Bioinspiration & Biomimetics

PAPER

Texture recognition and localization in amorphous robotic skin

To cite this article: Dana Hughes and Nikolaus Correll 2015 *Bioinspir. Biomim.* **10** 055002

Texture Recognition and Localization in Amorphous Robotic Skin

Dana Hughes and Nikolaus Correll
Department of Computer Science,
University of Colorado–Boulder,
UCB430, Boulder, CO 80309-0430
{dana.hughes, nikolaus.correll}@colorado.edu

June 8, 2015

Abstract

We present a soft robotic skin that can recognize and localize texture using a distributed set of sensors and computational elements that are inspired by the Pacinian corpuscle, the fast adapting, uniformly spaced mechanoreceptor with a wide receptive field, which is responsive to vibrations due to rubbing or slip on the skin. Tactile sensing and texture recognition is important for controlled manipulation of objects and navigating in one's environment. Yet, providing robotic systems or prosthetic devices with such capability at high density and bandwidth remains challenging. Each sensor node in the presented skin is created by collocating computational elements with individual microphones. These nodes are networked in a lattice and embedded in EcoFlex rubber, forming an amorphous medium. Unlike existing skins consisting of passive sensor arrays that feed into a central computer, our approach allows for detecting, conditioning and processing of tactile signals in-skin, facilitating the use of high-bandwidth signals, such as vibration, and allowing nodes to respond only to signals of interest. Communication between nodes allows the skin to localize the source of a stimulus, such as rubbing or slip, in a decentralized manner. Signal processing by individual nodes allows the skin to estimate the material touched. Combining these two capabilities, the skin is able to convert high-bandwidth, spatiotemporal information into low-bandwidth, event driven information. Unlike taxel-based sensing arrays, this amorphous approach greatly reduces the computational load on the central robotic system. We describe the design, analysis, construction, instrumentation and programming of the robotic skin. We demonstrate that a 2.8 square meter skin with 10 sensing nodes is capable of localizing stimulus to within 2 centimeters, and that an individual sensing node can identify 15 textures with an accuracy of 71%. Finally, we discuss how such a skin could be used for full-body sensing in existing robots, augment existing sensing modalities, and how this material may be useful in robotic grasping tasks.

Keywords: robotic skin, amorphous materials, tactile sensing, texture recognition

1 Introduction

In humans, the sense of touch is critical for a variety of tasks: grasping, manipulating and identifying objects, detecting collisions with the environment, and perception and control of the body [1, 2]. Developing a sense of touch in robotic devices has been explored for several decades. Within the last decade, tactile sensors have evolved from being located solely on a fingertip or hand to sensor arrays for full body sensing [3]. Robotic arms equipped with full-body tactile sensing become capable of navigating cluttered

environments and avoid damaging fragile objects [4]. This capability becomes even more important with direct human-robot interaction, such as with nursing robotic assistants or robotic companions [5]. Incorporating a sense of touch can improve the robustness of grasping tasks, such as when grasping with an end effector [6, 7] or with full-body manipulation of large objects [8]. In addition, autonomous mobile robots can utilize full-body tactile sensing in environments where vision may be limited by occlusion of obstacles, such as when exploring an object by manipulating it [9] or when navigating through bushes and trees or manipulating foliage during a foraging task. Here, the ability to not only sense touch, but also texture might add a whole new dimension of environmental awareness both during human-robot interaction and when navigating through or manipulating the environment.

Figure 1: A soft, amorphous texture-sensitive skin mounted on the back of a Baxter robot.

The move to full-body, multi-modal tactile sensing presents several engineering challenges which are not of concern with fingertip and hand sensors [10, 11]. Sensor arrays and networks suffer scalability issues as the number of sensors becomes large and their required bandwidth increases. Communication bandwidth and centralized processing of measured values are both bottlenecks in the system, limiting the number of sensors (and consequently, the sensor density) and individual sensor bandwidth. The ability to tessillage individual sensors into a large array places constraints on the shape of a sensor. Adhering sensors to complex robotic structures may result in gaps in the tessellation, severed communication channels and areas where sensors must be modified or omitted [12].

This paper presents a soft, autonomous sensing skin for localizing and identifying textures rubbed against the skin that can be manufactured in arbitrary shapes or sizes, only limited by the size and spacing of individual sensor nodes. We propose a design that considers the skin as an amorphous material capable

of processing stimuli within the material itself, rather than a matrix of densely packed individual sensors acting as taxels and communicating measurements directly to a central processor. The tight coupling between physical and computational properties has been explored recently in the broader context of robotic materials [13]. Collocating microcontrollers with sensors allows for local and distributed processing of sensor measurements. The skin only needs to communicate with external devices when an event of interest occurs (e.g., the skin rubs against an obstacle), reducing communication and processing bandwidth. This is particularly important when moving from binary tactile sensors to high-bandwidth sensors such as textures, which require information at bandwidths in the order of hundreds of Hertz. A preliminary version of this work has been presented in [14] and has been extended by error analysis as a function of spatial location of each sensing node.

The remainder of the paper is organized as follows. Section 2 provides information regarding the biological motivation for the approach presented, and related work in the robotics field. Section 3 describes the design and manufacture of the robotic skin and the distributed algorithm used to locate and identify textures. Section 4 describes vibration propagation within the skin, source localization and uncertainty analysis. Section 5 discusses in-network texture identification. Section 6 discusses the benefits to this approach, applications in robotics, and future work to improve upon the initial prototype.

2 Background

Work has been performed over the last several decades to provide human-like tactile sensing capabilities in robotic skin. Tactile sensitive skins in robotics are typically designed to mimic sensing modalities found in humans. In addition, there exists several engineering and application challenges once such a skin is designed. Pressure sensing, which most resembles the behavior of Merkel's discs and Ruffini corpuscles, has seen a majority of the research in robotic skins over the last several decades. Pressure can be detected by measuring changes in resistance [15, 16, 17], capacitance [18, 19], or optical properties [20, 21] of the material under pressure, or by measuring changes in magnetic or electric fields [22, 23]. Detecting vibrations, which mimics the role of Pacinian corpuscles, have been performed using accelerometers [7] and simple microphones [24, 25].

Early work into texture detection includes the development of a fingertip for identifying texture [26]. This fingertip combined slip detection, temperature, pressure sensors and vision to train a neural network to distinguish between 20 different textures with almost 100% accuracy. More recently, a robotic fingertip was developed which was capable of identifying textures from a database of 117 textures with an accuracy of 95.4% using a Bayesian classifier [27]. As comprehensively surveyed in [3], the last decade has seen tactile sensing evolve from sensors for fingertips and hands to sensor arrays suitable for full-body tactile sensing, particularly for force and pressure sensing [28, 29, 30, 31].

Robust texture recognition has several applications in robotics. Navigation using only tactile information of the terrain has been explored [32], where tactile information was used to orient a wheeled robot parallel to the boundary between a tiled and carpeted terrain. In humanoid robots, knowledge of contact forces in a robotic arm is useful when reaching into and navigating unknown, cluttered environments [4]. In such cases, even categorizing the compliance (rigid vs. hard) and mobility (movable vs. fixed) of objects the arm makes contact allows for efficient searching and mapping cluttered volumes [33]. In the domain of assistive robotics, such knowledge is critical to ensure safe interaction between humans and robots [5].

Understanding the shape and material properties of objects is also very beneficial in grasping. In humans, tactile feedback is used to control grasping pressure [1]. This approach has been mimicked for grasping with parallel-jaw grippers [7]. Tactile sensing also enables learning stable grasps for unknown objects, reducing the need for object models [34], or using the entire body of the robot for grasping large objects [35]. Finally, tactile sensing is an important capability in the emerging domain of soft robotics [36], possibly allowing high-degree of freedom soft manipulator to autonomously explore and conform to an object's surface properties and shape.

A central issue with full-body tactile sensing is transmitting the signals from the sensors to a communication sink for further processing. Arranging the sensors in a matrix is one solution, and is the approach used by [29] and [30], among others, for a capacitive pressure sensing skin. An alternative is to organize the sensors in a hierarchical bus. Using this approach, [37] constructed "cut-and-paste" tactile sensing sheets whose 1,024 sensors in groups of 32 can be connected to an SMBus. Also, [18] demonstrated a system with 192 capacitive pressure sensors that feed into a hierarchical bus with increasing bandwidth (I^2C to CAN) and can detect binary touch events at 50 Hz. This system has evolved to the ROBOSKIN project [31]. As the design maxed out the bandwidth of the CAN bus, any additional sensor sharing this communication channel would drastically reduce the sensor bandwidth. Finally, in [38] capacitive skin patches have been organized into slave nodes connected to a master node using an EtherCAT-based communication bus. This distributed architecture allows for a large pressure sensitive skin to be implemented on humanoid robots while ensuring timing constraints for real-time control using tactile feedback is satisfied. A distributed tactile architecture [35] provides a self-organizing, full body tactile feedback system which allows tactile stimulation to be transformed into reactive motion.

The focus of this paper is not to replicate the results of previous work described above, but to investigate an amorphous architecture and distributed algorithms that can integrate such high-bandwidth, multi-modal sensors into a stretchable skin in a scalable and robust fashion. We focus on the Pacinian corpuscles as this sense has the highest bandwidth requirement. The full-body arrays described in [3] have focused primarily on the development of the sensors themselves and transduction of the signal. There is still much need for suitable conditioning and processing of the signal within the sensor network, prior to passing the signal to a central processor [3]. Finally, while the works on distributed tactile architectures have improved on the communication requirements of large sensing skins, the works have not explored the capability of significant in-network processing of information, such as texture localization and recognition, as presented here.

3 Skin Design

Vibration sensitive skin is of particular interest in this paper. There currently exists a large body of literature for pressure sensitive skins, developing skins capable of detecting and processing vibrations can provide an additional sensing modality for robotic applications. Detecting vibrations provides unique engineering challenges, such as high-bandwidth signals, and opportunities, such as non-local detection of stimuli.

Glabrous skin, the hairless region of human skin, contains four different mechanoreceptors used for tactile perception, as shown in Figure 2 [39, 40, 41]. Individually, these receptors provide sensitivity to light touch and low frequency vibrations (Meissner's corpuscles), skin deformation and static force (Merkel's discs), skin stretching and tangential shear (Ruffini corpuscles) and high-frequency vibrations

(Pacinian corpuscles). This investigation is specifically interested in Pacinian corpuscles. These receptors have been shown to be the primary means of perception of various textures [42], and an equivalent tactile sensor for robotic skin can augment pressure sensitive skins. In addition, Pacinian corpuscles have a very wide receptive field, and are capable of responding to stimuli occuring several centimeters away from the receptor [40]. This motivates using a sparse network of sensors, rather than a dense array, for this type of perception.

By collocating computing elements with individual sensors, the skin itself may be programmed to process the signal and determine when a signal should be reported to an external device for further processing. Individual sensor nodes or local neighborhoods of sensors may also perform in-network processing of the signal, possibly reducing several data points to a single, predefined event. Combining distributed sensing and computation with known material properties shifts this approach of skin design from one of a sensor array to an amorphous material. The material may be considered amorphous, both in a physical and computational sense, in that there is no requirements on the final shape of the skin nor the location of the sensing nodes. The skin may be cut to a required shape after manufacturing, so long as the underlying network is not divided into two distinct region. Furthermore, while the sensing nodes are equally spaced in the prototype, the localization and texture classification algorithms work independently of node spacing, and only require the position of the node within the skin.

For this investigation, we constructed a prototype skin combining texture sensing and localized computation. The purpose of this prototype is to explore texture detection and identification, networking issues and stimulus localization. The skin prototype consists of a network of ten sensor nodes embedded in silicone rubber (Figure 1). The task of texture recognition demonstrates the ability of this type of skin to solve the computation and communication bottlenecks associated with sensor arrays, while stimulus localization demonstrates how material selection and skin design can be leveraged.

3.1 Design and Manufacturing

Individual sensor nodes, shown in Figure 3, left, are first placed on a flexible neoprene rubber mesh (McMaster) and then embedded into silicone rubber (Ecoflex Supersoft 0030). The communication bus wires are woven in the rubber mesh. The wires are woven in a spiraling pattern, providing strain relief for the wires and ensuring the resulting skin remains stretchable. When connected to the sensor nodes, the wires securely attaches the sensor nodes to the mesh. Connecting the nodes and communication wires to a rubber mesh before embedding in silicone is necessary, as these components would otherwise eventually tear out of the silicone during use of the skin. Figure 3, middle, shows the sensor network and mesh before embedding in silicone rubber. The sensor nodes are spaced 15 cm apart. Based on the resonant frequency of the Pacinian corpuscle, 250Hz is considered the mid-frequency of interest to be measured by the microphones. In the silicone rubber, 15 cm corresponds to a distance half of a wavelength of sound at 250 Hz. The overall size of the mesh is approximately 61 cm x 43 cm. For comparison,the spatial acuity of vibration on the human torso (and hence the expected spacing of Pacinian corpuscles), based on two-point discrimination experiments, is 2 to 3 cm [44].

Ecoflex is a two-part liquid rubber which cures solid. The sensor network and mesh are placed in a form, the bottom of which is covered with 60-grit aluminum oxide sandpaper to create a surface texture similar to a fingertip. The male human fingertip contains an average of 22.4 ridges per cm, or a distance of

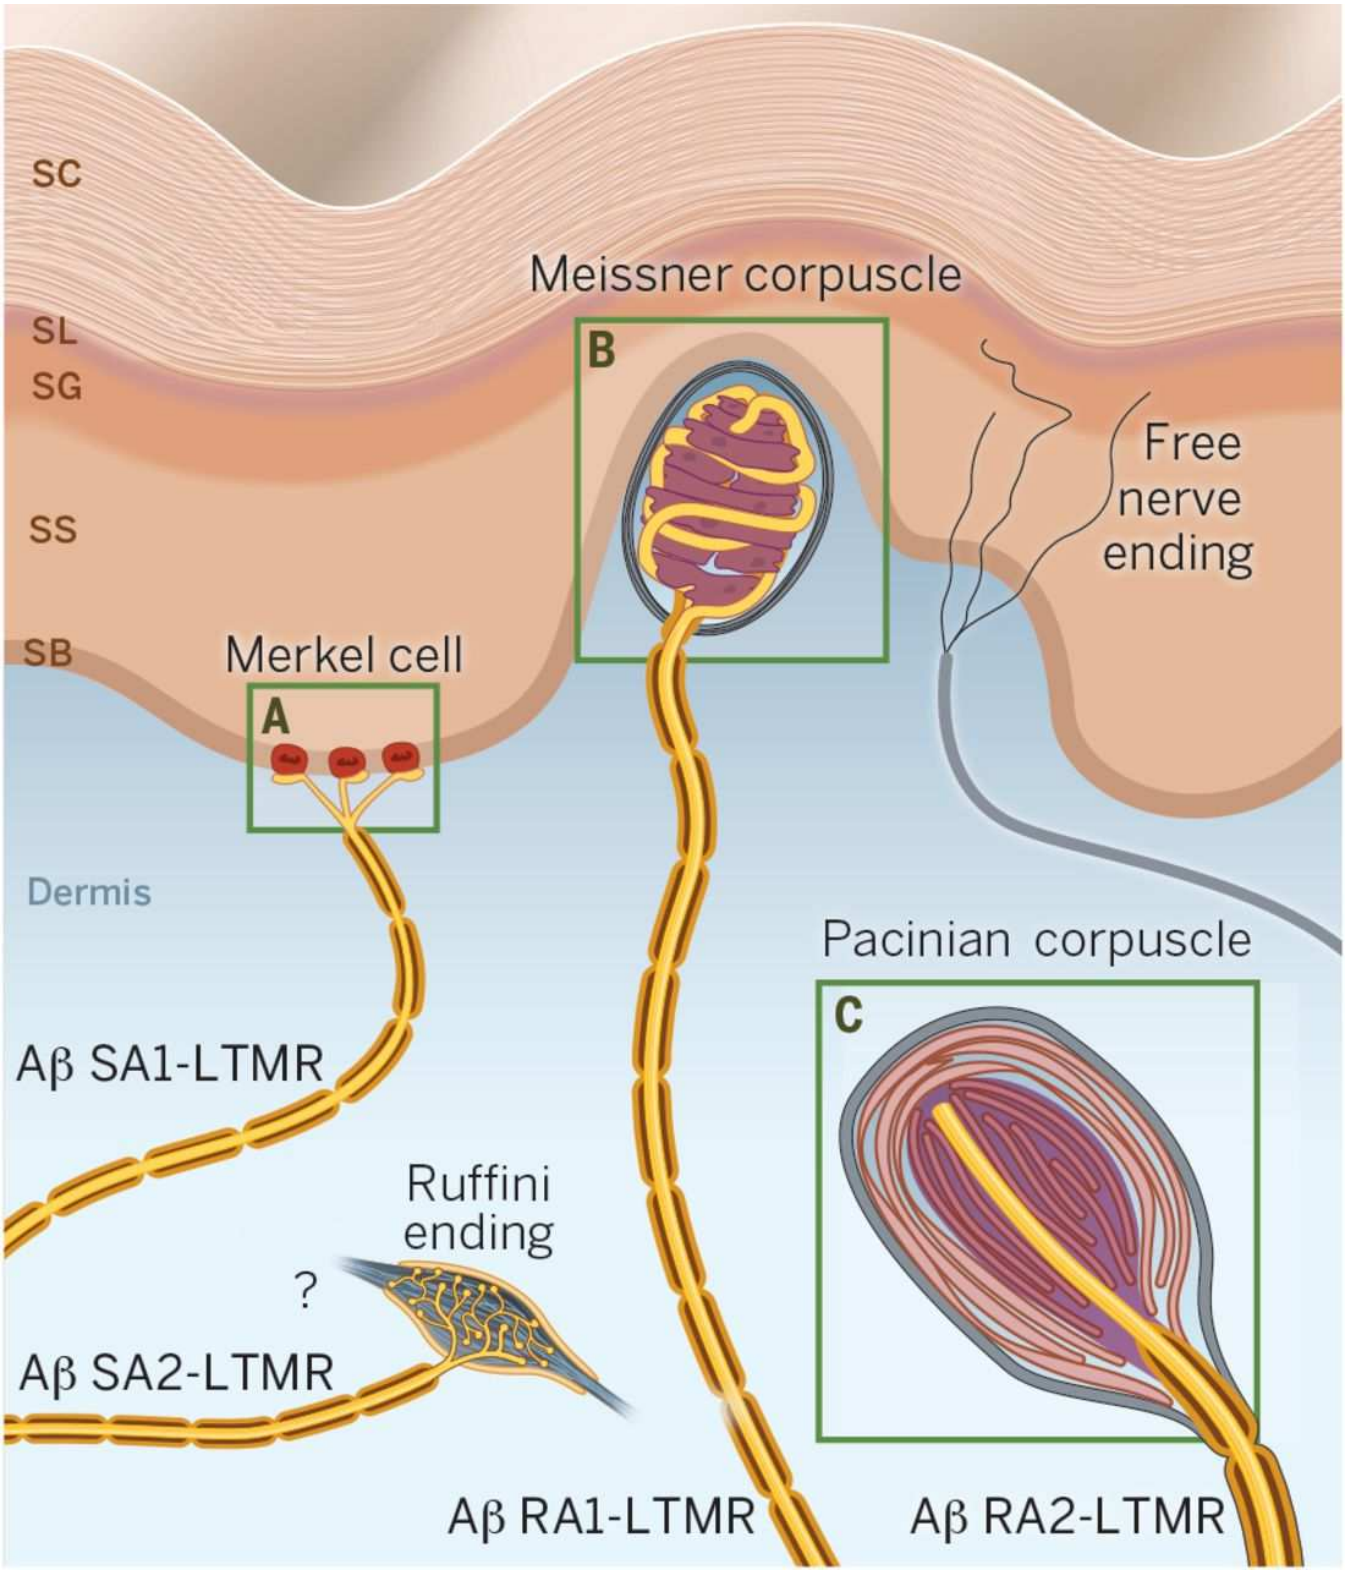

Figure 2: Relative location of mechanoreceptors in human skin. Image ©AAAS [43]

0.45 mm between ridges [45]. The grit size of 60-grit sandpaper is 0.25 mm [46], which roughly corresponds to the groove width of a human fingertip. Figure 3, right, shows the surface of the prototype skin after the sensor network is embedded in silicone. The resulting skin is approximately 1 cm in thickness. For

comparison, the thickness of human skin, measures at the forearm, ranges from 0.82 mm and 1.19 mm [47]. Mechanically, Ecoflex has a Young's modulus of 125 kPa [48], while human skin has a higher modulus of 420-850 kPa, depending on age [49]. Thickness and mechanical properties of the skin are dependent on the size of sensing components and available silicone rubber, respectively.

3.2 Sensor Node Network

Each sensor node is composed of an Atmel ATxmega128A3U microcontroller attached to a single WM-64K microphone, which serves as a vibration sensor. The WM-64K microphone is an omnidirectional microphone with a sensitivity of -45dB, and a signal-to-noise ratio of 58dB. The microphone is attached to the microcontroller's 12-bit analog-digital converter (ADC) through an operational amplifier (LM358). Each microcontroller can communicate with six neighboring nodes in the network using six hardware serial ports (USARTs) at 115kbps. The 115kbps data rate is the fastest serial communication rate available on the microcontroller. Two wire interfaces, such as I^2C, operate at a similar rate (100kHz or 400kHz), but may only address a single device at a given time, whereas the six USART ports can communicate in parallel. Using these six independent communication channels allows arranging the sensor nodes in a hexagonal lattice packing [50]. Communication channels consist of a four-wire bus consisting of power, ground, transmit and receive lines. A seventh serial port is optionally available on each node, and is only used for interfacing to a computer for initial programming and retrieving data at the sink node. Nodes can propagate a program in a viral fashion throughout the network. The absence of any central component allows for including or removing nodes at any time. Each node can calculate a unique ID based on manufacturing information (i.e., tray number, chip row and column, etc.) written into memory of the microcontroller during production.

Routine tasks are performed in each sensing node through the use of regularly scheduled interrupts. Communication is performed asynchronously through the use of direct memory access (DMA) channels. The DMA channels allows packets to be written into or read from predefined buffers in the microcontroller's memory independent of the central processing unit, and ensures that the main program does not need to pause measurements or computation to process communication packets. A clocked interrupt occuring every 10 μs processes received packets and transmits pending packets. Microphone samples are recorded into a 256-sample circular buffer at a rate of 1 kHz. When the buffer is full, the spectral energy of the signal is calculated using the Fast Fourier Transform (FFT), resulting in 128 spectral bins. Each spectral bin consists of the energy in a band of frequencies in the original signal. The FFT is a computationally efficient means to compute the discrete Fourier transform (DFT) of the signal. For a 256 sample signal, the FFT computes the energy spectrum 8 times faster than the DFT. The sensing nodes require 15.9 ms to calcualte the FFT of the sample buffer. Values for the ambient and transient spectral energy are updated from the calculated spectral energy. Each node keeps track of information from neighboring nodes with a table consisting of the neighboring node ID, position in the skin and the overall energy of the transient signal received at that node.

For the purpose of this work, information is propagated using a simple flooding algorithm. That is, whenever a sensor node detects a texture, results from classification are flooded through the network and can be collected anywhere. In order to limit communication to the immediate neighborhood during texture localization, we have implemented a Bloom-filter based multicast routing algorithm [51, 52].

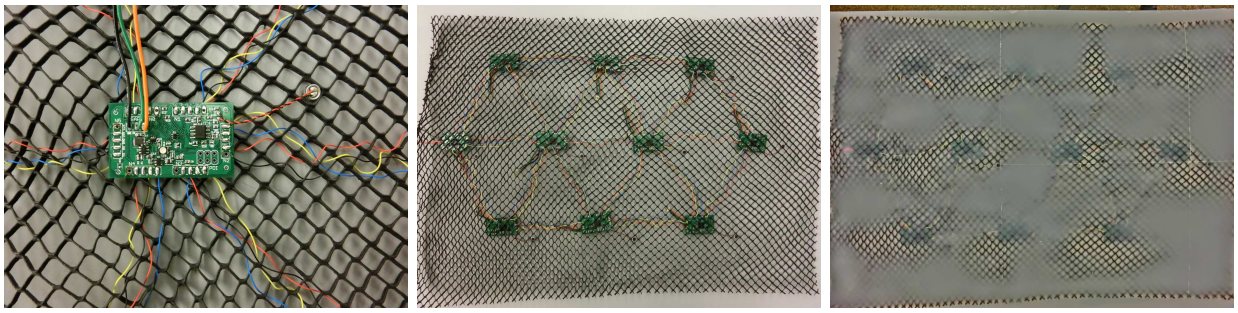

Figure 3: Left: Close-up on an individual sensor node. Middle: sensor network woven into a neoprene lattice. Right: Sensor node network embedded into EcoFlex^{TM} rubber. Finished skin with 60-grit surface texture.

To compare the reliability and robustness of the proposed network topology to that of a bus-like architecture, we conducted two sets of experiments. First, we sent 64 byte packets from node 2 to node 3 in Figure 4 (three hops) using a simple flooding algorithm with all adjacent nodes enabled. Figure 5 shows the percentage of received packets as a function of packet loss at each intermediate node for transport through the amorphous network (solid line) as a function of individual node failure rate (e.g., dropped packet, failed bus connection, etc.). In a second experiment, we removed all nodes labeled X from the network and sent packets from the node 1 to node 4 (6 hops) and measured throughput at each node for different packet loss (dashed lines). Results show that the amorphous network provides higher throughput than communication over one hop with nodes deliberately dropping 30% of the packets, and over two hops with nodes dropping 55% of the packets.

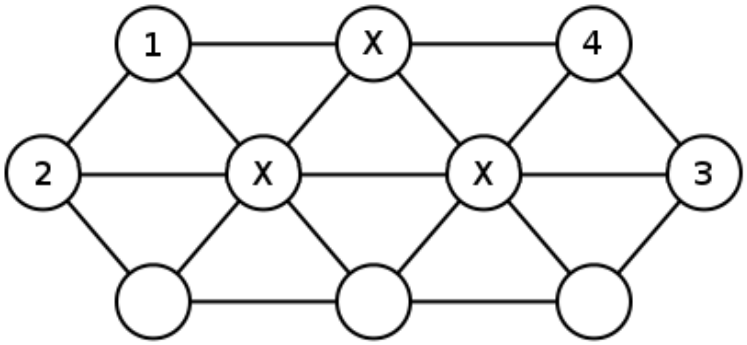

Figure 4: Source and sink nodes used in networking experiment

Individual sensing nodes are designed to operate in an event-based manner, such as detecting a transient signal or receiving packets from other nodes. To accommodate this, sensing nodes are treated as state machines with five distinct states, as shown in Figure 6. Nodes respond to stimuli and received packets differently depending on the current state. The behavior of nodes in each state, including which states nodes will transition to, are described below.

When nodes are initially started, they enter a CALIBRATION state. Calibration is used to verify the DC offset of the ADC, and estimate the initial level of the ambient signal. The node fills the microphone sample buffer and performs an FFT on the signal. If the DC offset is outside a predefined threshold, the

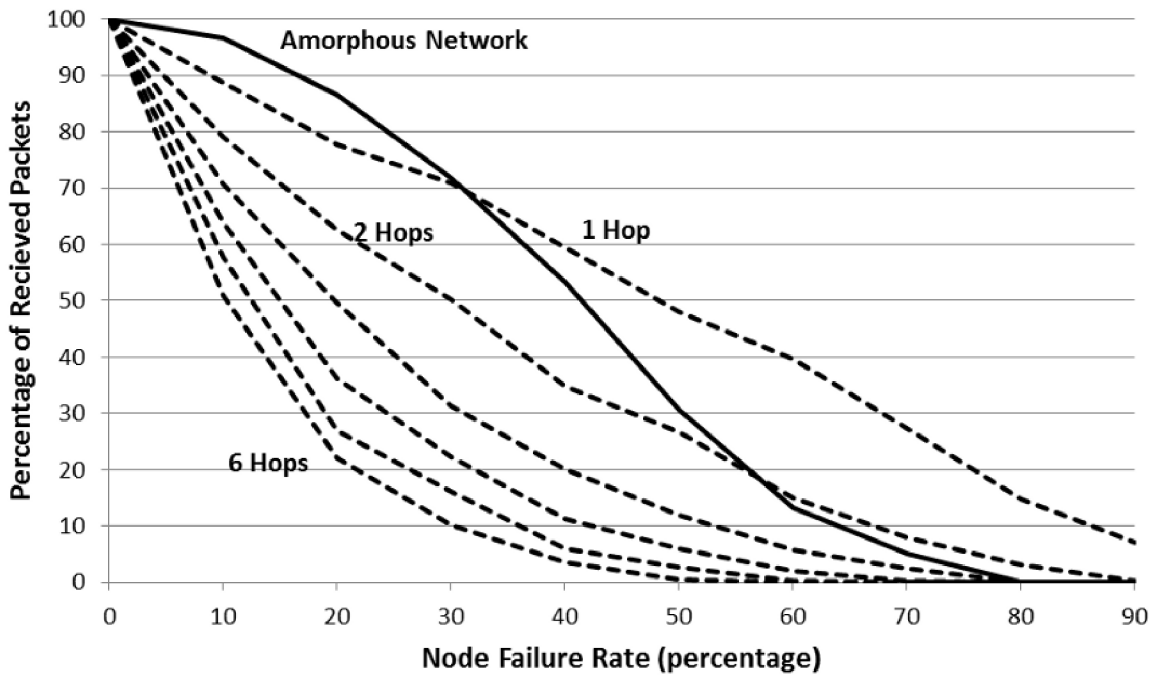

Figure 5: Packet loss in a network with unreliable nodes for different networking architectures

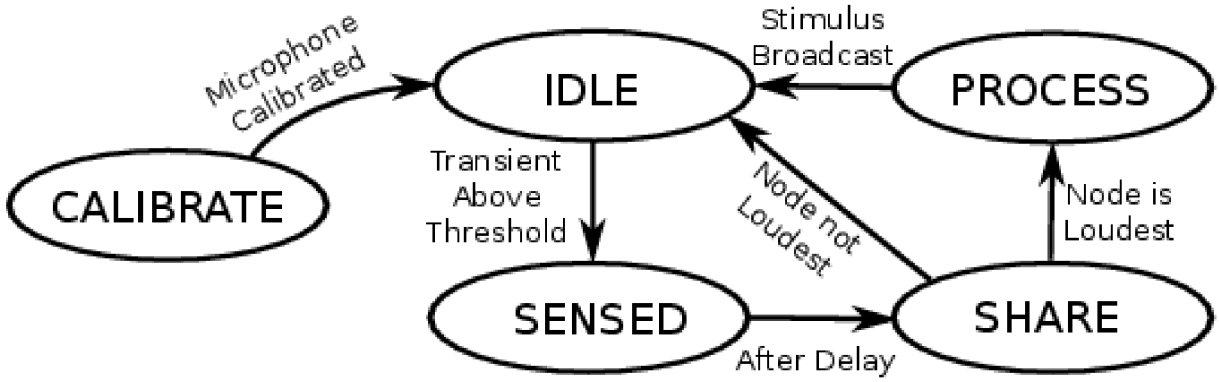

Figure 6: State machine of individual sensor node

measurement is discarded. This is repeated until a predefined number of valid measurements are obtained, the average of these are used as the initial estimate for the ambient signal. Once the valid measurements have been made, the node transitions into the IDLE state. While in the CALIBRATION state, nodes simply ignore received packets, except to forward received packets to neighboring nodes.

While in the IDLE state, a node will sample the microphone, periodically calculating the energy spectrum of the measured signal and updating the local ambient and transient energy spectra. Invidual microphone samples are stored in the microphone sample buffer. Once full, the buffer represents a single measurement window at time t. The energy spectrum of this time window $S(f,t)$, is calculated using the FFT. From this, the ambient energy spectrum can be updated by slow averaging the incoming measurement window with the existing ambient energy spectrum using the following equation

$$A(f,t) = \alpha S(f,t) + (1-\alpha)A(f,t-1) \tag{1}$$

where $A(f,t)$ is the estimated ambient noise at time slice t, and α is a factor representing the rate at which the ambient signal is updated. High values of α imply that the ambient estimation is sensitive to recent events, while lower values result in an ambient estimation which responds slower to changes in background noise. When the background noise is steady, there should be little change in $A(f,t)$ over time. For this investigation, a value of 0.05 was used as a smoothing constants, as suggested in [53].

The transient spectrum, $T(f,t)$ is calculated as the difference between the spectrum of the measured signal and the ambient noise, or

$$T(f,t) = min(S(f,t) - A(f,t), 0) \tag{2}$$

If the total energy in the transient signal, which is simply the sum over all frequency bins in the transient spectrum, exceeds the current ambient level by a threshold, the node transitions to the SENSED state. While in the IDLE state, nodes ignore received packets, simply forwarding to neighboring nodes.

The purpose of the SENSED state is to ensure that vibrations have propagated through the skin and has been recorded by neighboring nodes. The node records its measured transient energy value into its local table. The node then waits a random amount of time between 5 and 25 milliseconds, with 1 millisecond spacing. A vibration stimulation will be detected by multiple nodes on the skin. The random delay ensures that neighboring nodes have time to process a stimulus locally, and helps to balance communication load by ensuring that a group of nodes do not flood the communication channels simultaneously. If the timer expires, or if the node receives a packet from a neighboring node containing the neighbor's transient energy value, it transitions to the SHARE state.

In the SHARE state, the node waits 50 milliseconds to ensure that it receives information about a vibration from all neighboring nodes. When entering this state, the node broadcasts a packet containing its ID, location and transient energy value. During this phase, any received packets containing neighboring transient energy values are added to the node's local table. After the 50 millisecond time limit, the node checks its table to determine if its transient energy value is higher than all neighboring node values. If the node has the loudest transient energy value in its table, it transitions to the PROCESS state. Otherwise, the local table is cleared and it transitions back to the IDLE state.

A node which has transitioned into the PROCESS state has received the most energy from a detected source (i.e., the node is closest to the source). Theoretically, only one node should be in the PROCESS state for a given detection. In the PROCESS state, the node first attempts to determine the position of the source using the transient energies and locations in its local table. The node which performs processing for a given stimulus performs both localization of the stimulus and classification of the texture. Position is estimated first, and then the node classifies the texture of the material which caused the stimulation. Localization is described in detail in Section 4 and texture identification is described in detail in Section 5. Finally, the node broadcasts a packet containing the position of the source of the stimulation, and the classification of the texture. The node then transitions into the IDLE state.

4 Localization of Stimulus

Localization of a stimulus can either be performed by determining which sensor detected the most intense signal, as is the case in a densely packed sensor array, as in [18, 29, 54] or by utilizing the signals detected from a collection of sparsely located nodes and knowledge of how a signal propagates through the skin's material. Pressure detection in robotic skin is often performed using a sensor array, which mirror the small, well-defined receptive fields of Meissner corpuscles in human skin. Pacinian corpuscles exist in lower density in the skin, and have much broader receptive fields. As such, we developed a sparse sensing approach to localizing vibrations on the skin. The approach used in this paper is similar to sound source localization in [55], albeit sound propagation in this investigation involves a different propagation medium and structure.

Vibrations propagate through the skin according to the following equation

$$I(r,\omega) = \frac{I_0(F,\omega)}{r} \tag{3}$$

where I is the intensity of the vibration measured at a distance r from the source of the vibration, ω is the frequency of the vibration, and F is the displacement force of the vibration. The intensity of the vibration at the source, I_0, is dependent on the displacement force of the vibration and mechanical properties of the skin (i.e., thickness, density and rigidity), which are assumed to be constant. Details on the derivation of this equation are given in Appendix A. Performing calculations using a single frequency (as may be done for localization), the source intensity I_0 may be considered a constant.

A sound of sufficient intensity will be detected by various sensors throughout the skin. Figure 7 shows three nodes located at (x_1, y_1), (x_2, y_2) and (x_3, y_3), which measure a transient vibration signal with intensities I_1, I_2 and I_3, respectively. The position of the source of the vibration, (x, y) can be determined by minimizing the following equation with respect to the source position

$$\begin{aligned} x, y = \arg\min_{x,y} \{ & I_1^2((x-x_1)^2 + (y-y_1)^2) - I_2^2((x-x_2)^2 + (y-y_2)^2) + \\ & I_1^2((x-x_1)^2 + (y-y_1)^2) - I_3^2((x-x_3)^2 + (y-y_3)^2)\} \end{aligned} \tag{4}$$

Details on the derivation of this equation is given in Appendix B. Gradient descent is used to estimate the position of the source, using the mean of the locations of each node involved in the measurement as an initial guess. Based on the number of calculations performed, a single iteration of gradient descent requires 5.8 μs. Determining the position of the source may require up to 100 iteration for a very poor initial guess, which we set as the upper limit on the number of iterations to perform for localizaton. Thus, localization requires approximately 0.6 ms to calculate in the worse case.

4.1 Uncertainty Analysis

Equation 4 is convex with respect to the position of the stimulation source (x, y). Uncertainty in measured values, however, will result in variation of the estimate of the position of the vibration source. Treating measurements as samples from Gaussian random variables, error propagation can be used to determine the uncertainty in the estimated source position [56]. Details on determining the uncertainty of source position from noisy measurements are given in Appendix C.

To determine the effect of noisy measurements on the prediction of source location, uncertainty analysis

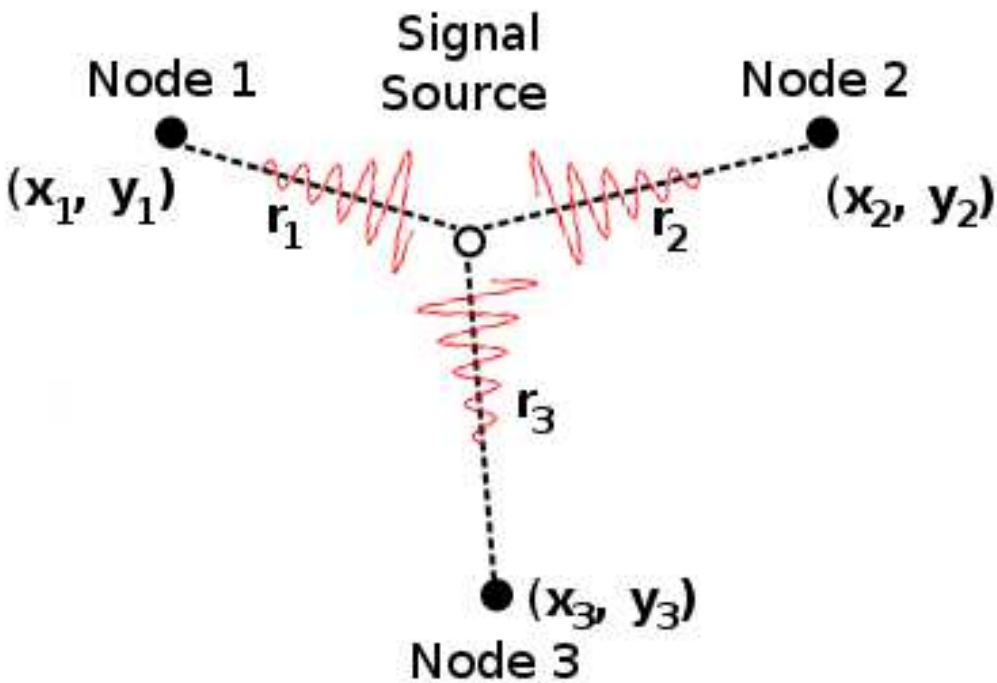

Figure 7: Propagation of sound to various sensors from an arbitrary point on the skin.

was performed for various conditions. Specifically, source position uncertainty is estimated as a function of sound intensity, sensor node spacing, and source position. The combination of these provide insight into determining node spacing given a desired skin acuity and microphone sensitivity.

4.1.1 Source Intensity

The uncertainty of the estimate of the source position with respect to the intensity of the source is given in Figure 8, left. The sensing nodes for this were spaced at 15 cm. This plot show a power-law relationship between source intensity and source position uncertainty, with uncertainty in source location decreasing as the intensity of the source increases. This is not surprising, given that the sound intensity decreases inversely as a function of distance, r, and the presence of a $r^{3/2}$ term in the Jacobian in Equation 20. This prediction also provides a lower bound on the expected spatial acuity–the lowest source intensity detectable by the microphone will give the highest location uncertainty expected in the skin.

4.1.2 Sensor Spacing

Figure 8, middle, shows the uncertainty of the source position estimate as a function of sensor spacing, for sensor node spacings ranging from 1 cm to 20 cm. As nodes are spaced further apart, the relative uncertainty in position appears to be bounded by a linear relationship with the node spacing. This relationship could be used to determine the node spacing necessary for a desired tactile acuity.

4.1.3 Source Location

To determine the effect on the relative position of the source with respect to the sensor nodes, the uncertainty in source position estimate was calculated for the source at various positions, with sensing nodes spaced 15 cm apart. Figure 8, right, gives the uncertainty of the position as a function of the location. The position of the source moves from the midpoint between two sensor nodes to a third sensor node, as shown in the inset of Figure 8, right.

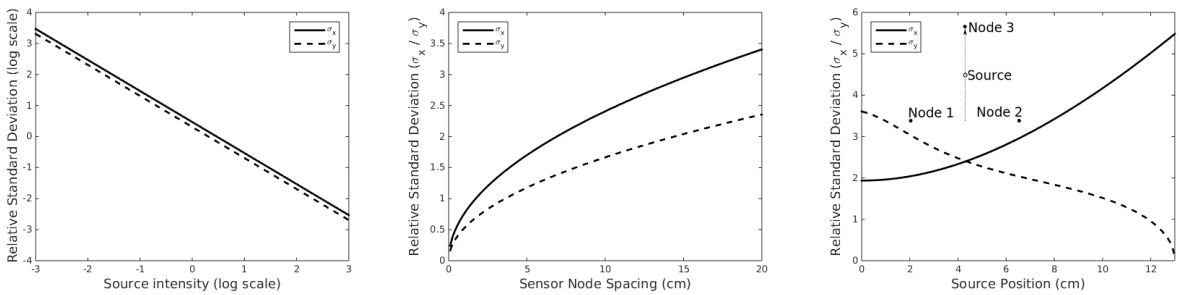

Figure 8: Left: Uncertainty of position as a function of source intensity. Middle: Uncertainty of position as a function of sensor spacing. Right: Uncertainty of position as a function of source location

4.2 Experimental Results

Two experiments were performed to verify the approach for localizing a vibrating stimulus on the skin. These experiments consisted of pressing a 2 mm thick vibration motor with a 1 cm diameter at various locations on the skin. The motor was placed so that the vibration footprint was a 1 cm diameter circle. The motor vibrated with a centripetal force of 0.887 N at a frequency of 150 Hz. The size of the motor is assumed to be much smaller than an area stimulated by rubbing, ensuring that the results from the motor experiment relate well to localizing texture signal sources in practice.

4.2.1 Sound Propagation

To validate sound propagation described in Equation 3, we pressed the vibration motor at specific distances from a microphone sensor against the skin. Figure 9 shows the intensity of the measured signal with respect to the distance between the vibration motor and microphone. At each point, 15 measurements were made. The smooth curve is a least-square fit to the expression I_0/r, where I_0 represents the unknown source intensity in Equation 3, and is the term varied to fit the measurement points. This figure demonstrates that Equation 3 is an accurate enough representation of the propagation of sound in the skin.

4.2.2 Localization

Initial localization experiments involved a 15 cm x 13 cm region of the skin shown in Figure 10, measuring the signal at three sensor nodes (Node 1, Node 2 and Node 3). The intensity of the vibration motor placed on a grid with 1 cm intervals were measured by each of these sensors. Figure 11 show the measured intensity of the transient signal measured by three sensor nodes located in the region. In all of these plots, the red areas (upper left and right corners, and lower center) indicate higher transient signal intensity, and darker blue indicates low signal intensity.

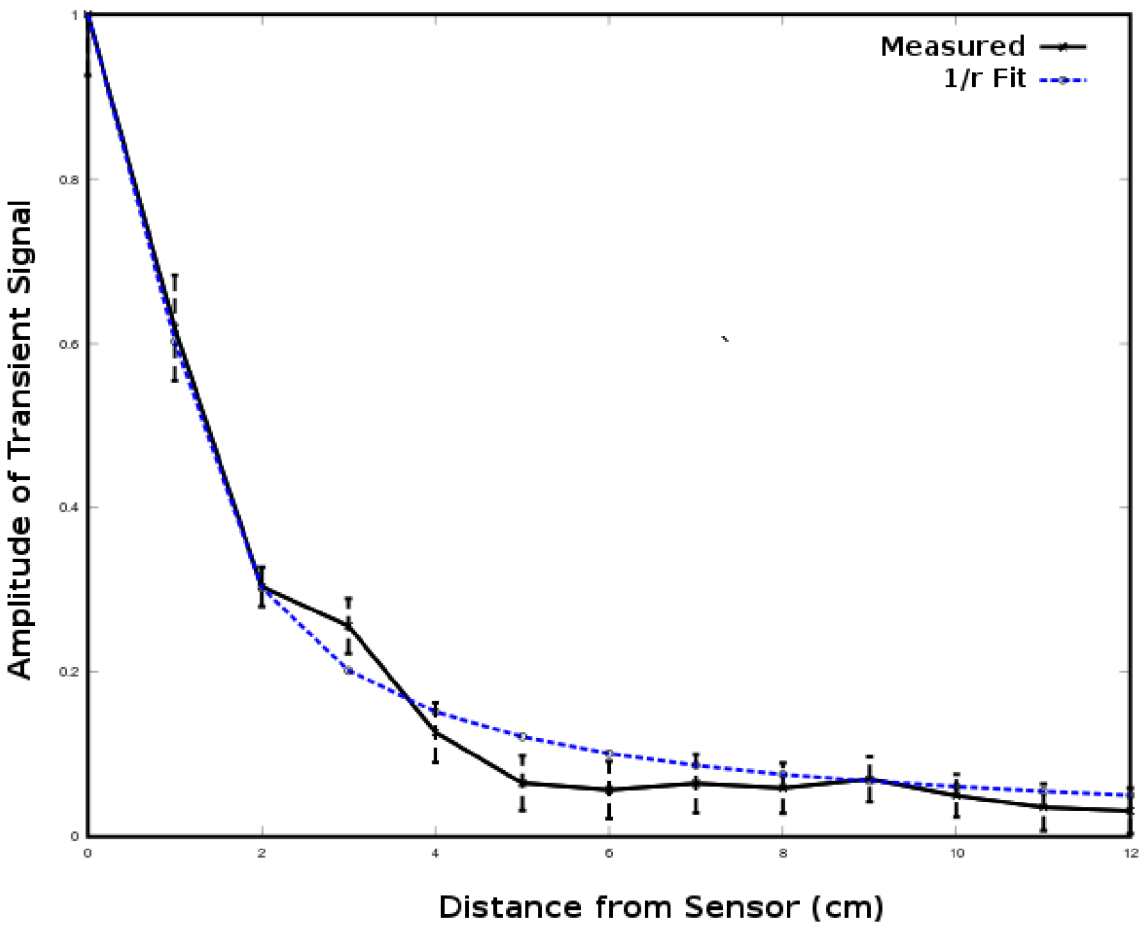

Figure 9: Amplitude of transient signal vs. Distance

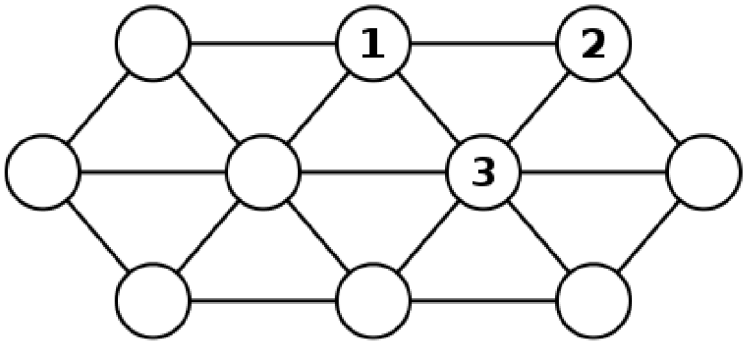

Figure 10: Region and sensor nodes used in localization experiment.

From these measurements, the location of the stimulus at each point was calculated as described in Section B. The error between the calculated and actual locations for each point is shown as a vector in Figure 12. The dashed lines define a triangle whose vertices are the three sensors which detected the signal. Outside this region, a different set of sensors would be used to determine the location of the source

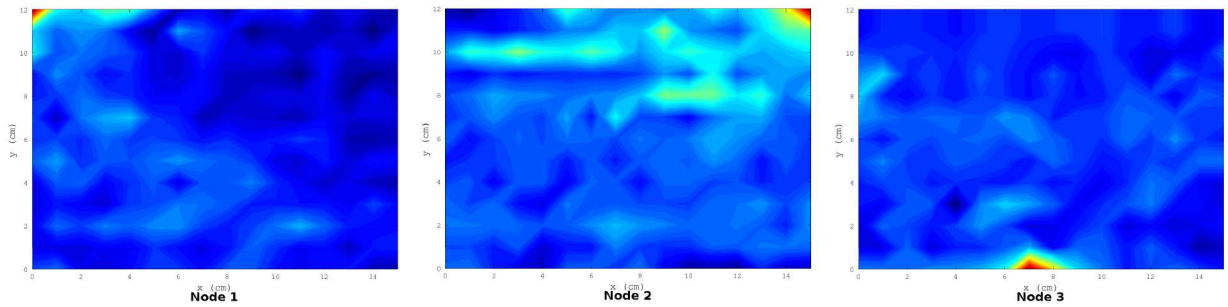

Figure 11: Normalized amplitude of transient signal from sensors nodes 1 (left), 2 (middle) and 3 (right).

and produce a more accurate result. Within the region of interest, the mean and standard deviation of the overall error was calculated as $\mu_r = 3.55$ cm and $\sigma_r = 1.96$ cm, respectively.

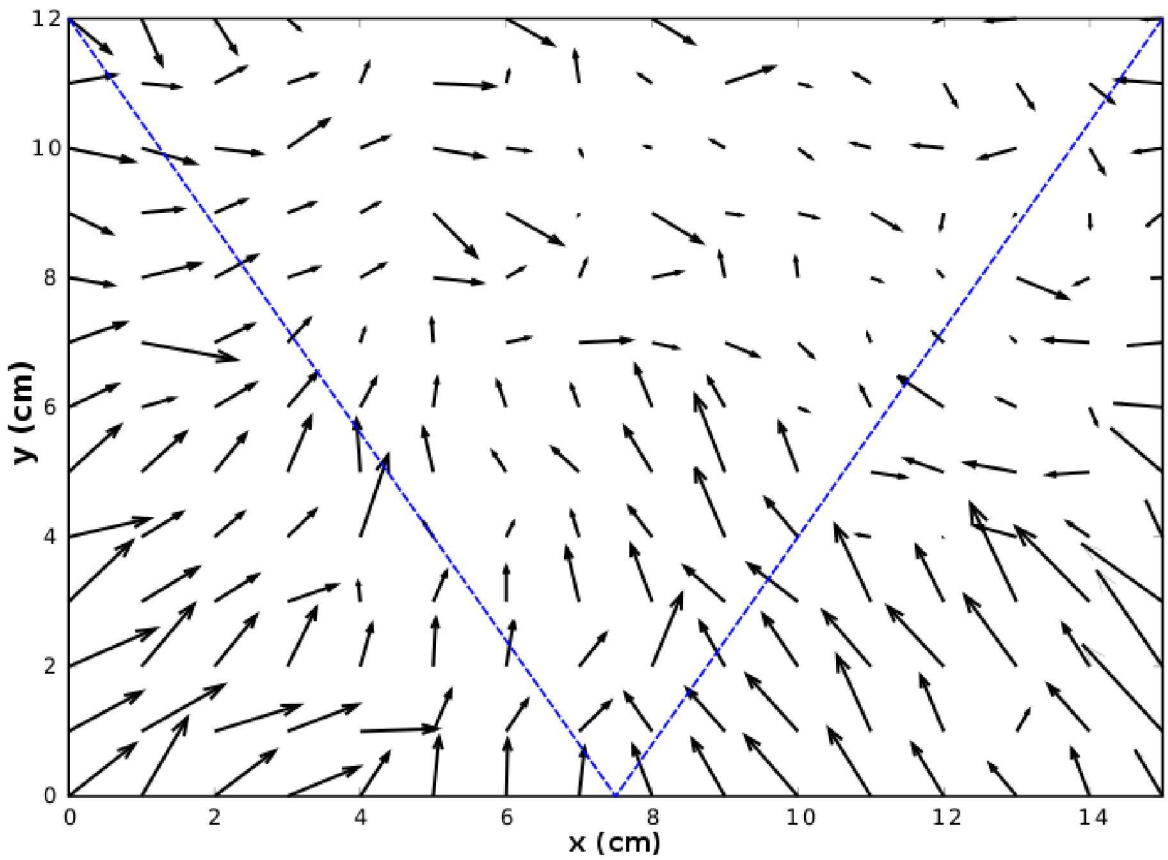

Figure 12: Error in calculated location of signal source.

5 Texture Identification

Once a stimulus is localized, the full spectrum of the source signal may be used to predict the texture rubbed against the skin. This task mimics one of the main tasks attributed to the Pacinian corpuscles in human skin. The spectrum of the source intensity will vary based on the mechanical properties of the material

rubbing against the skin [57], and the spectrum may be used as a feature for material classification. Several common machine learning approaches have been used in the past to perform this task. For example, texture sensitive fingertips have utilized artificial neural networks [26] and Bayesian classifiers [27]. As these models are very memory intensive, we implemented a logistic regression model to classify a detected stimulus as one of 15 predefined textures. Here, the likelihood that the spectrum of a signal, X, is produced by a given texture t is given by

$$y_t(X) = g(w_0 + w_1 X_1 + X_2 f_2 ... + w_n X_n) \tag{5}$$

where X_1 to X_n is the measured spectrum of the signal (in our application, the 128-bin Fourier spectrum), w_0 to w_n is a set of trained weights for the particular texture, $g(\circ)$ is the sigmoid function, and $y_t(\circ)$ is a value in the range $(0, 1)$ representing the likelihood that the texture t generated the measured signal. A logistic regression model was selected over a neural network in order to ensure that the model could be stored on each microcontroller in the skin. For $n = 128$ frequency components and 15 different textures, we need to store 1920 weights in the microcontroller, which requires a little less than 4kB of flash memory. With 128kByte of flash available on the Xmega platform, this approach can therefore scale to larger number of textures and more potent classifiers. Classification can be performed in 0.6 ms on the microcontroller.

5.1 Experimental Results

We performed a texture identification experiment using 15 textures, which are summarized in Table 1/Figure 13. We have cut each sample into a small patch of roughly one square inch size. For each texture, 100 samples were taken by rubbing the texture on the surface of the skin near a sensor by hand, and recording the transient signal measured by the node. We performed the rubbing by hand to introduce variation in pressure, speed and proximity to the sensor. The texture sample was rubbed over a region within 3 cm of the microphone, and the spectrum of the signal was sampled from the sensing node approximately every 5 seconds. The spectrum of the signal will vary based on the speed which the texture is rubbed (i.e., through shifting of the frequencies) and distance from the microphone (i.e., by attenuation of higher frequencies). As such, the variations introduced by generating the data in this manner ensures that a classification model is robust to these sorts of variations. A logistic regression model was trained on this data set using *Weka*, a machine learning library [58], and accuracy was assessed using 10-fold cross-validation. The logistic regression model was able to classify textures with an accuracy of 71.7%, which compares favorably to the expected accuracy of 6.7% from random guessing. To compare, a two-layer neural network was also trained using Weka. Classification was only slightly better with 73.1% accuracy.

Figure 14 shows the confusion matrix for the neural network classifier. The logistic regression confusion matrix is similar. As can be seen, errors are not specifically located at one specific spot on the off diagonal. The largest consistent misclassification is between cotton and dense foam. Cotton was misclassified as dense foam 14% of the time, and dense foam was misclassified as cotton 15% of the time. If these two classes are combined into a single class, the accuracy improves slightly (73.1% for the logistic regression model and 75.1% for the neural network model). The remaining errors are distributed relatively uniformly throughout the matrix.

To investigate the robustness of the logistic regression classifier, classification was performed with the

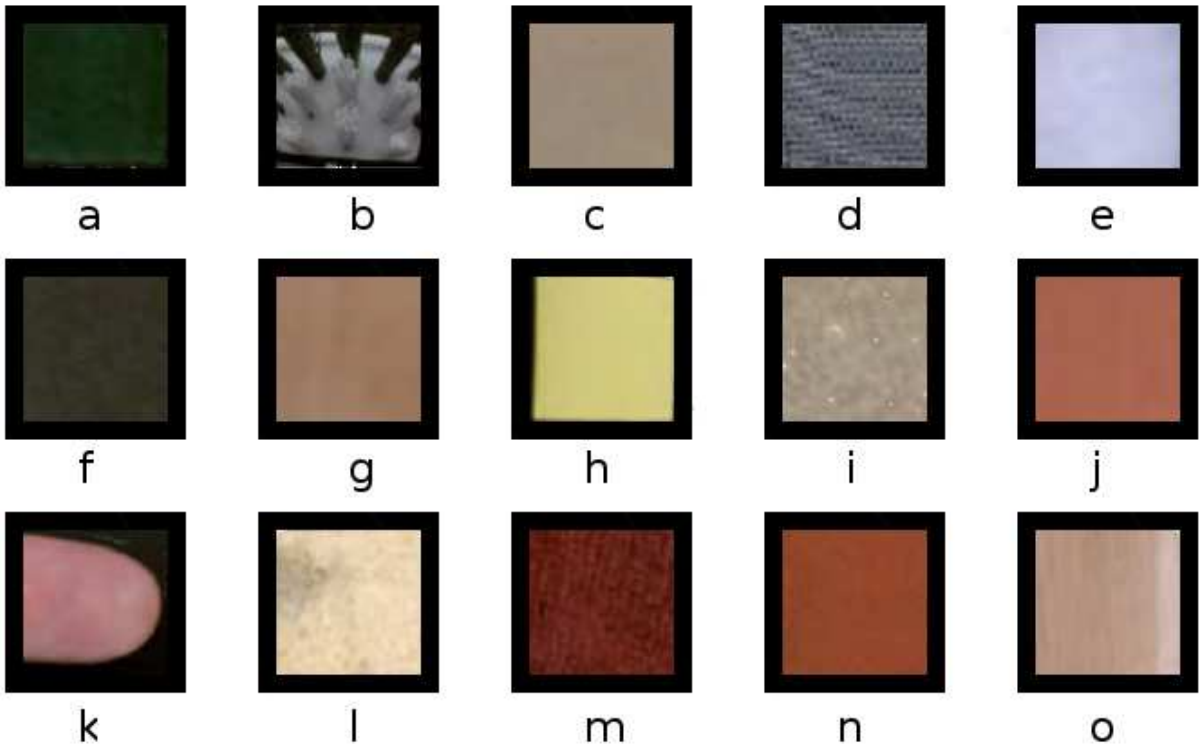

Figure 13: Textures used for training and validating the classifier.

Table 1: Textures used for training and validating the classifier.

a.	Brillo Pad	b.	Brush
c.	Cardboard	d.	Coarse Wire Mesh
e.	Cotton	f.	Dense Foam
g.	Fine Wire Mesh	h.	Plastic
i.	Sandpaper	j.	Silicone Foam
k.	Skin	l.	Sponge
m.	Terry Cloth	n.	Textured Silicone
o.	Wood		

trained classifier on data with added noise. The average energy level of each bin was determined for all samples. Noise was randomly generated with a mean value of zero, and a standard deviation a fixed percentage of the average energy of each bin. Classification was performed with noise with a standard deviation ranging from 0% to 50% of the average energy of each bin. For each standard deviation, 10 sets of data were produced by adding random noise. Figure 15 shows the average classifcation accuracy and standard deviation for the levels of added noise. Classification accuracy remains relatively high for relative noise levels up to about 15%. After this, the average accuracy drops by more than 5% of the original accuracy, and continues to decrease linearly. From Figure 9, the standard deviation of the measured signal reaches approximately 50% of the mean signal energy at 5-7cm. Therefore, for our prototype skin the expected classification accuracy at a single node is 43%. However, as the distance of the source is increasing from one node, and thus decreasing the expected classification accuracy, the distance to one or more other nodes is decreasing, and thus increasing the expected classification accuracy of those nodes. Combining predictions from several nodes may ensure a high classification accuracy.

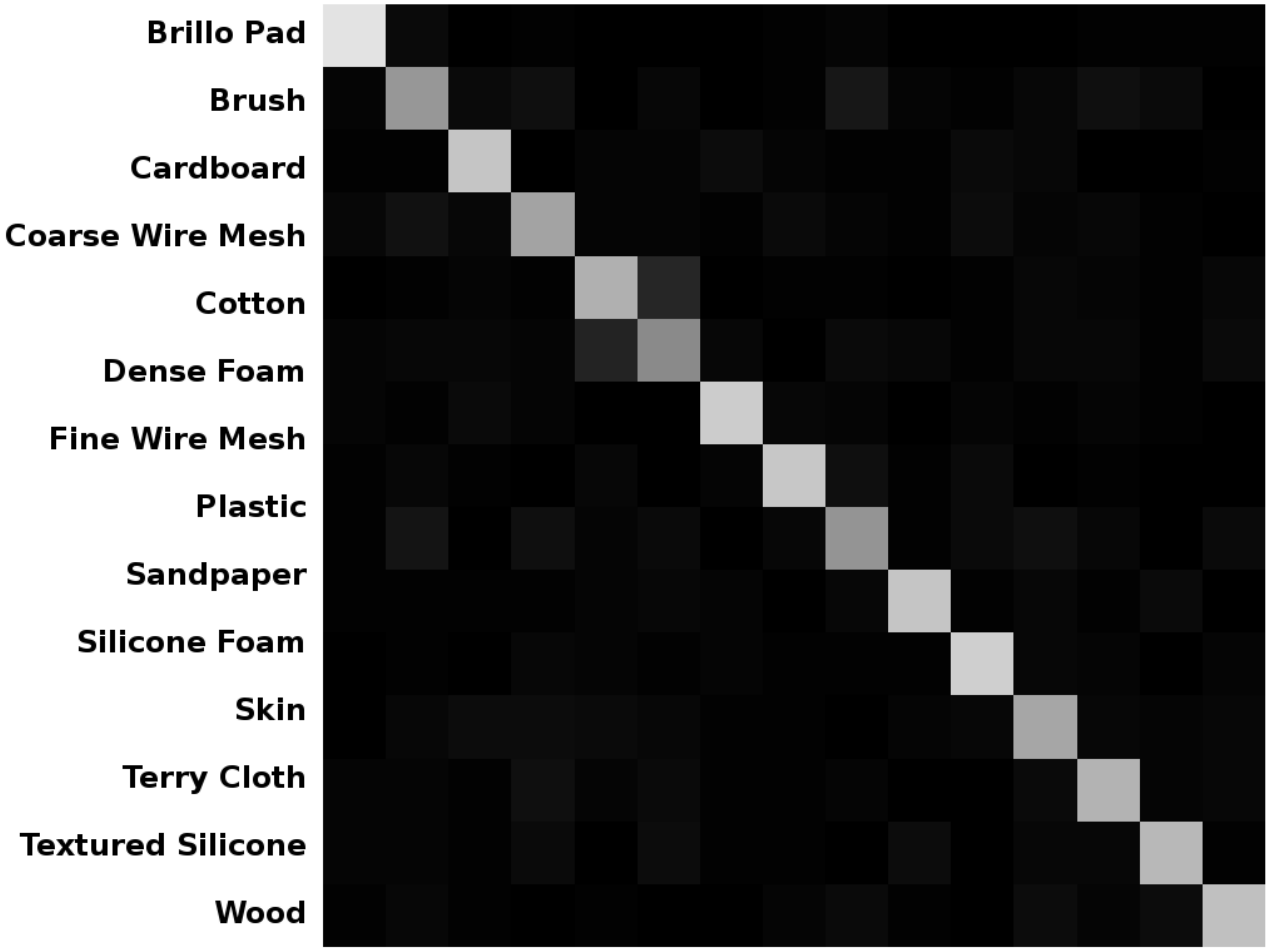

Figure 14: Confusion matrix for neural network classification of textures.

6 Conclusion and Future Work

We demonstrated the development and implementation of a soft, amorphous sensing skin that performs texture recognition, localization and event-driven data transport. Focusing on systems-level challenges in this paper, we limited this investigation to only one type of sensor (texture), localization of general sound sources, and data dissemination using a simple flooding protocol. With vibration due to texture having the highest bandwidth requirements among possible sensors (binary touch, pressure, temperature, capacitance, conductance, etc.) and the highest processing requirements (calculating a FFT vs. simply recording data or measuring changes), we believe that the proposed system can easily be extended to other sensor types and in-network processing algorithms, such as detecting patterns or shapes in a pressure profile. Our approach is enables skins to be formed to arbitrary shapes, and does not require uniform spacing of sensing nodes. However

The density of mechanoreceptors, and consequently the spatial resolution of the skin, vary at different locations in the human body. For example, the distance which the fingertip can correctly discriminate between two low-frequency vibrating signals is 0.8 to 1.2 mm, while the torso has a much coarser resolution of 2 - 3 cm [59]. Due to the large physical surface area, the torso provides twice as much tactile information as a fingertip. The prototype skin presented here provides similar spatial resolution as that of the human

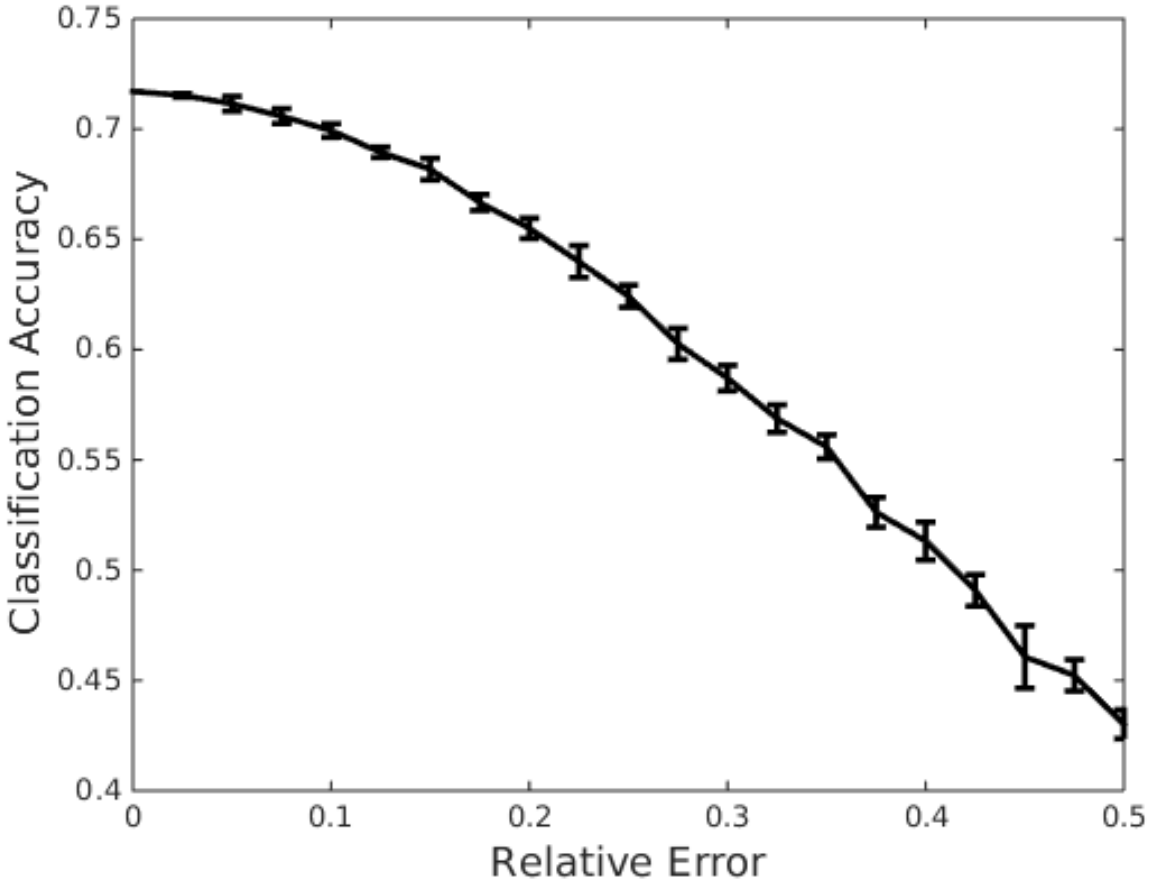

Figure 15: Accuracy of Logistic Regression classifier with noisy data.

torso, using only a few sensing nodes. To achieve the same resolution with an array of sensors capable of only local detection of stimulation would require thousands of sensing units.

Research into the mechanoreceptors in human skin indicate that the four mechanoreceptors provide different types of information to the brain–Merkel disks produce structural information of an object at high resolution, Meissner corpuscles provide motion signals and are critical for grip control and understanding the motion of objects the skin is touching, Pacinian corpuscles provide long distance vibration information transmitted through the skin or objects the skin is in contact, and Ruffini corpuscles give information about how skin is stretched, which may provide proprioceptive information such as joint angles [60]. These distinctions may provide valuable insight into the role artificial skins should play in robotics. High resolution pressure sensitive skins are common in robotic hands, and provide useful information about objects during grasping, similar to the role Merkel disks play. A multi-modal skin inspired by the mechanoreceptors in human skin could assist in other robotic tasks. In addition to texture recognition, the ability to detect vibration over large distances could be useful for controlling gait and determining the compliance of the surface walked on, through vibrations generated when a foot makes contact with the ground. Sensors mimicking Ruffini corpuscles could augment existing sensors for determining joint angles and robot poses. Understanding perceived information detected through the skin is typically performed subconciously–the human mind does not need to put conscious effort into determining texture, ensuring a stable grip, or

adjusting walking gait to account for changes in the ground. Similarly, allowing in-network processing of measurements from multiple sensors embedded in an amorphous skin allows the skin to provide high-level information to the central computer, enabling more computational time to be used for higher-level tasks.

The classification accuracy of textures (on the order of 70%) using simple logistic regression and a single sensor does not represent state-of-the-art accuracy. For example, recent results involving Bayesian classifiers [27] achieve an accuracy of 95.4% using a multi-modal sensing fingertip, combining pressure and vibration measurements. This approach included performing 36 exploratory measurements to determine texture, varying parameters such as touch pressure and velocity. Neural network classifiers using accelerometer data from a tactile sensitive fingertip [32] achieves an accuracy of 94.6%, though this requires a time window of 4 seconds, using features based on statistical values in the time and frequency domain of the accelerometer signal. Our approach, in constrast, respond to vibration only, and has no knowledge of the pressure or velocity used to measure the vibrations, and only relies on 0.25 second measurement windows. The required time to process a single measurement window is less than 20 ms. In addition, memory limitations of the sensing nodes provide a bottleneck to the complexity of any classifier model incorporated into a node. However, classification could be performed in a more controlled manner by simply measuring a particular sensor directly, once this information has been deemed important. Given an external controller, i.e., the robot on which the skin is attached, the skin could be rubbed over a texture in a controlled manner, with vibrations measured directly. Thus, the embedded classifier could be considered a preliminary means of detecting and discriminating textures, with the possibility of improved classification using external control for measurement and classification.

There may be several ways to improve classification performance within the network. The amorphous computing approach presented here allows for training different classifiers on each sensing node. This allows for implementing consensus classifiers [61] that compare texture signatures with those recorded by their neighbors, or through boosting by treating each sensor node which detected a texture as an individual classifier [62]. Automatic feature extraction is another area to be explored in texture recognition tasks. The approach described in this paper and [27, 32] relies on calculating the FFT of the measured signal. This requires significant computation in the sensing nodes, and result in features which may not be ideal to classify a particular set of textures. Shift invariant sparse coding [63] is an approach to represent a collection of time-series signals as the linear combination of a few sparse codes. While determining the codes is computationally intesive, this can be performed off-line using unlabeled texture data, and could result in more efficient feature extraction if the size and number of codes are sufficiently small. Similarly, self-taught learning [64] using autoencoders or generative neural network models (e.g., Restricted Boltzmann Machines) remain possible areas of exploration for more efficient feature extraction.

7 Acknowledgement

The authors would like to thank the reviewers for their helpful suggestions for improving this paper. The authors would like to recognize Nicholas Farrow for design of the sensing nodes. This work was supported by DARPA's WarriorWeb program, the Air Force Office of Scientific research under grant #FA9550-12-1-0145 and a NASA Early Career Faculty fellowship NNX12AQ47GS02.

A Sound Propagation in Robotic Skin

The derivation of a model of sound propagation is detailed here. For simplicity, the skin may be considered a thin vibrating plane. In practice, the skin is expected to be very large in terms of signal wavelength, thus it is assumed to be infinite in extent. This simplifies the analysis of sound propagation, as only a traveling wave radiating from a stimulus. A more complete analysis would consider reflections from the edge of the skin, connections, etc., and would consider various modes of vibration.

The displacement of an infinite plate vibrating at a single point is approximated by the equation

$$w(r,\omega) \simeq \frac{iF}{8\omega}\sqrt{\frac{2}{\pi\rho_s hDk_f r}}e^{i(k_f r-\pi/4)} \tag{6}$$

where F is the displacement force of the point source, ω is the frequency of the vibration, ρ_s, h and D are the density, thickness and rigidity of the plate, k_f is the wavenumber (ω/c), and r is the distance from the point source [65]. The sound intensity, $I(r,\omega)$, is proportional to the square of the displacement:

$$I(r,\omega) = \frac{1}{2}\rho_s c|w(r,\omega)|^2 \tag{7}$$

where c is the sound velocity in the skin. Substituting equation 6 into equation 7 yields

$$I(r,\omega) \simeq \frac{F^2}{64\omega\pi hDk_f^2 r} \tag{8}$$

Details of this derivation is available in [65].

The material properties of the skin (i.e., density, thickness and rigidity) may be considered constant. Isolating the terms associated with the material and stimulus results in

$$I(r,\omega) = \frac{I_0(F,\omega)}{r} \tag{9}$$

where $I_0(F,\omega)$ is related to the intensity of the vibration signal at the source of the vibration, and is given by

$$I_0(F,\omega) = \frac{F^2}{64\omega\pi hDk_f^2} \tag{10}$$

B Localization of Vibration Stimulus

Given measurements from multiple sensing nodes, estimating the location of the stimulus is easily expressed as an error minimization problem. The approach to estimating the location given here is based on measurements from three sensors, as shown in Figure 7 but may be extended to more sensors.

Consider a vibration induced in the skin at point (x,y), with source intensity $I_0(F,\omega)$, as described by equation 10. For purposes of localization, a single frequency in the spectrum of the vibration (i.e., the frequency bin of the FFT of the measured signal with the largest amplitude) is used. This simplifies the source intensity to a constant value, referred to as I_0 in this derivation.

As the signal propagates through the skin, it's intensity decreases inversely as a function of the distance from the source, as given by equation 3. The expected intensity of the vibration measured at node i, $\hat{I}_i$,

is simply

$$\hat{I}_i = \frac{I_0}{r_i} \tag{11}$$

where r_i is the distance from the source of the vibration to the sensing node

$$r_i = \sqrt{(x - x_i)^2 + (y - y_i)^2} \tag{12}$$

Assuming no noise, the difference between the intensity measured at node i, I_i, and $\hat{I}_i$ will be zero. Using measurements at three nodes, the position of the source (x, y) is given by minimizing the sum of the squared differences between expected and measured intensities at each node

$$x, y = \arg\min_{x,y} \sum_{i \in \{1,2,3\}} \left(\frac{I_0}{r_i} - I_i \right)^2 . \tag{13}$$

The displacement force F of the source of the vibration is not known. Threfore, the value of I_0 cannot be determined and must be removed from this equation. This can be done by combining the measurements at two nodes, whose intensity is expressed as

$$\begin{aligned} I_1^2 &= \frac{I_0^2}{(x - x_1)^2 + (y - y_1)^2} \\ I_2^2 &= \frac{I_0^2}{(x - x_2)^2 + (y - y_2)^2} \end{aligned} \tag{14}$$

where nodes 1 and 2 are located at positions (x_1,y_1) and (x_2,y_2), respectively, and have a measured intensity of I_1 and I_2. These two equations can be rewritten as

$$\begin{aligned} I_1^2((x - x_1)^2 + (y - y_1)^2) &= I_0^2 \\ I_2^2((x - x_2)^2 + (y - y_2)^2) &= I_0^2 \end{aligned} \tag{15}$$

The difference between these two equations results in the following single equation,

$$I_1^2((x - x_1)^2 + (y - y_1)^2) - I_2^2((x - x_2)^2 + (y - y_2)^2) = 0 \tag{16}$$

Given measurements from N nodes, it is possible to derive $N - 1$ independent equations of the above form. Minimizing the error (i.e., the deviation from zero) in these equations with respect to (x, y) can be performed using gradient decent. The equations are convex, so gradient decent is guaranteed to converge on a global minimum.

C Localization Uncertainty

Derivation of the error in the location of a source signal given noisy measurements is given here. Assume three sensing nodes have detected a signal. The signal intensities measured at each node, I_i, are sampled from Gaussian random variables with mean given by equation 8 and standard deviation σ_i

$$I_i(\omega, r) = \mathcal{N}\left(\frac{I_0(\omega)}{r}, \sigma_i \right) \tag{17}$$

The estimate of the source position was determined iteratively using Equation 4, and an equation directly estimating the source position as a function of measured intensities at the sensor nodes is not available. Equation 8 provides a function relating the measured intensities to source position. Using the explicit position of the source (x, y) and the measurement at three sensor nodes, this equation can be written as

$$\begin{bmatrix} I_1(x,y) \\ I_2(x,y) \\ I_3(x,y) \end{bmatrix} = \begin{bmatrix} \frac{I_0(x,y)}{\sqrt{(x-x_1)^2+(y-y_1)^2}} \\ \frac{I_0(x,y)}{\sqrt{(x-x_2)^2+(y-y_2)^2}} \\ \frac{I_0(x,y)}{\sqrt{(x-x_3)^2+(y-y_3)^2}} \end{bmatrix} \tag{18}$$

The frequency term ω is not explicitly represented in this equation for clarity, and analysis may be considered for a single frequency bin measurement. The uncertainty of the measurements, expressed using the covariance matrix Σ^I, can be calculated from the covariance matrix of the source position, Σ^{xy}, and the Jacobian matrix of Equation 18. The Jacobian matrix is defined as

$$J(x,y) = \frac{\partial I}{\partial x, y} \tag{19}$$

which is calculated as

$$J(x,y) = -I_0 \begin{bmatrix} \frac{(x-x_1)}{((x-x_1)^2+(y-y_1)^2)^{3/2}} & \frac{(y-y_1)}{((x-x_1)^2+(y-y_1)^2)^{3/2}} \\ \frac{(x-x_2)}{((x-x_2)^2+(y-y_2)^2)^{3/2}} & \frac{(y-y_2)}{((x-x_2)^2+(y-y_2)^2)^{3/2}} \\ \frac{(x-x_3)}{((x-x_3)^2+(y-y_3)^2)^{3/2}} & \frac{(y-y_3)}{((x-x_3)^2+(y-y_3)^2)^{3/2}} \end{bmatrix} \tag{20}$$

Using this, the uncertainty of a measurements due to uncertainty in the source position is approximated by

$$\Sigma^I \approx J\Sigma^{xy}J^T \tag{21}$$

Calculating the uncertainty of the source position from the measurements simply requires left and right multiplying equation 21 by the pseudoinverse of the Jacobian

$$\Sigma^{xy} \approx (J^TJ)^{-1}J\Sigma^I J^T(J^TJ)^{-1} \tag{22}$$

Measurement errors are assumed to be due to local variations in the sensing nodes (e.g., variations in microphones and amplifiers, quantification errors in the ADC, etc.). Therefore, measurement errors will be considered independent, and Σ^I given by

$$\Sigma^I = \begin{bmatrix} \sigma & 0 & 0 \\ 0 & \sigma & 0 \\ 0 & 0 & \sigma \end{bmatrix} \tag{23}$$

assuming all node measurements have the same standard deviation.

References

[1] R. S. Johansson and J. R. Flanagan. Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience*, 10(5):345–359, 2009.

[2] G. Robles-De-La-Torre. The importance of the sense of touch in virtual and real environments. *IEEE Multimedia*, 13(3):24–30, July–September 2006.

[3] R. S. Dahiya, G. Metta, M. Valle, and G. Sandini. Tactile sensing–from humans to humanoids. *IEEE Transactions on Robotics*, 26(1):1–20, February 2010.

[4] A. Jain, M. D. Killpack, A. Edsinger, and C. C. Kemp. Manipulation in clutter with whole-arm tactile sensing. *International Journal of Robotics Research*, 32(4):458–482, April 2013.

[5] V. Duchaine, N. Lauzier, M. Baril, M.-A. Lacasse, and C. Gosselin. A flexible robot skin for safe physical human robot interaction. In *Proceedings of the International Conference on Robotics and Automation*, pages 3676–3681, 2009.

[6] K. Hsiao, L. P. Kaelbling, and T. Lozano-Perez. Task-driven tactile exploration. In *Proceedings of Robotics: Science and Systems*, 2010.

[7] J. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K. Kuchenbecker. Human-inspired robotic grasp control with tactile sensing. *IEEE Transactions on Robotics*, 27(6):1067–1079, December 2011.

[8] P. Mittendorfer, E. Yoshida, T. Moulard, and G. Cheng. A general tactile approach for grasping unknown objects with a humanoid robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013.

[9] L. Ma, M. Ghafarianzadeh, D. Coleman, N. Correll, and G. Sibley. Simultaneous localization, mapping, and manipulation for unsupervised object discovery. In *IEEE International Conference on Robotics and Automation*, 2015.

[10] R. S. Dahiya, P. Mittendorfer, M. Valle, G. Cheng, and V. J. Lumelsky. Directions toward effective utilization of tactile skin: A review. *IEEE Sensors Journal*, 13(11):4121–4138, November 2013.

[11] R. S. Dahiya and M. Valle. *Robotic Tactile Sensing*. Springer, 2013.

[12] D. Anghinolfi, G. Cannata, F. Mastrogiovanni, C. Nattero, and M. Paolucci. On the problem of the automated design of large-scale robot skin. *IEEE Transactions on Automation Science and Engineering*, PP(99):1–14, April 2013.

[13] M. A. McEvoy and N. Correll. Materials that couple sensing, actuation, computation, and communication. *Science*, 347(6228):1261689, 2015.

[14] D. Hughes and N. Correll. A soft, amorphous skin that can sense and localize textures. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1844–1851, 2014.

[15] K. Weiss and H. Woern. Tactile sensor system for an anthropomorphic robotic hand. In *Proceedings of IEEE International Conference on Manipulation and Grasping (IMG 2004)*, pages 12–17, 2004.

[16] V. Maheshwari and R. F. Ravi. High-resolution thin-film device to sense texture by touch. *Science*, 312(5779):1501–1504, 2006.

[17] H. Zhang and E. So. Hybrid resistive tactile sensing. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 32(1):57–65, 2002.

[18] G. Cannata, M. Maggiali, G. Metta, and G. Sandini. An embedded artificial skin for humanoid robots. In *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pages 434–438, 2008.

[19] R. J. De Souza and K. D. Wise. A very high density bulk micromachined capacitive tactile imager. In *International Conference on Solid State Sensors and Actuators*, volume 2, pages 1473–1476. IEEE, 1997.

[20] J.-S. Heo, J.-H. Chung, and J.-J. Lee. Tactile sensor arrays using fiber bragg grating sensors. *Sensors and Actuators A: Physical*, 126(2):312–327, 2006.

[21] M. Ohka, H. Kobayashi, J. Takata, and Y. Mitsuya. Sensing precision of an optical three-axis tactile sensor for a robotic finger. In *IEEE International Symposium on Robot and Human Interactive Communication*, pages 214–219. IEEE, 2006.

[22] T. Nelson, R. vanDover, S. Jin, S. Hackwood, and G. Beni. Shear-sensitive magnetoresistive robotic tactile sensor. *IEEE Transactions on Magnetics*, 22(5):394–396, 1986.

[23] E. S. Kolesar, C. S. Dyson, R. R. Reston, R. C. Fitch, D. G. Ford, and S. D. Nelms. Tactile integrated circuit sensor realized with a piezoelectric polymer. In *IEEE International Conference on Innovative Systems in Silicon*, pages 372–381. IEEE, 1996.

[24] J. Edwards, J. Lawry, J. Rossiter, and C. Melhuish. Extracting textural features from tactile sensors. *Bioinspiration & Biomimetics*, 3(3):3–12, 2008.

[25] W. W. Mayol-Cuevas, J. Juarez-Guerrero, and S. Munoz-Gutierrez. A first approach to tactile texture recognition. In *IEEE International Conference on Systems, Man, and Cybernetics*, volume 5, pages 4246–4250. IEEE, 1998.

[26] D. Taddeucci, C. Laschi, R. Lazzarini, R. Magni, P. Dario, and A. Starita. An approach to integrated tactile perception. In *Proceedings of the International Conference on Robotics and Automation*, pages 3100–3105, 1997.

[27] J. A. Fishel and G. E. Loeb. Bayesian exploration for intelligent identification of textures. *Frontiers in Neurorobotics*, 6, 2012.

[28] T. Hoshi and H. Shinoda. A large area robot skin based on cell-bridge system. In *Proceedings of the 5th IEEE Conference on Sensors*, pages 827–830, 2006.

[29] S. Mannsfeld, B. Tee, R. Stoltenberg, C. Chen, S. Barman, B. Muir, A. Sokolov, C. Reese, and Z. Bao. Highly sensitive flexible pressure sensors with microstructured rubber dielectric layers. *Nature Materials*, 9:859–864, 2010.

[30] J. Ulmen and M. Cutkosky. A robust, low-cost and low-noise artificial skin for human-friendly robots. In *Proceedings of the International Conference on Robotics and Automation*, pages 4836–4841, 2010.

[31] A. Billard, A. Bonfiglio, G. Cannata, P. Cosseddu, T. Dahl, K. Dautenhahn, F. Mastrogiovanni, G. Metta, L. Natale, B. Robins, L. Seminara, and M. Valle. The roboskin project: challenges and results. In *Romansy 19–Robot Design, Dynamics and Control*, pages 351–358. Springer, 2013.

[32] P. Giguere and G. Dudek. A simple tactile probe for surface identification by mobile robots. *IEEE Transactions on Robotics*, 27(3):534–544, 2011.

[33] T. Bhattacharjee, J. M. Rehg, and C. C. Kemp. Haptic classification and recognition of objects using a tactile sensing forearm. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4090–4097. IEEE, 2012.

[34] J. Schill, J. Laaksonen, M. Przybylski, V. Kyrki, T. Asfour, and R. Dillmann. Learning continuous grasp stability for a humanoid robot hand based on tactile sensing. In *IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, pages 1901–1906. IEEE, 2012.

[35] P. Mittendorfer, E. Yoshida, and G. Cheng. Realizing whole-body tactile interactions with a self-organizing, multi-modal artificial skin on a humanoid robot. *Advanced Robotics*, 29(1):51–67, 2015.

[36] N. Correll, Ç. D. Önal, H. Liang, E. Schoenfeld, and D. Rus. Soft autonomous materialsusing active elasticity and embedded distributed computation. In *Experimental Robotics*, pages 227–240. Springer Berlin Heidelberg, 2010.

[37] Y. Ohmura, Y. Kuniyoshi, and A. Nagakubo. Conformable and scalable tactile sensor skin for curved surfaces. In *Proceedings of the International Conference on Robotics and Automation*, pages 1348–1353, 2006.

[38] E. Baglini, S. Youssefi, F. Mastrogiovanni, and G. Cannata. A real-time distributed architecture for large-scale tactile sensing. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2014)*, pages 1663–1669. IEEE, 2014.

[39] A. Vallbo and R. Johannson. Properties of cutnaeous mechanoreceptors in the human hand related to touch sensation. *Human Neurobiology*, 3:3–14, 1984.

[40] R. Johannson and A. Vallbo. Tactile sensibility in the human hand: Relative and absolute densities of four types of mechanoreceptive units in glabrous skin. *Journal of Physiology*, 286:283–300, 1979.

[41] J. Scheiber, S. Leurent, A. Prevost, and G. Debrégeas. The role of fingerprints in the coding of tactile information probed with a biomimetic sensor. *Science*, 323(5920):1503–1506, March 2009.

[42] S. Bensmaïa and M. Hollins. Pacinial representations of fine surface texture. *Perception & Psychophysics*, 67(5):842–854, 2005.

[43] A. Zimmerman, L. Bai, and D. D. Ginty. The gentle touch receptors of mammalian skin. *Science*, 346(6212):950–954, 2014.

[44] J. B. F. Van Erp. Vibrotactile spatial acuity on the torso: effects of location and timing parameters. In *Eurohaptics Conference, 2005 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2005. World Haptics 2005. First Joint*, pages 80–85. IEEE, 2005.

[45] S. B. Holt. Genetics of dermal ridges: the relation between total ridge-count and the variability of counts from finger to finger. *Annals of Human Henetics*, 22(4):323–339, 1958.

[46] Washington Mills. Particle size conversion chart. `http://www.washingtonmills.com/guides/grit-sizes-ansi/particle-size-conversion-chart-ansi/`. Accessed June 4, 2015.

[47] H. Alexander and D. L. Miller. Determining skin thickness with pulsed ultra sound. *Journal of Investigative Dermatology*, 72(1):17–19, 1979.

[48] M. A. McEvoy and N. Correll. Thermoplastic variable stiffness composites with embedded, networked sensing, actuation, and control. *Journal of Composite Materials*, page 0021998314525982, 2014.

[49] P. G. Agache, C. Monneur, J. L. Leveque, and J. De Rigal. Mechanical properties and young's modulus of human skin in vivo. *Archives of Dermatological Research*, 269(3):221–232, 1980.

[50] N. Correll, N. Farrow, and S. Ma. Honeycomb: a platform for computational robotic materials. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction*, pages 419–422. ACM, 2013.

[51] H. Hosseinmardi, R. Han, and N. Correll. Bloom filter-based ad hoc multicast communication in cyber-physical systems and computational materials. In *The 7th International Conference on Wireless Algorithms, Systems, and Applications (WASA 2012)*, volume 7405 of *Lecture Notes in Computer Science*, 2012.

[52] S. Ma, H. Hosseinmardi, N. Farrow, R. Han, and N. Correll. Establishing multi-cast groups in computational robotic materials. In *IEEE International Conference on Cyber, Physical and Social Computing*, 2012.

[53] R. Martin. An efficient algorithm to estimate the instantaneous snr of speech signals. In *Proceedings of the 2nd European Conference on Speech, Communication and Technology (EUROSPEECH '93)*, pages 1093–1096, 1993.

[54] H. Profita, N. Farrow, and N. Correll. Flutter: An exploration of an assistive garment using distributed sensing, computation and actuation. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*, pages 359–362. ACM, 2015.

[55] J.-M. Valin, F. Michaud, J. Rouat, and D. Létourneau. Robust sound source localization using a microphone array on a mobile robot. In *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, pages 1228–1233, 2003.

[56] A. A. Clifford. *Multivariate Error Analysis: A Handbook of Error Propagation and Claculation in Many-Parameter Systems*. John Wiley & Sons, 1973.

[57] S. Decherchi, P. Gastaldo, R. S. Dahiya, M. Valle, and R. Zunino. Tactile data classification of contact materials using computational intelligence. *IEEE Transaction on Robotics*, 27(3):635–639, June 2011.

[58] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I.H. Witten. The weka data mining software: an update. *SIGKDD Explorations*, 11(1), 2009.

[59] L. A. Jones, M. Nakamura, and B. Lockyer. Development of a tactile vest. In *Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2004. HAPTICS'04. Proceedings. 12th International Symposium on*, pages 82–89. IEEE, 2004.

[60] K. O. Johnson. The roles and functions of cutaneous mechanoreceptors. *Current opinion in neurobiology*, 11(4):455–461, 2001.

[61] J. A. Benediktsson and P. H. Swain. Consensus theoretic classification methods. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(4):688–704, July 1992.

[62] R. Meir and G. Rätsch. An introduction to boosting and leveraging. In *Advanced lectures on machine learning*, pages 118–183. Springer, 2003.

[63] M. S. Lewicki and T. J. Sejnowski. Coding time-varying signals using sparse, shift-invariant representations. *Advances in neural information processing systems*, pages 730–736, 1999.

[64] R. Raina, A. Battle, H. Lee, B. Packer, and A. Y. Ng. Self-taught learning: transfer learning from unlabeled data. In *Proceedings of the 24th international conference on Machine learning*, pages 759–766. ACM, 2007.

[65] M. C. Junger and D. Feit. *Sound, Structures, and Their Interaction.* MIT Press, 1972.