

Conductive Fur Sensing for a Gesture-Aware Furry Robot

Anna Flagg₁*

Diane Tam₁

Karon MacLean₁†

Robert Flagg

¹Department of Computer Science
University of British Columbia
Vancouver, Canada

ABSTRACT

Recent advances in artificial intelligence suggest that machines will soon be capable of communicating in ways previously considered out of their reach. For example, humans engage in sophisticated emotional communication through the language of touch. What technical capabilities would enable computers to do the same?

As our group examines this question in the context of emotional touch between a person and a furry social robot, we require sensors designed to detect and recognize subtle, nuanced touches. To this end, we demonstrate a new type of sensor based on conductive fur, which is sensitive to movements unavailable to conventional pressure sensors. The sensor captures motion by measuring changing current as the fur's conductive threads connect and disconnect during touch interaction. We then use machine learning to classify gestures from this time series. An informal evaluation with seven participants found 82% recognition of a 3-gesture set, showing promise for this approach to gesture recognition, and opening a path to emotionally intelligent touch sensing.

Index Terms: Gesture recognition, haptics, social robot, touch sensor, conductive fur, fur interface.

1 INTRODUCTION

The human brain is not purely rational; rather it carries out a complex combination of thinking and feeling. Picard [14] argues that therefore, a truly natural symbiosis between people and machines cannot exist without harnessing emotion. Early work in emotional computing has raised a range of controversial questions about the possible roles of emotion in computers, whether for artificial perception, expression, or even possession of emotion. What is clear is that the design of emotionally intelligent haptic experiences offers exciting and important possibilities. Touch-based social robots have been used for empathic communication, and are capable of providing emotional support and companionship. Affective touch is especially important for the development and well-being of the young, the old, the ill and the troubled. There are thus many valuable social and healthcare-related applications, including rehabilitation, education, treatment of cognitive disorders, and assistance for people with special needs [11, 4, 15]. Since there is considerable evidence that we find it natural to attribute emotional qualities to non-animate things [8, 5], lack of emotional intelligence may actually hamper a therapeutic robot's effectiveness, making it crucial that social robots become more literate in emotional communication.

Current emotion sensing exists in several forms, including touch, posture, speech, voice tonality, and physiological measures. While these are promising directions, none individually tells enough of the story to model emotion accurately [15, 6, 16, 3]. In the case of



Figure 1: Conductive fur sensor.

touch sensing, current social robots - which rely largely on force and electric field sensors - would not be able to classify gestures adequately even if used in combination (Section 2).

This suggests the need for an additional channel of information. In this work, we design, apply initial gesture recognition to and evaluate a new fur-based touch sensor based on above-surface hand motion information (Figure 1). We plan in future to fuse other complementary technologies with output from these conductive fur sensors, where the fur patches are placed to cover a whole robot body. The eventual goal of this work is creation of a gesture-sensitive furry robot that can accurately parse and respond to human emotion.

Our physical design is inspired by Perner-Wilson and Satomi's stroke sensor [13, 12] (Figure 2), which closes a circuit composed of conductive threads sewn into a fur-like configuration (Section 2.4). We expand upon this seed idea by 1) integrating multiple touch-sensitive layers of conductive threads into animal-like fur, 2) developing a method to extract time-series hand motion information, and 3) modeling that information to classify gesture. A preliminary evaluation suggests this design could contribute to gesture recognition.

2 RELATED WORK

2.1 Touch-Sensitive Social Robots

Huggable, PARO, Aibo and Probo are other affective robots that are sensitive to touch. Huggable [16] is the most relevant due to its full-body sensitive skin. This furry teddy bear companion has a wide range of sensors around its body, and its initial gesture recognition engine identifies nine classes of touch: tickle, poke, scratch, slap, pet, pat, rub, squeeze and contact. The interactive robotic seal PARO [15] can also recognize patterns in its environment, including common phrases like greetings, praise and its own name, and it has a long-term memory of owner touch behavior. Aibo [6] is an interactive robot dog that grows gradually from a puppy personality to a mature dog through interaction with people and its environment. It is able to process its environment and human behavior patterns in many ways, including recognizing its owner, locating its charging station, learning tricks, etc. Finally Probo [7] is an elephant-like

*e-mail:me@annaflagg.com

†e-mail:maclean@cs.ubc.ca

social robot used to ease anxiety in hospitalized children. Equipped with a large variety of sensors, it infers seven classes of emotion: anger, disgust, fear, happiness, sadness, surprise, and sleepiness.

The goal of this work is to contribute to the gesture recognition capabilities of another touch-sensitive affective robot, the Haptic Creature [19]. An animal-like but deliberately non-representational robot, the Haptic Creature senses the world through touch alone, with a focus on identifying human emotional states from touch gestures. Early experimentation with gesture recognition models *stroke*, *slap*, *poke*, *pat*, *pinch* and *squeeze*, with accuracy related to density of sensor coverage [3].

2.2 Force, Electric Field, and Temperature

Current touch sensing in affective systems is largely focused around force sensors such as Force Sensitive Resistors (FSRs) and Quantum Tunneling Composites (QTCs). PARO, Aibo and the Haptic Creature all use FSRs alone to identify touch [15, 6, 3]. Huggable uses capacitive sensors, which measure proximity of a body via electric field, as well as temperature and FSRs [16].

While these directions are promising, they are still in early stages of gesture recognition, and none will have the needed sensing scope. Huggable contains the most advanced recognition engine, but it uses over 1500 sensors, relies partly on location of touch to define gesture, and does not have complete recognition capabilities [16]. The capacitive sensors may be vulnerable to interference and are quite expensive [1]. FSRs are inexpensive, but don't function well on curved surfaces, and production scales poorly to continuous coverage. They are also insensitive to light touches, including those that interact with the fur above the "skin" surface. QTCs are less affected by curved surfaces and potentially more sensitive to lighter touches [10], but we have found they exhibit intractable nonlinearity, and they are not easily available at this time.

It is therefore of interest to investigate alternate sensor types that could improve recognition accuracy by providing a different channel of information for affective touch.

2.3 Gestures to Recognize

To specify sensor requirements, we reference Yohanan's dictionary of affective gestures that a furry social robot should be able to recognize, collected through video observation of touching scenarios [20]. Consider the example of a scratch and a rub: both involve similar pressure and proximity, and it is the subtleties of hand position and motion over time that defines each gesture. A rub involves moving the pads of the fingers along the skin to exert force, and a scratch involves the nails ruffling the fur and making high-pressure contact against its length, especially at the roots. Sensing force, temperature and proximity would not provide this information. Similarly, the hand motion involved in a pinch and a poke is very different, but might involve about the same hand force and proximity. A tap and a weak slap involve different hand positions but could exert roughly similar force.

There are many such examples. Thus, tasked with improving the Haptic Creature's sensing capabilities using inexpensive, computationally-efficient and hardware-light methods, we explore above-surface hand motion information as an input to touch sensing.

2.4 Sensing Above-Surface Hand Motion

Perner-Wilson and Satomi describe the concept for a low-tech stroke sensor [12]: a circuit made up of conductive threads in insulating fabric is sewn vertically like fur, and connected to an LED for output. When stroked, the threads brush against each other, closing the circuit and providing power to turn on the light. When the hand moves away, the threads return to their vertical state where they are no longer touching, breaking the circuit and turning off the light.

This creates a fur-like interface capable of communicating binary information corresponding to whether or not it is being stroked.

Perner-Wilson and Satomi's sensor is configured to indicate just one of two states (circuit open/closed), although this is not an inherent limitation of the concept as we show in our extension, assuming sufficient parallelism of threads for multiple circuit closings.

This version of the stroke sensor is most successful with one particular gesture - a long stroke from one end to the other. A long stroke has the property of moving a long line of threads together, which is required to make a physical connection between the two wired ends of the fabric. This also restricts the physical configuration of the threads, which must be positioned far enough apart to avoid extraneous, non-stroked touching. Low thread density then limits sensitivity.

When only binary open/closed information is used, sensor output is also not rich enough to differentiate the subtleties of different types of gestures; such analysis was indeed not part of the original presentation.

In the present research, we apply the idea in Perner-Wilson and Satomi's thread sensor to a new, considerably more sensitive above-surface fur sensor design from which we can collect and analyze continuous hand motion data for gesture classification (Sections 3-5).

2.5 Gesture Recognition Technologies

The use of machine learning for touch gesture recognition in affective systems is in early stages. The designers behind both Huggable and PARO have experimented with supervised neural networks using feature-based sensor data [16, 15]. The Haptic Creature team has also made use of features, with an eventual probabilistic structure in mind [3].

One approach is the use of learning schemes specifically designed for data mining of time series; to our knowledge, it is unexplored for gesture recognition, a surprising gap given the time-dependent nature of gestures. For instance, the use of 1-nearest neighbor coupled with dynamic time warping does not seem to have been investigated for gesture recognition, yet it is considered possibly the most effective method of classifying time-series data [18]. Another promising direction would be fast Fourier transforms or wavelets.

In summary, time-series-specific classification is an unexplored but conceptually promising approach to gesture recognition.

3 APPROACH

Our goal is to assess the value of augmenting existing sensing technologies with the conductive fur sensor: specifically, are there classes of gesture that the fur can model more accurately? As a proof of feasibility, we focus on identifying a few key gestures from Yohanan's touch dictionary [20], leaving inference of the emotional content to later work. We have selected *stroke*, *scratch* and *light touch* on the basis of crucial affective content [14], inadequate differentiation by existing sensor technology, and a potentially good match to the fur-based sensor.

These gestures are defined as follows by Yohanan [20]:

stroke: moving one's hand gently over the fur, often repeatedly,

scratch: rubbing the fur with one's fingernails,

light touch: touching the fur with light finger movements (referred to as *tickle* in the Yohanan paper [15]).

Our recognition approach is based on the observation that during a touch interaction between a human and a furry animal, the hand disturbs the configuration of the animal's fur, with an arguably distinctive pattern. We are interested in capturing physical changes in the fur for visibility into the gesture space.

In the sensor concept which we have adopted from Perner-Wilson and Satomi [12], a stroking motion brushes the vertically-sewn conductive threads together. When a pair of adjacent threads




Figure 2: Circuit for our design. Touches change the fur configuration and consequently the net fur resistance, R_{fur} . The resulting fluctuating current= $f(\text{time})$ is sampled at 144 Hz (I_{sense}).

do not touch, they present infinite resistance to the circuit, and a finite resistance when they do touch. This circuit is effectively made of many resistors connected in parallel; its total resistance drops as more connections are made, and hence measurable current increases (Figure 2).

3.1 Physical design and data collection

We build upon Perner-Wilson and Satomi’s idea in several ways, including both changes to the physical design, and an added analytical component. Changes to the physical design include:

Thick, animal-like fur: We sew conductive threads into a sample of the fur that is used in the Haptic Creature, creating a circuit with the conductive threads functioning as wires. Thick fur (rather than Perner-Wilson and Satomi’s individual sensor threads) maintains realism and visual, tactile attractiveness (Figure 3).

Sampling current over time: Rather than sampling a single *stroke* or *no stroke* state, we sample current over time, $I(t)$. Current fluctuates according to connections between the threads, i.e., changes in the fur’s physical state that occur during a touch.

Dense threads: Using $I(t)$ allows us to position the threads more densely, because we are no longer restricted to maintaining a broken circuit when the threads are not being stroked. Dense thread configuration is desirable because it increases the touch-sensitive coverage of the fur, making an arbitrarily-placed small touch more likely to be noticed.

Optimizing thread patterns for sensitivity: We iterate on thread patterns to increase sensitivity without saturating the circuit (discussed in Section 4.3). A saturated signal $I(t)$ would be insensitive to gesture nuances.

Distinguishing touch types with layering: In our most successful design, the fur includes two integrated “layers” of conductive thread lengths. One is the same length as the Creature’s regular fur, and the other about a third as long, activated only with more aggressive, deeper disturbances. Their combined data reflects the differences between gestures that activate different parts of the fur. Details follow in Section 5.

3.2 Analysis:

As a first step in evaluating both the basic sensor concept (previously untested) and our enhancements, we collect and analyze data for our three test gestures, focusing on feature-based machine learning methods.

Feature extraction: We extract several standard statistics as features for training classifiers, making use of analytical methods as




Figure 3: Perner-Wilson and Satomi’s conductive thread stroke sensor (left) [12], our conductive fur touch and gesture sensor (right).

well as visual analysis of density curves to select the most powerful features (see Section 5).

Classification: We train a logistic regression model on our extracted feature data, and measure its accuracy on a test set. The results, as well as our preliminary comments about the sensor’s suitability as a potential data source for classification, are described in Section 5.

4 SENSOR DESIGN

4.1 Objectives

The conductive fur sensor is designed with many physical and performance-related requirements in mind, including:

Practicality: The underlying hardware must be small and lightweight to facilitate the system’s intended functions as a laptop. Since the application involves an anthropomorphic design, the sensor must be realistic, pleasant and natural to touch, and functional on curved surfaces. In the interests of practicality for everyday use, it should also be flexible, washable, not easily worn out, and inexpensive to construct.

Sensitivity: To ensure performance, the sensor should be highly sensitive to even light touches. Its spatial coverage should be close to continuous, so as to miss as few touches as possible.

Accuracy: To be of value, the sensor and gesture recognition engine must together be able to distinguish several relevant touches unavailable to conventional sensors, to a reliability required by the specific application. Here, $\geq 80\%$ is a good initial target.

Efficiency: Excess wiring can interfere with signal accuracy, so a single-circuit sensor will ideally be able to cover a large block of space on the creature body. At the same time, each analog line must yield rich, differentiable data.

4.2 Architecture

A LilyPad Arduino microprocessor samples $I(t)$ at 144 Hz through a connection from one of its six analog inputs to one of our sensor prototype’s two strips of conductive fabric. This strip is also connected through $R_{resistor}$ (Figure 2) to the LilyPad’s ground port. The fur’s second strip of conductive fabric is wired up to power, 5V (Figure 4).

The LilyPad itself receives power from a USB connection to a laptop, on which the Arduino host program stores the sampled data for gesture analysis, currently performed offline (Section 5).

4.3 Construction and Materials

Materials influence sensor performance. Conductive threads must be thin enough to avoid unravelling or fraying from repeated use. The insulating base fabric requires some degree of friction, so threads do not easily pull out. The insulating fur must be pleasant to touch, and convincingly animal-like (Table 1).

We prototyped many thread layouts for the sensor. Here we describe the iteration with the best sensitivity, long-term practicality and tactile attractiveness. This design is made up of a 10cm square patch of the Creature’s fur sewn onto an insulating piece of fabric,




Figure 4: Setup: the fur sensor prototype (1) is connected through one of its two strips of conductive fabric to the LilyPad's 5V power (+). The other conductive strip is wired both to ground (-) through a 1-kOhm resistor (2), and to the LilyPad Arduino (3) through an analog input (4).




Figure 5: Illustration of a 1cm-square cross section of the conductive thread pattern. The pattern consists of one 3cm-long layer of conductive threads sewn in groups of four and singly-sewn resistive threads, and a shorter 1cm layer of conductive threads sewn in groups of three.

with two additional 1cm-wide strips of conductive fabric attached at either end. Conductive and resistive threads are then sewn into the patch in loops, with the ends of each thread forming two hairs in the fur (Figure 5). A resistance change in the circuit thus occurs when two threads from different loops connect or disconnect.

Per square centimeter, there is one long 3cm layer made up of eight conductive thread hairs (a loop of four threads), and two resistive hairs (a loop of a single thread). The shorter 1cm-long layer adds 6 more conductive hairs (a loop of three threads). The two thread lengths are used to help differentiate gestures that occur at different positions in the fur. A gentle *stroke* mostly disrupts the tips of the fur and not its roots, leading to relatively small resistance changes. A *scratch*, on the other hand, activates both layers of conductive threads, resulting in greater resistance changes. In this way position-sensitive gesture types are reflected in the $I(t)$ trajectory.

To reproduce our results, it is not necessary to replicate the precise thread positions in Figure 5, as long as the density of conducting threads/cm is roughly the same. More or fewer thread crossings simply add a small constant to the baseline circuit current level,

Table 1: Conductive Fur Sensor Materials

Material	Specifications
Conductive thread	Silver plated nylon, 117/17, 2 ply Resistance: 30 ohms/10cm Visible length: 3cm, 1cm Where to buy: LessEMF (USA)
Resistive thread	66 Yarn 22+3ply 110 PET Resistance: less than 1000 ohms/10cm Visible length: 3cm Where to buy: LessEMF (USA)
Conductive fabric	Surface resistivity: 0.5 ohm/sq Where to buy: LessEMF (USA)
Resistor	Resistance: 1 kohm Where to buy: Available at most electronics stores
Neoprene fabric	Insulating, low-friction Where to buy: Outdoor Innovations (Canada)
Fur	Insulating, animal-like Length: 3cm Where to buy: Fabricana (Canada)




Figure 6: 2-second $I(t)$ samples by experimenter for *stroke*, *scratch* and *light touch* $I(t)$. Scale is constant; axis values omitted to focus on curve shapes.

which will not affect time series analysis. For example, we have ascertained that the following are effective alterations to the Figure 5 pattern: 1) allowing 2-3 crossings in the threads to occur, 2) insisting none of the threads cross, 3) increasing the resistive-to-conductive thread ratio to 1-1, 4) further increasing the resistive-to-conductive ratio to 5-1.

Many other effective patterns likely exist. The most effective density is the highest before circuit saturation readily occurs - a saturated circuit will result in generally high current flow, masking the effect of any touches to the fur.

5 SENSOR DATA AND GESTURE RECOGNITION

5.1 First Iteration Fur Sensor Data

We begin our analysis with a data set made of 30 2-second samples each of *stroke*, *scratch* and *light touch* gestures, performed by one of the experimenters (Figure 6). We used this data to identify analytical techniques that might work to differentiate gesture, and later evaluated more objectively (Section 5.3).

The resulting time series current curves show how a gesture can be described by its effect on the physical state of the fur. A *stroke* is a fluid motion that pushes the fur against itself, creating smooth, high-flow curves. A *scratch* is a vigorous, disruptive movement, resulting in high-frequency data that bounces between states of high flow and zero current. A *light touch* is gentler and smaller in scope, causing small rifts in the flow (Figure 6).

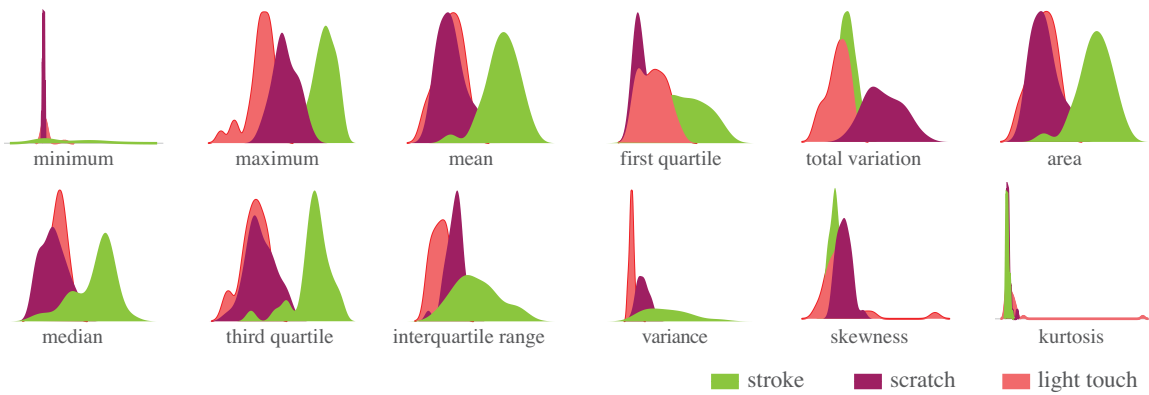


Figure 7: Distribution of features over the 30 2-second gesture samples. Scale is consistent for all data.

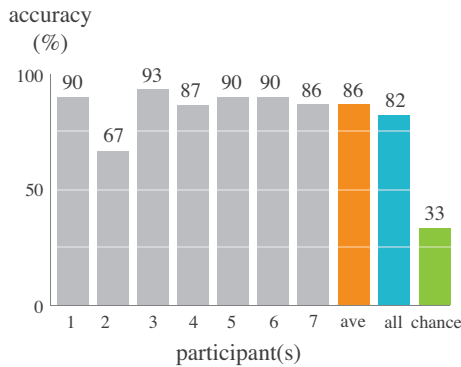


Figure 8: Gesture recognition accuracy for all participants, and for the combined model based on all participant data.

5.2 Gesture Recognition

The data curves in Figure 6 suggest the morphological parameters that an automatic system could use to distinguish different gestures. Frequency is drastically different for *stroke* and *scratch*; it can be described by summing the absolute values of the differences between each consecutive point, resulting in an approximation to the total variation of the sequence. Area under the curve is much larger for *stroke* than for the other two gestures; and maximum in general is large for *stroke*, small for *light touch*, and in the middle for *scratch*.

These statistics are descriptions that could potentially be used to categorize unknown data based on previous domain knowledge. In the context of a machine learning algorithm, this is termed a feature, and the previous domain knowledge comes from the training data [18]. For instance, consider frequency as a feature in our initial training set. Since it is higher for *scratch* and *light touch* than for *stroke*, a touch with high frequency is likely to be either a *scratch* or a *light touch*. Then to further distinguish between *scratch* and *light touch*, we would need another feature.

Feature Selection: Crucial to classification performance, a good feature distinguishes between two different gestures; i.e., its ranges for the two gestures do not overlap (or overlap very little). To extract the most valuable features, we compare model performance for several standard statistics. We also plot the distribution of feature values over the training data to visualize why the model performs better with some than others. We experiment with the following statistics: minimum, maximum, mean, first quartile, total variation, area, median, third quartile, interquartile range, variance,

Table 2: Classification confusion matrix for the combined model trained on *stroke*, *scratch* and *light touch* gesture data ($n=7$).

	Classified as stroke	Classified as scratch	Classified as light touch
Stroke	61	2	7
Scratch	1	60	9
Light touch	7	11	52

skewness, and kurtosis (Figure 7).

From the graphs in Figure 7, we see that skewness seems to have little predictive value, with considerable overlap among the three gestures. Maximum is able to distinguish all three gestures quite well, with relatively little overlap. Total variation is also powerful, because the ranges of *light touch* and *scratch* are well separated. Inclusion of these features did indeed improve model performance.

Therefore, after experimenting with different feature choices and analyzing the distribution graphs, we proceeded with initial feature selections of minimum, maximum, total variation, area, median and variance, and continued to the recognition phase.

Recognition: We uploaded the selected features' data to Weka [17], a software framework for rapid comparison of many standard machine learning algorithms. We tested several learning schemes using leave-one-out cross-validation, including a Bayesian network [2], multilayer perceptron, and logistic regression [9]. We found that a logistic regression model accurately classifies 98% of this initial experimenter-supplied set, miscalculating in only one case.

5.3 Informal Validation of Results

Based on the positive results of our initial data analysis, we conducted an informal evaluation to include a broader data set, and thus access a more realistic view of the sensor's potential for use in gesture recognition.

Seven volunteers (two female, all university students in their twenties) were asked to contribute 30 examples of *stroke*, *scratch* and *light touch* gestures, 10 of each. Prior to performing the gestures, the participants were shown written definitions of each gesture from Yohanan's touch dictionary (as mentioned above, our term "light touch" is the same as his term "tickle," and received his definition) [20].

We trained a linear regression classifier on each of our seven independent data sets, and also used an eighth combined model that was trained on all the data. This was done to test the classifier's ability to personalize to differentiate an individual's gestures, as

well as its potential for modeling variable touch behaviors across a group of people. Leave-one-out cross-validation was again used to evaluate the resulting recognition accuracy.

The model was able to distinguish gesture type for six participants with at least 86% accuracy, and 67% for one participant. The combined model trained on all participant data achieved 82.4% accuracy (chance=33% accuracy) (Figure 8). The combined model's corresponding confusion matrix (Table 2) suggests that *light touch* is the hardest of the three gestures to model accurately, confused primarily with *scratch*. *Stroke* and *scratch* can also be misclassified as *light touch*, but are well distinguished from each other.

6 DISCUSSION

Using our prototype fur sensor and first-pass gesture recognition algorithms, we are able to recognize some broad categories of gesture across several people with > 80% accuracy (our loose initial application-derived design specification). This indicates that further effort to recognize more gesture types, and more touch patterns within a gesture, is well placed. The results for a single participant's data are typically quite accurate, especially so when the person has learned to work with the sensor. This suggests both that the sensor might have the ability to personalize to an individual, and that the human participant also learns from the interaction, much as in the case of a real human/animal relationship.

Our preliminary evaluation has several limitations. It sampled seven individuals, too little to capture the variation among large numbers of users in a real-world setting. Our tests were conducted in a lab environment with the fur immobile on a desk surface, so signals arising from picking up the robot, carrying it around, etc. are not yet taken into account. Perhaps most importantly, we considered only three gestures. While these were selected on the basis of importance to social touch rather than ease of distinction, for a social robot to be truly emotionally intelligent, it must distinguish many more gesture types. Having not attempted recognition of other gestures, we do not know our sensor's limits even in its present prototype form.

7 CONCLUSIONS AND FUTURE WORK

This work began with the hypothesis that hand motion data is an important description of gesture. We have demonstrated a new type of touch sensor that captures hand motion information from changes in conductive fur configurations. We also described an approach to recognizing gestures from this sensor using machine learning techniques on time-series circuit data. We reported the results of an initial informal evaluation of the model's performance on participant gestures, which we believe are promising for future work in this direction.

Our next steps include moving the offline recognition process to a real-time system, integrating time-series force and accelerometer data into the model, and combining multiple patches of conductive fur together to cover larger curved surfaces along an entire robot body, allowing for a combined position-aware system. We will also widen our test set to include more people and more categories of gesture, as well as experiment with different learning techniques, possibly including time-series-specific analysis.

If successful, this work will be integrated into the Haptic Creature to improve gesture recognition. Better gesture recognition in the Creature will provide a better understanding of emotion, which will allow for more intelligent emotional responses. We hope in this way to contribute to the therapeutic power of emotion-aware furry social robots.

8 ACKNOWLEDGEMENTS

We would like to thank the GRAND NSERC Network Center for Excellence, who provided partial support for this work.

REFERENCES

- [1] Russell Bik. Electric field proximity sensor. www.bik.com, May 2011.
- [2] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [3] J. Chang, K. MacLean, and S. Yohanan. Gesture recognition in the haptic creature. In *Generating and Perceiving Tangible Sensations (Proceedings of 2010 International Conference on Haptics)*, 2010.
- [4] Kerstin Dautenhahn. I could be you - the phenomenological dimension of social understanding. *Cybernetics and Systems Journal*, 28(5):417–453, 1997.
- [5] B. R. Duffy. Anthropomorphism and robotics. In *AISB '02 (Proceedings of The Society for the Study of Artificial Intelligence and the Simulation of Behaviour)*, 2002.
- [6] Batya Friedman, Peter H. Kahn Jr., and Jennifer Hagman. Hardware companions?: what online aibo discussion forums reveal about the human-robotic relationship. In *CHI '03*, pages 273–280, 2003.
- [7] K. Goris, J. Saldien, Innes Vanderniepen, and D. Lefebber. The hug-gable robot probot, a multi-disciplinary research platform. In *Eurobot '08*, 2008.
- [8] V. Groom, L. Takayama, P. Ochi, and C. Nass. I am my robot: The impact of robot-building and robot form on operators. In *HRI 2009 (Proceedings of the Human-Robot Interaction Conference), Nominated for the best paper award*, 2009.
- [9] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, second edition, 2009.
- [10] Ambrose Martin. Tactile gloves for autonomous grasping with the nasa/darpa robonaut. In *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, 2004.
- [11] Allison Okamura, Maja J Mataric, and Henrik I. Christensen. Medical and health-care robotics. *IEEE Robotics and Automation Magazine*, pages 27–37, September 2010.
- [12] H. Perner-Wilson. Stroke sensor. Plusea, www.plusea.at, May 2011.
- [13] H. Perner-Wilson and M. Satomi. Diy wearable technology. In *ISEA 2009 Wearable Materialities Panel*, 2009.
- [14] Rosalind Picard. *Affective Computing*. MIT Media Laboratory, 1997.
- [15] T. Shibata, K. Inoue, and R. Irie. Emotional robot for intelligent system: Artificial emotional creature project. In *Proceedings of IIZUKA*, pages 43–48, 2006.
- [16] W. Stiehl and C. Breazeal. Design of a therapeutic robotic companion for relational, affective touch. In *Ro-Man-05 (Proceedings of Fourteenth IEEE Workshop on Robot and Human Interactive Communication)*, pages 408–415, 2005.
- [17] I. H. Witten and E. Frank. *Data mining: Practical machine learning tools and techniques*. Morgan Kaufmann Publishers, second edition, 2005.
- [18] Zhengzheng Xing, Jian Peo, and Eamonn J. Keogh. A brief survey on sequence classification. *SIGKDD Explorations*, 12(1):40–48, 2010.
- [19] S. Yohanan and K. MacLean. The haptic creature project: Social human-robot interaction through affective touch. In *AISB '08 (Proceedings of The Reign of Katz and Dogz, 2nd AISB Symp on the Role of Virtual Creatures in a Computerized Society)*, pages 7–11, 2008.
- [20] Steve Yohanan and Karon E. MacLean. The role of affective touch in human-robot interaction: Human intent and expectations in touching the haptic creature. *International Journal of Social Robotics (SORO), Special Issue on Expectations, Intentions, and Actions*, accepted 2011.