucher:

Love. Imagine this: You and the person sitting next to you are spending the afternoon together. You’re walking

to get a bite, the weather is amazing, and you’re catching up on everything in the way that friends do. You look at them, and it suddenly strikes you how much this friendship means to you, that life is so much easier and better with them around. Reach out and touch this person to express your love for them.

Our protocol was designed in collaboration with IDEO (Palo Alto, CA, US <https://www.ideo.com/>). Behavioral scientists at IDEO developed storyboards of the scenarios with the goal of creating an immersive experience and eliciting authentic touches. Various scenarios were iterated in preliminary tests to be natural and compelling. Through these iterations, a number of design choices were made, including: manually advancing scenarios by the experimenter when the users returned their hands to their laps; placing temporary walls around the participants so they felt alone in the room, comfortable and fully immersed; and telling participants that they could make facial expressions, but not communicate verbally.

The touch data were recorded using a custom soft capacitive sensor sleeve created by Pressure Profile Systems (Los Angeles, CA, US). The sensor wraps around the upper arm and forearm, and records pressure data in 1 in² cells at 20 Hz. The device has a range from 0psi to 2.96psi with 0.004 ± 0.001psi measured resolution. It has calibrated linearity of 99.9 ± 0.1% and signal-to-noise ratio of 732 ± 219. We found when the sensor was laid flat with no contact, pressure readings did not exceed ±0.08psi. We used two sensor sizes, with 120 or 142 sensor cells (Fig. 1b). The sensors had different lengths and circumferences, with the sensor chosen based on fit for each participant’s arm. Fig. 1c shows a single example data frame from the sensor. We also recorded the interaction from two different views using RGB cameras.

Due to recording errors, the data for three participants was lost, resulting in 37 participants’ worth of recorded data. A total of 661 instances were recorded (5 were removed due to additional recording error). Each touch can be described as a series of short gestures (such as “poke” or “stroke”). The experimenters visually inspected the video and time-synchronized pressure sensor waveforms to annotate the recorded data with gestures based on the gesture descriptions in [15], using a subset of gestures from that work. Fig. 3a shows the set of gestures with total number of times each gesture was seen across all instances of each scenario, and Fig. 3b-d shows statistics about touch location, pressure, and duration. To test the importance of the relationship between the toucher and receiver on touch motion, we checked the significance of relationship type (“friends” or “romantic partners”) for each of these statistics. We found that touch duration was significantly long for those who identified as romantic partners using the Kruskal-Wallis test ($\chi^2(1, N = 222) = 37.23, p = 1.0e-9$). Those in romantic relationships may also touch the lower arm more often ($\chi^2(1, N = 222) = 3.76, p = 5.2e-2$). We chose the Kruskal-Wallis test as responses were independent, but the data was non-normal and not homoscedastic.

Before the experiment, each participant completed a survey about their experience and preferences in expressing social touch (Supplementary Fig. S1). Survey data was obtained for all 40 participants. Raw data of all sensor

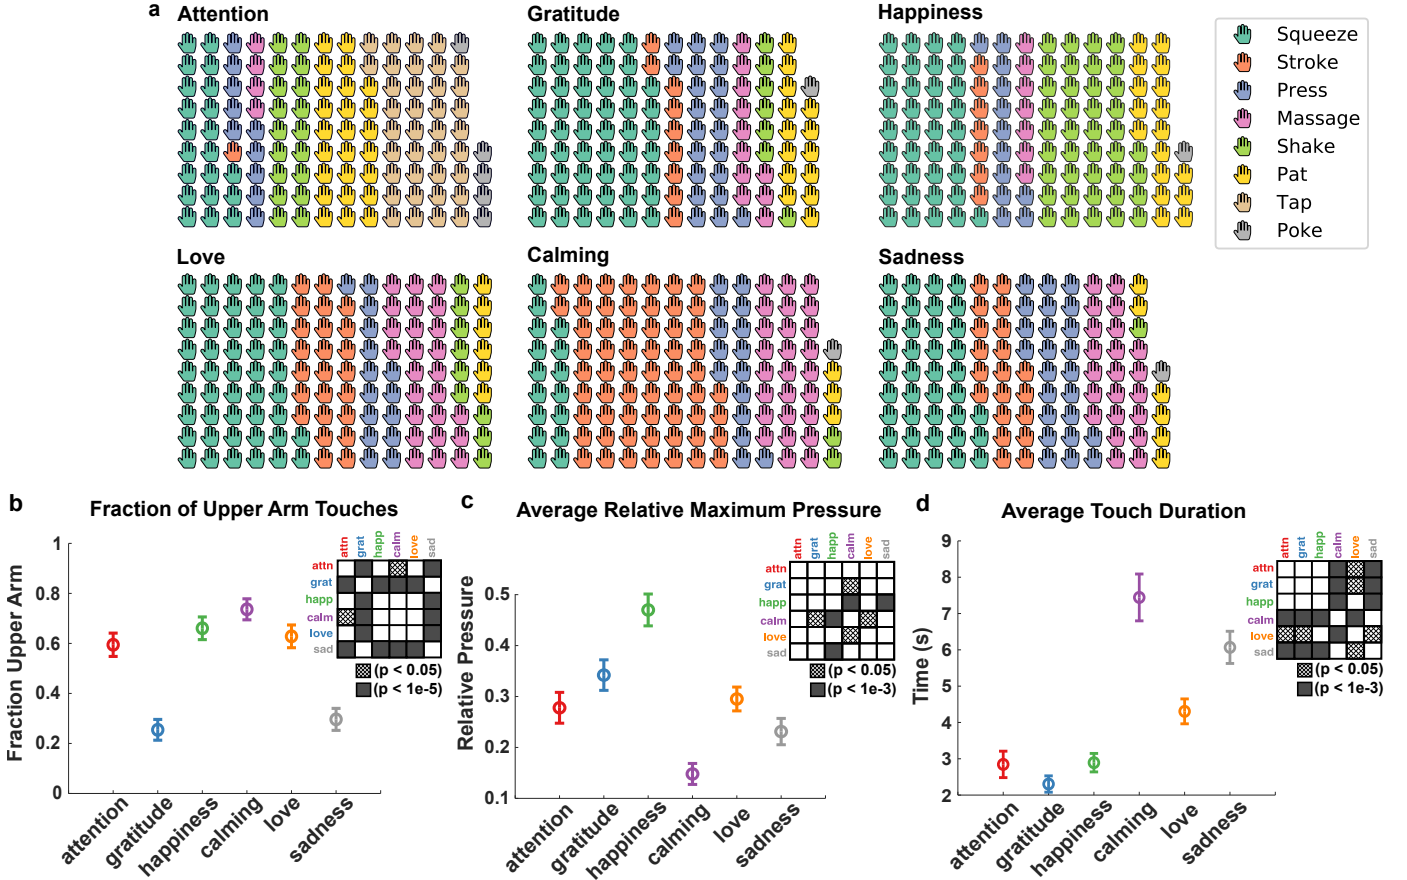


Fig. 3. Dataset gesture counts and per scenario metrics. (a) Gesture counts. For each touch, we manually labelled the gestures apparent in the touch. This plot shows how many times each gesture appeared for each touch. Patterns appear in the data, such as a large number of taps and pats for attention, strokes for calming, and shakes for happiness. (b) Touch location plot. The fraction of touches for which more frames have maximum value on upper arm as opposed to lower arm. Using Friedman test we find the relationship between scenario and upper arm count was significant ($\chi^2(5, N = 666) = 98.06, p = 1.4e-19$). The associated significance matrix shows the probability two scenarios differ and was calculated using pairwise Friedman tests. Note that Friedman is equivalent to Cochren Q and McNema tests for the overall and pairwise data respectively [25]. (c) Relative maximum pressure plot. For each participant we determined the maximum pressure for each of their 18 touch instances. We then normalized these values to obtain a relative maximum pressure for each instance for each participant. Using Friedman test we find the relationship between scenario and relative maximum pressure was significant ($\chi^2(5, N = 666) = 83.01, p = 2.0e-16$). The associated significance matrix shows the post-hoc pairwise probability two scenarios differ and was calculated using the Dunn-Sidak test. (d) Touch duration plot. Participants were free to take as long as desired when touching. Using Friedman test we find the relationship between scenario and touch duration was significant ($\chi^2(5, N = 666) = 285.81, p = 1.1e-59$). The associated significance matrix shows the post-hoc pairwise probability two scenarios differ and was calculated using the Dunn-Sidak test. The error bars on (b), (c), and (d) show standard error. We selected our statistical tests as we found the data as non-normal, not homoscedastic, and does not have independence (because each subject responds to every condition)

recordings and gesture annotations is provided in a free online dataset (Supplementary Dataset SD2) to be leveraged for future research into both naturalistic remote human-human communication and human-robot interaction.

3 DATA TO ACTUATOR MAPPING

We developed an algorithm to map recorded touch data to signals that can be rendered on a haptic device (Fig. 4). This algorithm assumes the haptic device consists of an array of actuators to render touch information, and the number of actuators is smaller than the number of sensors. Thus, we reduce the sensor data to the dimensionality of the actuators by finding consistent, high pressure trajectories through the data, and render a selection of those trajectories with the actuator array. For this project we used this algorithm to map to a 2×4 array of voice coil actuators (Fig. 1d, Sec. 3.1). Our algorithm is akin to a multi-object tracking problem (Step 1), with actuator specific restrictions (Step 2).

Step 1 is to find the most salient, contiguous signal within the data. This step is independent of the haptic device and produces a set of signal trajectories through the data. We base this on a multi-object tracking algorithm [19]. This step relies on two assumptions. First, we assume that for a set of data frames the intensity of a point increases monotonically with the probability of interaction with the device. This assumption allows us to capture trajectories with higher intensity. Second, we assume that when points close in time are of similar intensity, they are more likely to represent a sustained interaction with the device. The multi-object tracking algorithm works by solving an optimization problem for which longer, greater intensity paths are weighted more highly. Fig. 4c shows a sampling of frames from a tracked path. Fig. 4b shows corresponding frames sampled from the raw data.

Step 2 is to find the optimal set of trajectories to render on the haptic device. This step requires each actuator to have a defined workspace, making it dependent on the actuator

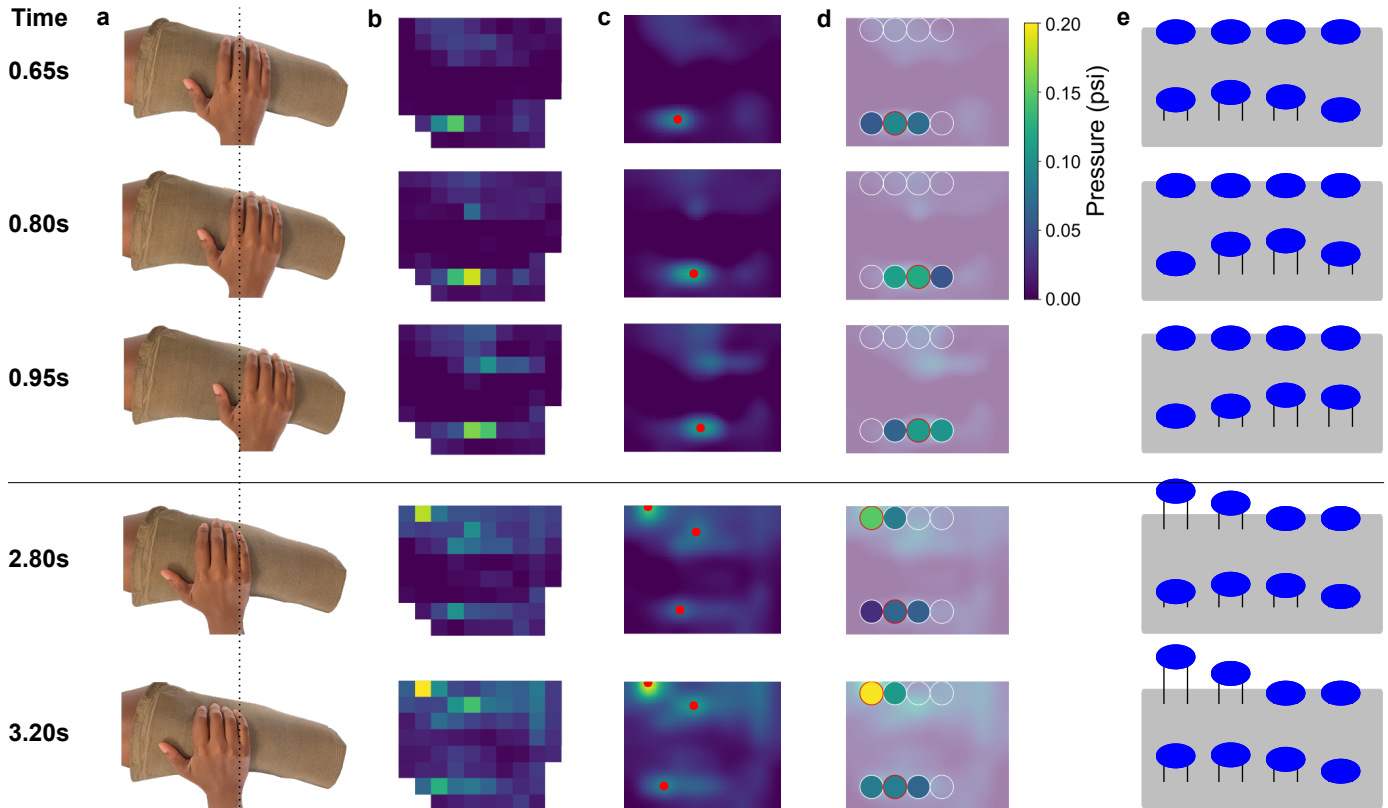


Fig. 4. Data mapping algorithm. Here we show representative key frames from our data mapping algorithm. The actual data was processed at 20 Hz. (a) A user interacting with the pressure sensor sleeve. The top three frames show a sampling of a stroking motion, while the bottom two frames show a sampling after the user transitions to a squeezing motion. (b) Example data frames from the sensor sleeve. This is sensor data from a ‘love’ example in our dataset – it is not drawn from (a) itself. (c) The data from (b), after preprocessing and multi-object tracking. The red dots represent tracked paths. This is the result of Step 1 of our algorithm. (d) Actuator restricted tracking. The circles represent the optimal areas to render, given the tracked paths. The fill color represents the intensity to render. The red circle contains the trajectory and renders its intensity. This is the result of Step 2 of our algorithm. We also render the maximum intensity within the circles which neighbor the trajectory for continuity. The pressure scale for (b-d) is shown to the right of column (d). The algorithm is run only once on the data and generates both the stroke and squeeze motions. (e) Example voice coil activity. This is a representation of our 2×4 voice coil sleeve, with each blue circle being a normal indentation actuator. The coils have height proportional to the intensities of each workspace in (d).

hardware being used. Workspaces are specified as 3D geometric shapes, for which we require that only one trajectory pass through a given workspace at a given time. For our actuators these are cylinders, which we represent as circles in Fig. 4d. To optimize the rendering, we determine a placement of the actuator workspaces within the trajectory space. We determine the placement by testing different translations of the workspaces in the trajectory space, optimized by a scoring function based on the probability of elements in a trajectory and their locations in the workspace. Depending on the actuators and desired use case, the workspaces could be tested under other transformations, such as rotation, scaling, or separation. Fig. 4d shows the trajectories tracked from Fig. 4c, optimally mapped to the workspaces for our haptic device. Fig. 4e shows an example set of actuators rendering those trajectories. We see the algorithm captured both a stroking motion (top) and squeezing motion (bottom).

We could have merged Step 1 and Step 2 into a single optimization where the trajectories are selected with information about the actuator workspaces provided. However, we separated these procedures so that Step 1 represents the trajectories strictly as a function of the data, and therefore is a more accurate representation of the data itself. We also render on actuators that neighbor the actuators for which the trajectory is specified, and perform other processing for

smoothness. Below we provide pseudocode and further explanation. Supplementary Information: Mapping Algorithm Formalization provides a full algorithm formalization and Supplementary Dataset SC1 provides the code.

3.1 Actuation Hardware

Our actuator sleeve consists of eight voice coils (Tectonic Elements TEAX19C01-8), arranged in a 2×4 pattern (Fig. 1d, Supplementary Mov. SM2). Coils were placed approximately 37 mm apart in each row, and 50 mm apart in each column. The sleeve has an inelastic fabric backing, with softer elastic fabric covering the backs of the voice coils. The sleeve was wrapped around the forearm and affixed using velcro strips, so the voice coils sat on the dorsal side of the arm. The coils were controlled with a PCI board (SENSORAY Model 826) at 1000 Hz, through a linear current amplifier with a gain of 1 A/V via an LM675T op-amp.

3.2 Algorithm Pseudocode

Algorithm 1 shows a pseudocode of the mapping algorithm. The algorithm takes 3 inputs: a series of pressure frames F , a list of actuator workspaces A , and a list of actuator transforms M . Each frame $f_i \in F$ consists of 2D coordinates and 1D pressure identical in structure to a single channel image where the pixel value is the pressure. We will refer

Algorithm 1 Pressure Data to Actuators Trajectories

```

1: procedure MAPPINGALGORITHM( $F, A, M$ )
2:                                     ▷ Pressure Data Frames  $F$ , Actuator Workspaces  $A$ , Actuator Transforms  $M$ 
3:    $\mathcal{G} \leftarrow \{\}$ 
4:   for all  $f_t \in F$  do                                     ▷ Get local maxima points across frames.
5:      $\mathcal{G} \leftarrow \mathcal{G} \cup \text{FindMaxima}(f_t)$ 
6:   end for
7:    $\mathcal{J}, \mathcal{R} = \text{MultiObjTracking}(\mathcal{G}, F)$                  ▷ Find trajectory set  $\mathcal{J}$ ; determine associated scores  $\mathcal{R}$  based on trajectory element intensities.
8:   for all  $m \in M$  do                                       ▷ For each transform, find best trajectories and associated scores
9:      $\mathcal{J}_m, \mathcal{R}_m \leftarrow \text{ConstrainedTrajs}(\mathcal{J}, \mathcal{R}, m(A))$    ▷ when restricted to one trajectory per actuator per timestep.
10:  end for
11:   $\mathcal{J}^* \leftarrow \mathcal{J}_i$  s.t.  $\sum_{R_i \in \mathcal{R}_i} R_i \geq \sum_{R_j \in \mathcal{R}_j} R_j, i \neq j$    ▷ Select best scoring restricted transformed trajectories.
12: end procedure                                             ▷ In case  $\sum_{R_i \in \mathcal{R}_i} R_i = \sum_{R_j \in \mathcal{R}_j} R_j, \mathcal{J}^*$  chosen arbitrarily.

```

to these coordinates as pixels. These frames do not need to be rectangular. Each $a_i \in A$ is a 3D shape representing the workspace of actuator i . In our work the workspaces are cylinders (seen as circles in 2D). The actuator transforms $m \in M$ are functions which inform how we can place the actuator workspaces relative to the pressure frames. They can consist of operations such as translation, rotation, or scaling, but we only use translation in this work.

Lines 3-5 the algorithm find the set of all local maxima coordinates \mathcal{G} for the data frames F . These are the pixels which have higher pressure than each of their spatial neighbors for the frame in which they are located. This is so the following step is computationally tractable.

Line 6 addresses the multi-object tracking algorithm portion. This is Step 1 described above. Multi-object tracking algorithms are used in computer vision to maintain a consistent labeling of identified objects in a scene across multiple frames [19]. These assume a mechanism already exists to identify the objects. For example, in video of multiple people walking in a busy street, a detector may be used to identify people, then multi-object tracking could be used to track each individual across frames. In our work we seek to find “trajectories,” which are sequences of pixels through time that represent sustained, high pressure data. We assume these represent the movement of fingers or the palm on the sensor. As an input, we take the local maxima coordinates \mathcal{G} and data frames F . We output salient trajectories \mathcal{J} in the data, which we use to represent the movement of the fingers or palm. Each trajectory also has an associated score R , with the set of scores denoted \mathcal{R} for associated trajectories \mathcal{J} . We create a set of trajectories \mathcal{J} that are tracked through the data frames $f_t \in F$. A single trajectory $J \in \mathcal{J}$ consists of a sequence of pixels, p_1, \dots, p_l , through frames f_s, \dots, f_{s+l-1} . $p_i \in f_{i+s-1}$. l represents the length of the trajectory, and s the time of the first frame in the trajectory. A trajectory consists of only contiguous frames, and a single trajectory cannot consist of more than one pixel in a frame. Additionally the pixels are only those in \mathcal{G} – those which were local maxima from lines 3-5.

In order to find the trajectories in line 6, we adapt a multi-object tracking algorithm [19] with the following assumptions. First, we assume that for a set of data frames the intensity of a point increases monotonically with the probability of interaction with the device. This assumption allows us to capture trajectories with higher frame intensity. Second, we assume that when points close in time are of similar intensity, they are more likely to represent a sus-

tained interaction with the device. The multi-object tracking algorithm works by solving an optimization problem for which longer, higher intensity trajectories are weighted more highly. The scores $R \in \mathcal{R}$ are also based on the length and intensity of the trajectories (see Supplementary Information: Algorithm Formalization for exact calculation).

Lines 7-8 address the actuator workspace restriction, the first part of Step 2 of our algorithm. This step finds the best sets of trajectories for which each actuator is constrained to act on only one trajectory at a time. For each call inside the loop, we take as input the trajectories \mathcal{J} , scores \mathcal{R} , and the actuator workspaces. For each iteration of the loop, a single transform $m()$ is applied to all workspaces in A . Considering the trajectories as tracking movement in Euclidean space, the actuator workspaces are placed in this space such that the tracked trajectories describe movements for the actuators. Each transform $m()$ then represents a transformation, such as translation, of the set of actuators in the space, so the actuators will render different parts of the trajectories based on the transform. The algorithm then calculates the best subset of the trajectories for that transform, \mathcal{J}_m , for which no two trajectories exist in the same workspace at the same time, using the sum of trajectory scores. The trajectory scores are affected by a small weighting based on the actuator workspaces, as explained in Supplementary Information: Algorithm Formalization.

Line 9 then completes Step 2 by selecting the trajectories output by the transform on line 8 which resulted in the highest summed trajectory score. We take this opportunity to highlight why it is important to test different actuator transforms. As an example, consider four consecutive workspaces, and one trajectory long enough to pass through all of them. If the trajectory starts in the first actuator, it would render through all four. However, if the trajectory started in the second actuator, the last segment of the trajectory would not be rendered, since it would be placed in the space beyond the actuators. This problem becomes more difficult with multiple trajectories and the restriction of one trajectory per actuator per timestep, necessitating the development of Step 2 of our algorithm.

3.3 Data Processing

For our system, to preprocess the recorded data, we first upsample the data by a factor of seven. We then perform a 2D Gaussian blur with standard deviation of three.

After the algorithm determines which actuators render the trajectory points at each time, we also examine the

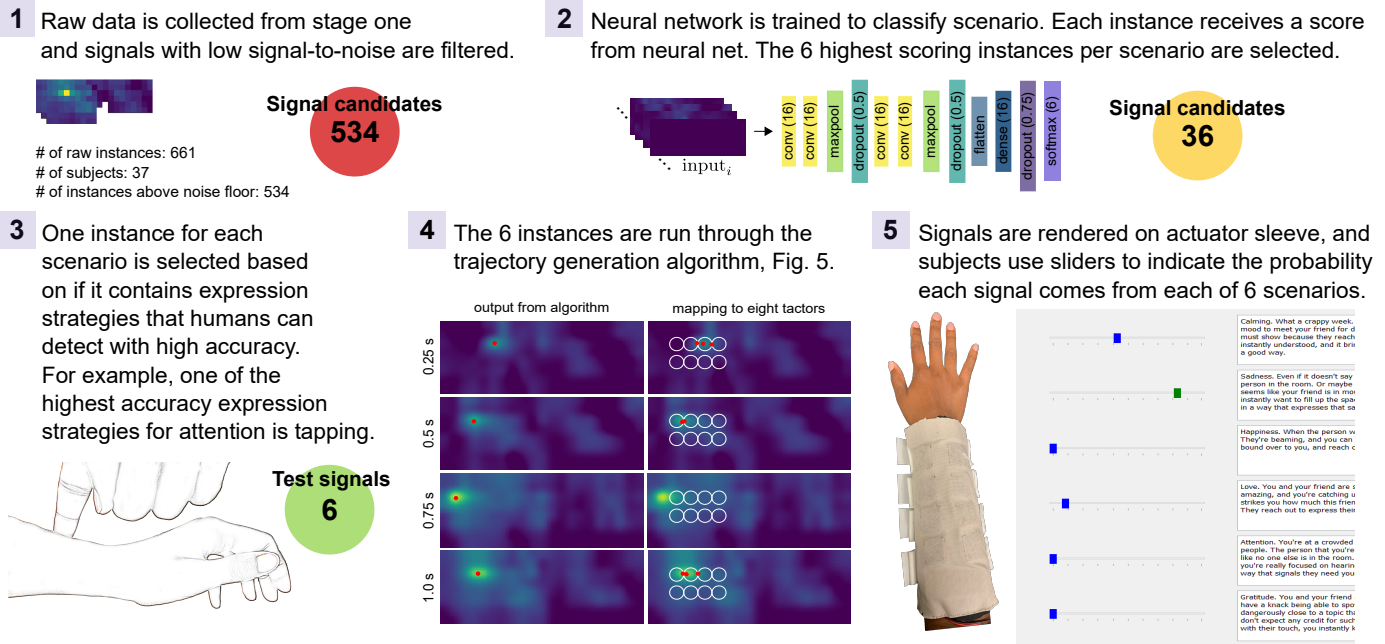


Fig. 5. Process to integrate data, mapping algorithm, and hardware for scenario classification experiment. In step 1, we removed sequences for which 5 or fewer total readings exceed .08 psi.

row-wise neighboring actuators. Each neighboring actuator which is not already rendering a trajectory will render the maximum sensor intensity within their workspace area. This is to effectively smooth the signal, which creates continuous sensations for gestures such as stroking for our system.

After producing the the signals for each actuator to render, we remove sharp peaks in the data. We do this by finding triplets of points that rise and drop by $\frac{1}{8}$ the total pressure range between each pair of frames and setting them to the mean of the first and third point. We then filter the whole signal with a 4 Hz, 5th order Butterworth lowpass filter. During informal pilot testing prior to implementing these measures, participants found it difficult to parse the signals due to jitter, and we saw limits on the speed of the voice coils. This filtering removes the jitter while still maintaining signal information.

4 SCENARIO CLASSIFICATION EXPERIMENT

4.1 Methods

We conducted an experiment to determine whether participants could classify haptic signals generated by applying our algorithm to our dataset, felt through the voice coil sleeve hardware (Fig. 5). Six signals were chosen and used in the experiment (Sec. 4.2), one for each of the six scenarios: attention seeking, gratitude, happiness, calming, love, and sadness. For each signal subjects selected which scenario it was derived from. Participants gave informed consent, and the protocol was approved by the Stanford University Institutional Review Board (protocol #22514). 30 participants engaged in the experiment (14 female, 16 male, ages 19-35 years old). Participants in the social touch dataset collection were excluded from this experiment.

During the experiment, participants listened to the toucher audio for each of the six scenarios. They then passively felt each of the six haptic signals once, with no

indication of what the signals represented, so participants were acclimated to the range of signals they would feel. They were then presented each signal a total of three times, in randomized sets of each of the six signals. For each signal, participants were asked to provide the probability that they believed the sensation was drawn from each of the six scenarios using a slider bar. Participants were informed that the values would be normalized to one, and were required to not have an exact tie for the largest value. During this phase of the experiment, subjects were shown text transcriptions of the scenarios to serve as a reminder of the audio. While they were provided the scenario identifier term, they were asked to focus on the scenario itself. In the last phase of the experiment, participants felt each signal once more and were asked to rate the valence and arousal of each signal, using the Self-Assessment Manikins [29]. Raw data of participant’s responses is provided for free online (Supplementary Dataset SD3). The toucher audio prompts were presented to the participants instead of the receiver audio prompts, because the participant’s task was to determine the intent of toucher (device), which is not necessarily encoded in the receiver audio prompts.

To test the number of subjects required for a moderate effect size significant relationship between provided scenario and response, we used a χ^2 test with 25 degrees of freedom. To get moderate effect size of $w = 0.3$ with power $1 - \beta = 0.95$ and error probability $\alpha = 0.05$, we would require 370 samples. We collected approximately double this number of (720) scenario classifications.

4.2 Instance selection

For this experiment we selected instances from the social touch dataset to render on our hardware that exemplified common gestures, as determined by our gesture annotations (Fig. 3a) and previous work [14]. Because our dataset had over 500 sequences, we used a 3D convolutional neural

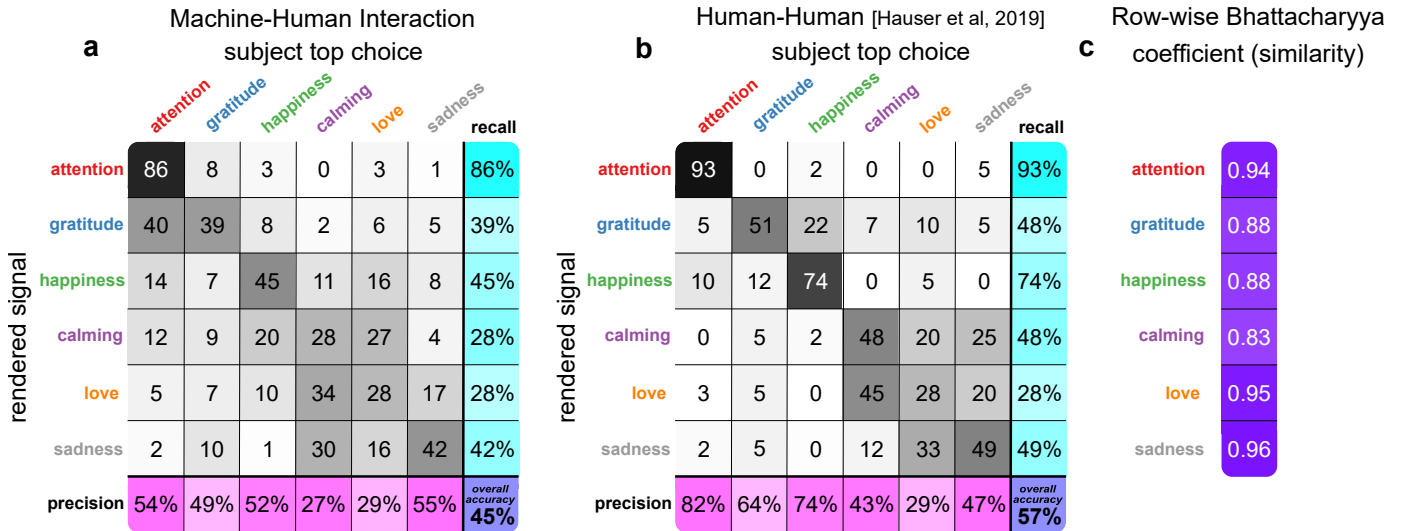


Fig. 6. Scenario classification top choice results. (a) The confusion matrix for the scenario that participants thought was most likely displayed by the rendered signal. This provides information on the participants’ overall accuracy in identifying the scenario the rendered signal was drawn from, how participants confused signals, and participants’ precision and recall for their top scenario choices. Recall ($\frac{\text{True Positive Rate}}{\text{True Positive Rate} + \text{False Negative Rate}}$) gives a sense of how often the rendered signal type was correctly identified when presented, and precision ($\frac{\text{True Positive Rate}}{\text{True Positive Rate} + \text{False Positive Rate}}$) gives a sense of how often that signal type was actually presented when it was chosen. Rows are normalized to 100, with 120 samples per row in raw data. We fit a general linear mixed effects (GLME) model with a binomial distribution and a logit link function where scenario is represented as a 6-level factor and subject as a random effect. The Analysis of Deviance (Type II Wald χ^2 test) shows that the scenario has a significant effect on accuracy ($\chi^2(5) = 87.7$, $\text{Pr}(> \chi^2) < 2.2e - 16$), implying that scenario affected response. We checked linearity and normality of residuals by inspecting Pearson residuals and the q-q plot respectively. On a per scenario basis, binomial tests show that subjects did not choose randomly (at a significance level of $p < 0.002$) for each scenario. (b) Human-human interaction results from [14] for comparison. (c) Row-wise Bhattacharyya coefficients. For each rendered signal row we normalize the values to 1 and then calculate the Bhattacharyya coefficient [28] between the machine-human and human-human row values. This provides comparison of the similarity of responses per scenario. Higher value indicates more similarity, coefficient has range of [0, 1].

network (Fig. 5.2) to classify the instances by scenario, in order to prune the dataset for instances that exemplified the scenario and make instance selection more tractable. We ran 5-fold cross validation to get a classification score for each instance for each scenario. Each sequence was scored based on the fold where it was part of the test set.

The longest sequence we input is 705 data frames. For all sequences shorter than this, we prepended frames with all values set to 0, so the sequence length becomes 705. To address the difference in size of the two pressure sensors, for the shorter sleeve we repeated the column immediately prior to the elbow joint and two columns immediately after the elbow joint. This prevented disruption of the movement along each arm segment. Additionally, because the sensor arrays are not rectangular, each frame was padded with 0s to form a rectangular shape. Therefore, the input data to the neural net was of dimension (705, 8, 21) (Fig. 1b).

Each convolutional layer was 3D and had filter size (1, 2, 2) with stride length (1, 1, 1) and a ReLU activation. Each max pooling layer had pool size (2, 2, 2) and stride length (2, 2, 2). Categorical cross entropy loss was used. The Adam optimizer [30] with learning rate 0.001 was used (along with all other Keras defaults) for 30 epochs.

We found that the neural net attributed high classification scores to touch instances with the gestures we expected, so for each scenario we chose one instance from the top six highest scoring instances for that scenario. These instances were selected based the experimenters’ intuition about which signals would be most interpretable by the subjects. After selecting one expression instance per scenario, we processed each with our data mapping algorithm. We chose

a time range to render for each signal, but otherwise the processing was identical – no parameters varied between signals. This selection process resulted in a set of six total signals, one per scenario (Supplementary Figs. S6-S7).

4.3 Results

Fig. 6a shows a confusion matrix of the highest probability (“top choice”) scenario selected by participants for each trial in the scenario classification experiment. We found participants identified the scenario the signal was drawn from with overall accuracy of 45% (2.7 times the rate of chance). For comparison, Fig. 6b shows the results of a prior study on single-choice scenario classification based on direct human-human touch interaction [14], which had overall accuracy of 57% (3.4 times the rate of chance). For each row of each confusion matrix we normalize the values to 1 to calculate the Bhattacharyya coefficient [28] between the machine-human and human-human row values (Fig. 6c). The attention, love, and sadness responses have larger coefficients than the gratitude, happiness, and calming responses. Additionally, Supplementary Information: Scenario Classification Experiment provides scenario classification top choice results confusion matrix for the first round of responses only for comparison.

Fig. 7 shows the valence/arousal results in a circumplex affect model graph [31]. The medians for sadness, calming, and happiness exist in the same quadrants as the foundational work in the area [31] (love, attention, and gratitude do not appear in [31]).

Each participant assigned a probability that the perceived touch was drawn from each of the six scenar-

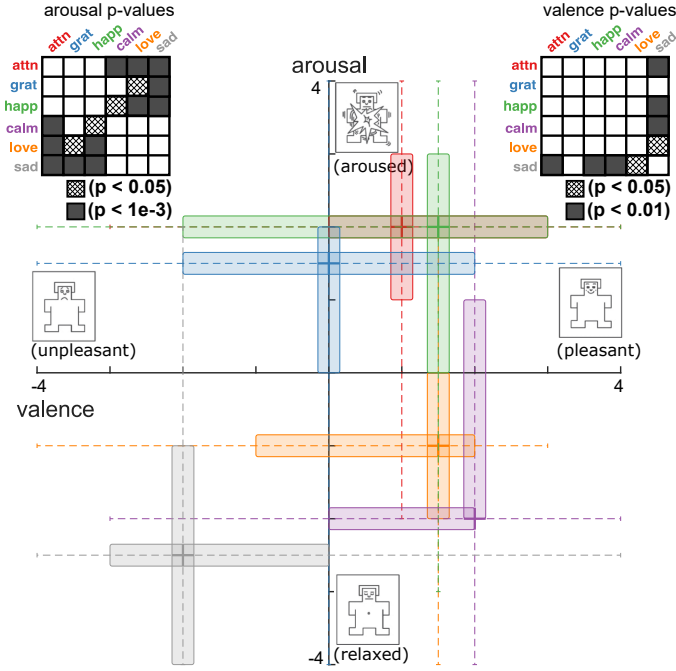


Fig. 7. The valence and arousal ratings for each scenario across participants as box-and-whisker plots. The + symbol for each scenario represents the point of median valence and arousal ratings. To determine differences in arousal, we performed a Kruskal-Wallis test comparing responses per scenario. Scenario was a significant factor ($\chi^2(1, N = 180) = 23.26, p = 3.0e-4$). A similar test was performed for valence, and scenario was a significant factor ($\chi^2(1, N = 180) = 64.16, p = 1.7e-12$). The two tables show post-hoc p-value levels for each pair of scenarios using the Dunn-Sidak test. We choose our tests as the data was not normal or equivariant, but samples were independent.

ios. Fig. 8a shows a data clustering based analysis using the probability assignments. Each assignment can be represented as a 6-dimensional vector with unit L1 norm. We used k-means clustering to find groups of similar response vectors across participants. From both the Calinski-Harabasz [32] and Silhouette [33] clustering evaluation criteria, we found the optimal number of clusters to be seven. Clusters 1, 3, 5, and 7 indicate that a large number of participants assigned each of happiness, attention, gratitude, and sadness with high probability when presented with the respective signal. Cluster 4 shows that some participants assigned high probability for love when the true signal was calming or love, and to a lesser extent happiness or sadness. Cluster 6 shows that some participants assigned high probability for calming when the true signal was calming, love, or sadness. Cluster 2 represents high uncertainty. These clusters indicate that while many participants were certain of attention, gratitude, and happiness, participants also responded that each of calming, love, and sadness were likely to be calming or love in many cases. Fig. 8b shows a 2D t-Distributed Stochastic Neighbor Embedding (t-SNE) [34] plot of the scenario vectors, with a hull of each cluster’s points drawn. This plot provides another means to visualize data from Fig. 8a, as well as intuition on which clusters border each other. Similar clustering analysis is provided on Self-Assessment Manikin data in Supplementary Fig. S5.

The results show that skin stimulation using only eight one-degree-of-freedom actuators can convey social touch scenario information. Our cluster analysis shows where

distinguishing information between signals is possibly lost.

5 DISCUSSION

5.1 Similarity of Machine-Human Scenario Classification with Human-Human Interaction

With no training, humans achieved high scenario classification accuracy using our system: 45% for 6 classes, as compared to 57% accuracy of human-human interaction shown in prior research [14] and 16.7% for chance. In contrast to [14], our experiment was performed using subjects who had no known relationship to the person from whom the signal was generated. Research has shown that individuals in romantic relationships are more accurate at distinguishing emotions than pairs of strangers [35], so signals generated from a partner could potentially achieve higher accuracy, and a lower accuracy than [14] is not unexpected. Additionally, we rendered all signals on the forearm, including those sensed from the upper arm, which may have hindered the naturalness of some signals.

Our results show that it is possible to create signals using natural human expressions for detailed haptic communication. Using Fig. 6 we can compare patterns of recognition between our results and those of human-human touch [14]. We use the Bhattacharyya coefficient [28] as a measure to compare machine-human and human-human response similarity (Fig. 6c). We observe that the response patterns for love, sadness, and attention have higher Bhattacharyya coefficient than the response patterns for gratitude, happiness, and calming. In particular, for true signal gratitude, our results involve less relative confusion between happiness and gratitude, and a higher rate of confusion between attention and gratitude. Our particular prompts likely influenced these confusion patterns. Our gratitude scenario (Supplementary Information: Social Touch Dataset), prompted subjects to touch surreptitiously. We believe this is why the measured gratitude signals are short (Fig. 3d), resulting in short press/squeeze actions which could be confused with attention. This highlights the importance of providing context to more fully understand social touch data and create signals for remote communication.

Cluster analysis can determine common scenario probability assignments for each signal. If all participants drew from the same distribution of assignments, we would expect one cluster for each scenario, with weights equal to the corresponding row of the confusion matrix. However, we see cases where the cluster analysis contains information not in the top choice confusion matrix (Fig. 6a). For example, in top choice analysis for true class gratitude, we observe a nearly equal number of examples for predicted class gratitude and attention. In the cluster analysis (Fig. 8a), for true class gratitude, there is a cluster with high probability assignment for gratitude, but a relatively low probability assignment for attention. This implies that there was a class of people who were confident that gratitude, indeed, was gratitude. Based on the lack of a cluster with a large number of true class gratitude with high probability assignment for attention, we do not see a group of participants who were confident that gratitude was attention. This analysis demonstrates the importance of looking beyond top choice analysis, enabling detailed analysis of where subject confusions exist.

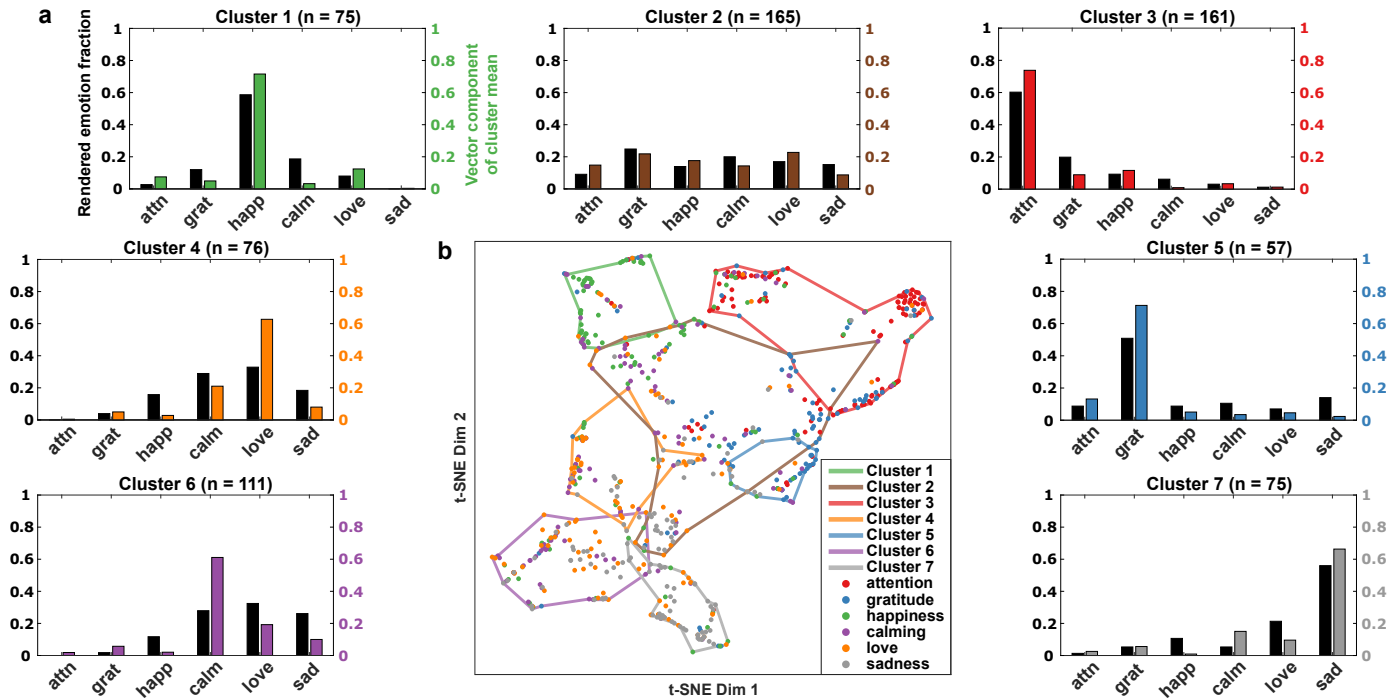


Fig. 8. Probability assignment cluster analysis. Each participant assigned the probability that they felt a given touch was drawn from each of the 6 scenarios displayed to them, providing a 6-dimensional vector with unit L1 norm. Outer figures (a) K-means analysis. The points were clustered to find consistent perspectives across participants. The colored bars show the vector value of the mean of the cluster. This can be interpreted as the black bar showing the presented signal, and colored bar showing the likelihood participants in that cluster attributed the signal to that scenario. (b) t-SNE visualization. We performed t-SNE visualization on all data points to map them to 2D. The point color represents the presented signal. We then drew a hull of the points in each cluster from (a). This provides a two-dimensional visualization of how clusters border each other, in order to provide intuition for what the graphs in (a) represent.

While our task was forced-choice, our probability assignment measure allows us to measure similar information to a non-forced-choice task. In the k-means clustering, Cluster 2 informs us of a group of participants that are likely to be uncertain of the best choice. In Supplementary Information: Scenario Classification Experiment we provide information on top-choice accuracy as a function assigned scenario probabilities to more explicitly simulate non-forced-choice.

5.2 Implications of Sparse Representation

The above results were achieved with a sparse display of haptic signals, using only eight, fixed-location, one-degree-of-freedom actuators placed at larger distances than the threshold for discrimination by afferents in the skin, based on direction discrimination tests [20]. Our results support human the spatial sensitivity of afferents is higher than required to classify social touch signals.

Our work shows that short, sparse haptic signals can convey social touch information. Thus, such signals could be used to create a set of haptic icons, similar to text emojis, for sending social messages. We expect that these signals would be combined with other means of communication, such as video chat, video of a partner interacting via a virtual avatar, or integration with text communication. Given that touch can feel more or less pleasant depending on who administers the touch [36], our signals may be able to convey even more salient affective significance through combination with vision and sound. In addition, individuals can use the algorithm to create haptic icons specific to their partners, for more personalized affective expression.

5.3 Data-Driven Rendering

Our work shows the benefit of a data-driven approach to social communication on simple array hardware. These contributions could lead to cheap, lightweight, wearable consumer social touch devices, or integration of such devices into more complex systems with minimal extra cost and weight. This is in contrast to current consumer haptic devices such as successors to the PHANTOM haptic device [37] (3D Systems, Rock Hill, SC, US) and the HaptX glove [38] (Seattle, WA, US), which are heavy, have limited workspaces, and are expensive. The algorithm is not specific to our dataset or hardware, so it can be applied to other types of actuators, including microfluidic haptic arrays [39], pin arrays [40], and ultrasonic devices [41]. Our dataset is also freely available to be applied for further research.

5.4 Use of Scenarios as Social Touch Prompt

To generate social touch data, we chose to provide specific scenarios rather than discrete affective category labels for the following two reasons: First, we wished to maximize the probability that the subjects performed authentic touches. Genuine emotion has been shown to be distinguishable from non-genuine emotion for facial expressions [42]. Based on verbal feedback, informal pilot testing indicated that subjects were engaged by the scenarios. Second, we wished to minimize the impact of cultural perception, language ambiguity, or personal history in the terms we used. Different cultures express different levels of valence and arousal to the same emotional keywords [43]. In general, these terms also

have broad definitions, capable of encompassing a spectrum of affective states that could depend heavily on each user at a given time. Therefore, we posit that in our dataset the exact nature of what someone is trying to express is clearer than in datasets which assume that all participants have a similar interpretation of a single-word cue. For these reasons, we propose that the choice of scenarios over discrete affective category labels resulted in a less ambiguous task for those in the classification experiment. Further analysis on the impact of single-word cues compared to scenarios is an interesting direction for future work.

We acknowledge that such scenarios are limited in their ability to capture a complex affective state. As mentioned, one use of this work could be to create haptic icons similar to emojis, which are limited in their precise expressive capability, but understandable in context.

ACKNOWLEDGMENTS

Research reported in this publication was supported by Facebook, Inc. The authors acknowledge additional support for S.R.W. and C.M.N. from the National Science Foundation Graduate Research Fellowship Program. We thank the behavioral scientists at IDEO who contributed to study design.

REFERENCES

- [1] M. McPherson, L. Smith-Lovin, and M. E. Brashears, "Social isolation in america: Changes in core discussion networks over two decades," *American Sociological Review*, vol. 71, no. 3, pp. 353–375, 2006.
- [2] J. Holt-Lunstad, T. B. Smith, M. Baker, T. Harris, and D. Stephenson, "Loneliness and social isolation as risk factors for mortality: a meta-analytic review," *Perspectives on Psychological Science*, vol. 10, no. 2, pp. 227–237, 2015.
- [3] S. M. Kissler, C. Tedijanto, E. Goldstein, Y. H. Grad, and M. Lipsitch, "Projecting the transmission dynamics of sars-cov-2 through the postpandemic period," *Science*, vol. 368, no. 6493, pp. 860–868, 2020.
- [4] H. Seifi and K. E. Maclean, "A first look at individuals' affective ratings of vibrations," in *Proc. IEEE World Haptics Conference*, 2013, pp. 605–610.
- [5] G. Wilson, D. Dobrev, and S. A. Brewster, "Hot under the collar: mapping thermal feedback to dimensional models of emotion," in *Proc. CHI Conference on Human Factors in Computing Systems*. ACM, 2016, pp. 4838–4849.
- [6] M. Y. Tsalamalal, N. Ouarti, J.-C. Martin, and M. Ammi, "Haptic communication of dimensions of emotions using air jet based tactile stimulation," *Journal on Multimodal User Interfaces*, vol. 9, no. 1, pp. 69–77, 2015.
- [7] H. Nakanishi, K. Tanaka, and Y. Wada, "Remote handshaking: touch enhances video-mediated social telepresence," in *Proc. ACM Conf. on Human Factors in Computing Systems*, 2014, pp. 2143–2152.
- [8] F. Mueller, F. Vetere, M. R. Gibbs, J. Kjeldskov, S. Pedell, and S. Howard, "Hug over a distance," in *Proc. CHI Extended Abstracts on Human Factors in Computing Systems*, 2005, pp. 1673–1676.
- [9] D. Tsetserukou, "Hapthug: A novel haptic display for communication of hug over a distance," in *Proc. EuroHaptics Conference*, 2010, pp. 340–347.
- [10] E. Eichhorn, R. Wettach, and E. Hornecker, "A stroking device for spatially separated couples," in *Proc. ACM Int. Conf. on Human Computer Interaction with Mobile Devices and Services*, 2008, pp. 303–306.
- [11] H. Culbertson, C. M. Nunez, A. Israr, F. Lau, F. Abnoui, and A. M. Okamura, "A social haptic device to create continuous lateral motion using sequential normal indentation," in *IEEE Haptics Symposium*, 2018, pp. 32–39.
- [12] E. Knoop and J. Rossiter, "The tickler: a compliant wearable tactile display for stroking and tickling," in *Proc. ACM Conf. Extended Abstracts on Human Factors in Computing Systems*, 2015, pp. 1133–1138.
- [13] A. Israr and F. Abnoui, "Towards pleasant touch: vibrotactile grids for social touch interactions," in *Proc. ACM Conf. Ext. Abstracts on Human Factors in Computing Systems*, 2018, pp. 1–6.
- [14] S. C. Hauser, S. McIntyre, A. Israr, H. Olausson, and G. J. Gerling, "Uncovering human-to-human physical interactions that underlie emotional and affective touch communication," in *Proc. IEEE World Haptics Conference*, 2019, pp. 407–412.
- [15] M. J. Hertenstein, R. Holmes, M. McCullough, and D. Keltner, "The communication of emotion via touch," *Emotion*, vol. 9, no. 4, p. 566, 2009.
- [16] J. Rantala, K. Salminen, R. Raisamo, and V. Surakka, "Touch gestures in communicating emotional intention via vibrotactile stimulation," *International Journal of Human-Computer Studies*, vol. 71, no. 6, pp. 679–690, 2013.
- [17] G. Huisman, A. D. Frederiks, B. Van Dijk, D. Hevlen, and B. Kröse, "The tast: Tactile sleeve for social touch," in *2013 World Haptics Conference (WHC)*. IEEE, 2013, pp. 211–216.
- [18] C. M. Nunez, S. R. Williams, A. M. Okamura, and H. Culbertson, "Understanding continuous and pleasant linear sensations on the forearm from a sequential discrete lateral skin-slip haptic device," *IEEE Transactions on Haptics*, vol. 12, no. 4, pp. 414–427, 2019. [Online]. Available: <https://doi.org/10.1109/TOH.2019.2941190>
- [19] L. Zhang, Y. Li, and R. Nevatia, "Global data association for multi-object tracking using network flows," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 2008, pp. 1–8.
- [20] R. Ackerley, I. Carlsson, H. Wester, H. Olausson, and H. Backlund Wasling, "Touch perceptions across skin sites: differences between sensitivity, direction discrimination and pleasantness," *Frontiers in Behavioral Neuroscience*, vol. 8, p. 54, 2014.
- [21] M. M. Jung, M. Poel, R. Poppe, and D. K. Heylen, "Automatic recognition of touch gestures in the corpus of social touch," *Journal on multimodal user interfaces*, vol. 11, no. 1, pp. 81–96, 2017.
- [22] A. Flagg and K. MacLean, "Affective touch gesture recognition for a furry zoomorphic machine," in *Proc. ACM Int. Conf. on Tangible, Embedded and Embodied Interaction*, 2013, pp. 25–32.
- [23] M. M. Jung, X. L. Cang, M. Poel, and K. E. MacLean, "Touch challenge'15: Recognizing social touch gestures," in *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*, 2015, pp. 387–390.
- [24] J. T. Suvilehto, E. Glerean, R. I. Dunbar, R. Hari, and L. Nummenmaa, "Topography of social touching depends on emotional bonds between humans," *Proceedings of the National Academy of Sciences*, vol. 112, no. 45, pp. 13 811–13 816, 2015.
- [25] I. S. Statistics, "Tests for several related samples test types," IBM <https://www.ibm.com/docs/en/spss-statistics/24.0.0?topic=samples-tests-several-related-test-types>, 2016.
- [26] S. E. Jones and A. E. Yarbrough, "A naturalistic study of the meanings of touch," *Communications Monographs*, vol. 52, no. 1, pp. 19–56, 1985.
- [27] S. Guest, J. M. Dessirier, A. Mehrabyan, F. McGlone, G. Essick, G. Gescheider, A. Fontana, R. Xiong, R. Ackerley, and K. Blot, "The development and validation of sensory and emotional scales of touch perception," *Attention, Perception, & Psychophysics*, vol. 73, no. 2, pp. 531–550, 2011.
- [28] A. Bhattacharyya, "On a measure of divergence between two multinomial populations," *Sankhyā: The Indian Journal of Statistics*, vol. 7, no. 4, pp. 401–406, 1946.
- [29] M. M. Bradley and P. J. Lang, "Measuring emotion: the self-assessment manikin and the semantic differential," *Journal of Behavior Therapy and Experimental Psychiatry*, vol. 25, no. 1, pp. 49–59, 1994.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [31] J. A. Russell, "A circumplex model of affect," *Journal of Personality and Social Psychology*, vol. 39, no. 6, p. 1161, 1980.
- [32] T. Caliński and J. Harabasz, "A dendrite method for cluster analysis," *Communications in Statistics—Theory and Methods*, vol. 3, no. 1, pp. 1–27, 1974.
- [33] P. J. Rousseeuw, "Silhouettes: a graphical aid to the interpretation and validation of cluster analysis," *Journal of Computational and Applied Mathematics*, vol. 20, pp. 53–65, 1987.
- [34] L. v. d. Maaten and G. Hinton, "Visualizing data using t-SNE," *Journal of Machine Learning Research*, vol. 9, pp. 2579–2605, 2008.
- [35] E. H. Thompson and J. A. Hampton, "The effect of relationship status on communicating emotions through touch," *Cognition and Emotion*, vol. 25, no. 2, pp. 295–306, 2011.

- [36] V. Gazzola, M. L. Spezio, J. A. Etzel, F. Castelli, R. Adolphs, and C. Keysers, "Primary somatosensory cortex discriminates affective significance in social touch," *Proceedings of the National Academy of Sciences*, vol. 109, no. 25, pp. E1657–E1666, 2012.
- [37] T. H. Massie, J. K. Salisbury *et al.*, "The phantom haptic interface: A device for probing virtual objects," in *Proc. ASME Sym. on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 1994, pp. 295–300.
- [38] J. A. Rubin, R. S. Crockett, M. Y. Goupil, J. D. D'Amelio, B. L. Rojanachaichanin, K. C. Sjöberg, P. Piller, and N. J. Bonafede, "Haptic feedback glove," 2018, uS Patent US20180335842A1.
- [39] N. Besse, "Large array of shape memory polymer actuators for haptics and microfluidics," Ph.D. dissertation, École Polytechnique Fédérale de Lausanne, 2018.
- [40] A. F. Siu, E. J. Gonzalez, S. Yuan, J. B. Ginsberg, and S. Follmer, "shapeShift: 2D spatial manipulation and self-actuation of tabletop shape displays for tangible and haptic interaction," in *Proc. ACM CHI Conference on Human Factors in Computing Systems*, 2018.
- [41] T. Carter, S. A. Seah, B. Long, B. Drinkwater, and S. Subramanian, "Ultrahaptics: multi-point mid-air haptic feedback for touch surfaces," in *Proc. ACM Sym. on User Interface Software and Technology*, 2013, pp. 505–514.
- [42] P. Gosselin, G. Kirouac, and F. Y. Doré, "Components and recognition of facial expression in the communication of emotion by actors." *Journal of personality and social psychology*, vol. 68, no. 1, p. 83, 1995.
- [43] E. Bann and J. Bryson, "Measuring cultural relativity of emotional valence and arousal using semantic clustering and twitter," in *COGSCI 2013*, aug 2013, pp. 1809–1814.



Mike Salvato received their S.B. and M.Eng. in computer science from the Massachusetts Institute of Technology, where they conducted research with the Marine Robotics Group. They are currently a PhD candidate in mechanical engineering at Stanford University. Prior to their PhD they worked as a research engineer at the Allen Institute for Artificial Intelligence. Their research on haptics and computer vision aims to improve sensory feedback during interactions in virtual and augmented reality.



Sophia R. Williams received her B.S. from Harvey Mudd College in 2015 and her M.S. and Ph.D. from Stanford University in 2019 and 2021, respectively. Her research interests include wearable haptic interfaces, psychophysics, design, robotics, and human-robot interaction. Sophia is a recipient of the National Science Foundation Graduate Research Fellowship, Stanford Robotics Center fellowship sponsored by FANUC, and Watson Fellowship. She now works at Auris Health.



devices for medical applications, human-robot interaction, virtual reality, and STEM education.

Cara M. Nunez received a B.S. in biomedical engineering and a B.A. in Spanish from the University of Rhode Island in 2016, the M.S. degree in mechanical engineering from Stanford University in 2018, and a Ph.D. in bioengineering at Stanford University in 2021. She was a National Science Foundation Graduate Research Fellow at Stanford and a DAAD Graduate Research Fellow at the Max Planck Institute for Intelligent Systems. Her research interests include haptic perception, cutaneous feedback, and wearable



Xin Zhu received his BSc in Computer Science from Beijing Jiaotong University, Beijing, China, and MS in Computer Science from University of Southern California in 2019. He is currently a PhD student in HaRVI lab, USC. Xin's research interests include haptics technology with social touch, virtual reality and HCI.



Ali Israr received his B.Sc. in Mechanical Engineering from University of Engineering & Technology, Lahore, Pakistan, and M.S/ and Ph.D. in mechanical engineering from Purdue University in 2004 and 2007, respectively. He was at Rice University as postdoc and has since led haptic projects in Disney Research and Facebook, leading efforts to expand the utility of haptic enabled technologies for entertainment, educational, social, therapeutic and ARVR settings.



Frances W. Y. Lau received the B.A.Sc. degree in electrical engineering from the University of Toronto in 2005, and the M.S. degree and Ph.D. degrees in electrical engineering from Stanford University in 2007 and 2013, respectively. Since 2016, she has been working at Facebook on research and development of systems that act as interfaces between humans and technology, using methods that range from haptics to brain-computer interfaces.



Keith Klumb completed his BS at Purdue University and is currently a Technical/Research Program Manager. He has worked on research and new technology initiatives including BCI, Soli Radar, haptics psychophysics and HR monitoring at various companies including Google, Facebook and Fitbit.



Freddy Abnoui is the Head of Health Technology at Facebook and a practicing Interventional Cardiologist. He also serves as Innovation Advisor to the American College of Cardiology, Professor Adjunct at Stanford University School of Medicine, and Assistant Professor Adjunct at Yale University School of Medicine. He completed Fellowships in Cardiovascular Medicine and Interventional Cardiology, as well as Residency in Internal Medicine at Stanford University Medical Center. He was previously a resident surgeon at the University of California, San Francisco. He completed his MD at Stanford University School of Medicine, MBA from Oxford University, and MSc in Health Policy, Planning, & Financing from the London School of Economics.



Allison M. Okamura received the B.S. degree from the University of California, Berkeley, in 1994, and the M.S. and Ph.D. degrees from Stanford University, Stanford, CA, in 1996 and 2000, respectively, all in mechanical engineering. She is currently a Professor of mechanical engineering at Stanford University, Stanford, CA. Her research interests include haptics, teleoperation, medical robotics, virtual environments and simulation, neuromechanics and rehabilitation, prosthetics, and engineering education.



Heather Culbertson Heather Culbertson is an assistant professor of Computer Science and Aerospace and Mechanical Engineering at the University of Southern California. She received the M.S. and Ph.D. degrees in the department of mechanical engineering and applied mechanics (MEAM) at the University of Pennsylvania in 2013 and 2015, respectively. She earned the B.S. degree in mechanical engineering at the University of Nevada, Reno in 2010. Prior to joining USC she was a research scientist at Stanford University. Her research focuses on the design and control of haptic devices and rendering systems, human-robot interaction, and virtual reality.