

**Sensing and Recognizing Affective Touch
in a Furry Zoomorphic Object**

by

Anna Flagg

B.Sc., University of Toronto, 2010

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Master of Science

in

THE FACULTY OF GRADUATE STUDIES
(Computer Science)

The University Of British Columbia
(Vancouver)

August 2012

© Anna Flagg, 2012

Abstract

Over the last decade, the surprising fact has emerged that machines can possess therapeutic power. Due to the many healing qualities of touch, one route to such power is in haptic emotional interaction, which in turn requires sophisticated touch sensing and interpretation. We explore the development of affective touch gesture recognition technologies in the context of a furry artificial lap-pet, with the ultimate goal of creating therapeutic interactions by sensing human emotion through touch. We design, construct, and evaluate a low-cost, low-tech furry lap creature prototype equipped with 2 types of touch-sensing hardware. The first of these hardware types is our own design for a new type of touch sensor built with conductive fur, invented and developed as part of this research. The second is an existing design for a piezoresistive fabric pressure sensor, adapted to three dimensions for our robot-body context. Combining features extracted from the time-series data output of these sensors, we perform machine learning analysis to recognize touch gestures. In a study of 16 participants and 9 key affective gestures, our model averages 94% gesture recognition accuracy when trained on individuals, and 86% accuracy when applied generally across the entire set of participants. The model can also recognize who out of the 16 participants is touching the prototype with an accuracy of 79%. These results promise a new generation of emotionally intelligent machines, enabled by affective touch gesture recognition.

Preface

The study described in this thesis was conducted under the approval of the University of British Columbia’s Behavioural Research Ethics Board (BREB), certificate B01-0470: Low Attention and Affective Communication Using Haptic Interfaces.

The conductive fur sensor described in Chapter 3 was originally designed in collaboration with Diane Tam, a fellow graduate student at UBC, under the supervision of UBC professor Karon MacLean. The subsequent machine learning analysis of sensor data was performed by myself and my father Robert Flagg, and I did further design and development of the sensor under Karon MacLean’s supervision. Excluding Section 3.6, a version of Chapter 3 has been published: Flagg, A., Tam, D., MacLean, K., and Flagg, R. *Conductive Fur Sensing for a Gesture-Aware Furry Robot*, in IEEE Haptics Symposium (2012). Section 3.6 describes additional unpublished work done to demonstrate the work at the Haptics Symposium conference.

I constructed early prototypes of the piezoresistive fabric pressure sensor described in Chapter 4 in collaboration with UC Berkeley professor Adrian Freed, according to his and UC Berkeley graduate student Andrew Schmeder’s original design. I later adapted this design to three dimensions using materials supplied by Freed. Under the supervision of Karon MacLean, I then combined this sensor with other hardware into the final prototype as described, ran the study to collect human gesture data, and performed machine learning analysis. A version of Chapter 4 will be submitted for publication: Flagg, A., and MacLean, K. *Affective Touch Recognition for a Furry Therapeutic Machine*. Manuscript in preparation.

Table of Contents

Abstract	ii
Preface	iii
Table of Contents	iv
List of Tables	vii
List of Figures	viii
Acknowledgments	xi
Dedication	xii
1 Introduction	1
1.1 The Haptic Creature	2
1.2 Approach and Contributions	4
1.3 Thesis Organization	5
2 Related Work	7
2.1 Haptic Affective Robots	7
2.2 Touch Sensing	8
2.3 Hand Motion Sensing	9
2.4 Pressure Sensing with Piezoresistive Fabric	11
2.5 Affective Touch Gesture Recognition	12

3	The Smart Fur Sensor	14
3.1	Origins	14
3.2	Objectives	16
3.3	Approach	16
3.3.1	Advances in Physical Design and Data Collection	17
3.3.2	Architecture	19
3.3.3	Construction and Materials	20
3.4	Analysis	22
3.4.1	First Iteration Fur Sensor Data	22
3.4.2	Gesture Recognition	23
3.4.3	Informal Evaluation	25
3.4.4	Results	26
3.5	Discussion	27
3.6	Demo Prototype	28
3.7	Next Steps	29
4	A Creature Object: Combining Smart Fur and Fabric Pressure Sensing	30
4.1	Objectives and Approach	31
4.2	Integrated Sensor Construction and Electronics	32
4.2.1	Fur Sensor Improvements	32
4.2.2	Piezoresistive Fabric Pressure Sensor	32
4.2.3	Construction and Adaptation	35
4.2.4	Sensor Fusion	35
4.2.5	Construction Costs	37
4.3	Gesture Evaluation and Analysis	38
4.3.1	Gesture Data Collection	38
4.3.2	Machine Learning Analysis	39
4.3.3	Results	40
4.4	Discussion	45
4.5	Future Work	48
5	Conclusions and Future Work	51

Bibliography	53
Appendices	55
A Model Parameters	56
B Study Documents	58

List of Tables

Table 3.1	Conductive fur sensor materials	20
Table 3.2	Classification confusion matrix for the combined model trained on <i>stroke</i> , <i>scratch</i> and <i>tickle</i> gesture data (n=7).	27
Table 4.1	Estimated manufacture costs for a 20cm x 15cm gesture sensing prototype	37

List of Figures

Figure 1.1 The result of this research: a low-cost, low-tech zoomorphic touch and gesture-sensing prototype. 2

Figure 1.2 The Haptic Creature [28]. 3

Figure 2.1 MIT Media Lab’s Huggable [25]. 8

Figure 2.2 Perner-Wilson and Satomi’s conductive thread stroke sensor [13]. 10

Figure 2.3 Variants of Schmeder and Freed’s piezoresistive fabric pressure sensor [20]. 11

Figure 3.1 Our conductive “smart fur” sensor. 15

Figure 3.2 Circuit for our design. Touches change the fur configuration and consequently the net fur resistance. The resulting fluctuating voltage is sampled at 144 Hz. 17

Figure 3.3 Perner-Wilson and Satomi’s conductive thread stroke sensor [13] (left), our conductive fur touch and gesture sensor (right). . . . 18

Figure 3.4 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). 19

Figure 3.5	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 3.	21
Figure 3.6	2-second voltage samples by experimenter for <i>stroke</i> , <i>scratch</i> and <i>tickle</i> . Scale is constant; axis values omitted to focus on curve shapes.	23
Figure 3.7	Distribution of features over the 30 2-second gesture samples. Scale is consistent for all data.	24
Figure 3.8	Logistic regression gesture recognition accuracy for all participants, and for the combined model based on all participant data.	26
Figure 3.9	Demo prototype responding to a gentle stroke (top left), a playful scratch (top right), and a breath (bottom).	28
Figure 4.1	Hand stroking the touch-sensing creature prototype. Prototype is equipped with conductive fur and piezoresistive fabric pressure sensing, and sized ~15 x 20 cm.	31
Figure 4.2	Adaptation of Schmeder and Freed’s piezoresistive fabric pressure and position sensor [20].	33
Figure 4.3	Construction summary of the creature prototype hardware. Top down from left to right: (a) creature body consisting of styrofoam and plastic “skeleton” and “skin”, (b) first layer of piezoresistive fabric, (c) 3 mesh standoff layers, (d) second layer of piezoresistive fabric, and (e) touch-sensing conductive fur on top.	34
Figure 4.4	Sample fur sensor, pressure, x, and y data curves for stroke, scratch and pat gestures. (Axis scales excluded for clarity.) . .	36

Figure 4.5	Gesture recognition results for random forests, neural networks, logistic regression, and Bayesian networks. Models were trained and tested on each of 16 study participants individually, using 100-fold cross-validation. (A graph range of 80-100% is used to maximize detail.)	41
Figure 4.6	Gesture recognition results for random forests, neural networks, logistic regression, and Bayesian networks. Models were trained and tested on a combined set of all 16 participants, using 100-fold cross-validation. (50-100% displayed graph range).	42
Figure 4.7	Hinton diagram of the confusion matrix corresponding to the random forests gesture recognition model, evaluated using 100-fold cross-validation on the combined data set of 16 study participants. Size of grey squares represents classification: for example, the first row shows how many examples of “stroke” are recognized as <i>stroke</i> , how many are mislabeled as <i>scratch</i> , and so on for all gesture types.	43
Figure 4.8	Contributions of individual sensor channels to gesture classification performance of random forests model on combined 16-participant data set, evaluated using 100-fold cross validation. Upper curve shows performance when using features from all data channels <i>except</i> conductive fur, then <i>except</i> x position, <i>except</i> y position, <i>except</i> combined x and y position, and <i>except</i> pressure. Lower curve shows accuracy when using features from <i>only</i> conductive fur, then <i>only</i> x position, <i>only</i> y position, <i>only</i> combined x and y position, and <i>only</i> pressure. (50-100% displayed graph range).	44
Figure 4.9	Results of evaluating a random forests model’s ability to recognize people from their touch data. Model was trained and tested on <i>stroke</i> data for all participants, <i>scratch</i> data for all participants, etc for all gestures. Model was also evaluated on combined set of participant data for all gestures. All tests used 100-fold cross-validation.	45

Acknowledgments

I was very lucky to be supervised in this project by Dr. Karon MacLean. Working with Karon was a valuable experience, from which I took away a great deal, and I could not have asked for a better supervisor. I continue to look up to her as a great teacher, and an inspiring mentor.

I'd like to thank my second reader Dr. Tamara Munzner, whose valuable feedback, ideas, and encouragement made this thesis much better than it would have been. Thanks too to Hasti Seifi, my student reader, who also provided much needed feedback, support, and friendship.

My friends Diane Tam, Oliver Schneider, John Harris, and Mona Haraty supported me throughout this process, and made everything more fun. I'd especially like to mention two of my best friends, Arthur Lo and Shathel Haddad, who can always make me laugh. I don't think I could have made it through the last year without them.

I'd also like to acknowledge Dr. Nando de Freitas, Dr. Machiel Van der Loos, and Kerem Altun, all of whom provided valuable support and ideas that contributed to this research. In particular, Dr. Adrian Freed was very generous with his time and knowledge, as well as his building materials. His help greatly advanced my progress.

Many thanks to the people at the GRAND NSERC Network Center for Excellence, who partially funded this work.

Finally, words come short to describe all that my family has given me. My parents Bob Flagg and Geeta Ramani, my sister Shanti Flagg, my aunt Usha Ramani, and my boyfriend Moiz Syed: their love and support continue to get me through each day.

Dedication

To my dad, *Bob*
who taught me that math is better than a boyfriend
and to *Moiz*
who disagrees.

Chapter 1

Introduction

I just want to be your teddy bear.
— Elvis Presley (1957)

Machines are traditionally thought of as purely rational, unemotional things. However Picard [14] argues that due to the complex interplay of thought and emotion in the brain, truly natural communication between humans and computers can not exist without considering emotion in design. Due to the emotional nature of touch, this idea is especially applicable within the design of haptic systems.

Although affective computing has gained traction in the HCI community, at the moment, affective *touch* is largely ignored. One reason for this gap is that creating touch experiences involves building physical hardware and feedback, a more complicated design and implementation task than supporting screen interaction. More fundamentally, the idea of machines providing even a limited level of the type of support that comes from emotional touch is widely believed to be an impossibility, even in the field of affective computing.

However, as this thesis will demonstrate, we do not agree. We are motivated by the fact that affective touch is a crucial part of human development and well-being, especially for the young, the elderly, and the ill. Since natural forms of touch therapy, such as interaction with trained animals, are often unavailable in hospitals, homes of disadvantaged individuals, and other crucial situations, an artificial system capable of providing even partial support would have many valuable applications. Empowering machines with affective touch could lead to a whole new




Figure 1.1: The result of this research: a low-cost, low-tech zoomorphic touch and gesture-sensing prototype.

range of potential uses in therapy, rehabilitation, education, treatment of cognitive disorders, and assistance for people with special needs. As this thesis will argue, we believe these prospects may in fact be possible.

1.1 The Haptic Creature

It has long been known that pets can have a positive effect on their owner's emotions. In particular, petting an animal can release endorphins and lower blood pressure. Could this phenomenon ever occur in interactions between people and artificial systems?

Exploring this direction, our group has developed the Haptic Creature [32], a furry lap-sized social robot that communicates with the world through touch (Figure 1.2). Sensitive to touch and movement, the Creature can modulate its ear stiffness, breathing rate, and purring patterns. Focusing on employing a human/animal relationship analogy, the Creature system aims to create comforting experiences through touch-based interactions. Specifically, the goal is to sense human emotion and respond intelligently, particularly in the case of an anxious human user in need



Figure 1.2: The Haptic Creature [28].

of a calming interaction.

In a recent study, our group found that such calming interactions are indeed possible. During the study, the Creature was equipped with biometric sensors measuring galvanic skin response, heart rate, heart rate variability, and respiration rate as measures of anxiety. Using these metrics, the Creature was found capable of reducing anxiety in individuals who experience its active breathing [21].

This important result motivates further development of the Creature's therapeutic capabilities. In particular, the current method of sensing emotion through wearable biometric sensors is too intrusive to be acceptable in the long run. The next question therefore becomes: how can we sense emotion less intrusively?

Posture, speech, touch, voice prosody, and physiological measures have been explored as possible paths to emotion sensing [17]. While these are promising directions, modelling emotion accurately remains very much an open problem. Given the emotional nature of the human/pet relationship, we suggest that the zoomorphic and highly emotive, touchable form of the Creature encourages emo-

tional expression through touch. In particular, we propose that the way a person touches the Haptic Creature is a window into her emotional state.

This assertion is backed up by a recent study, in which our group found that different types of touch interaction with the Creature can indeed be associated with different human emotions [31]. Based on this result, the focus of our further research is investigating the Creature as a platform for sensing the emotional content of touch. Specifically, our goal is to sense human emotional state through analysis of touch interaction with the Creature.

This endeavour consists of two major research areas: first, the sensing and recognition of touch gestures, and second, the inference of emotional state from those touch gestures. This thesis is concerned with the first of these steps: touch sensing, and the recognition of affective touch gestures in the context of a furry robotic creature.

1.2 Approach and Contributions

In this thesis, we present the combination of physical design work and artificial intelligence methods that enable our new design for a low-cost, low-tech touch sensing system to achieve recognition rates competitive with existing gesture recognition systems that use more costly sensing technologies (Section 4.2.5). Our design consists of a small (~ 20 cm long) furry sensor prototype, in shape resembling a small lap animal, with two types of touch sensors embedded (Figure 1.1). We used this creature-like prototype as an interface to collect human touch data for 9 key emotional gestures, which we then classified using machine learning analysis.

Our contributions include:

- 1) design and development of a new type of touch sensor built with conductive fur, the physical principle of which is inspired by Perner-Wilson’s thread stroke sensor [13]
- 2) data collection for three sample touch gestures in an initial, informal 7-participant study, and subsequent machine learning analysis to classify gesture, evaluating viability of the sensor, and advisability of further development
- 3) adaptation of Schmeder and Freed’s piezoresistive fabric pressure sensor [20] to arbitrary non-flat surfaces, done to cover our three-dimensional robot body con-

text

4) creation of furry animal-like prototype, and integration of the above two low-cost sensing technologies, outputting synchronized touch position, pressure, and fur sensor data

5) data collection for 9 key emotional gestures in a 16-participant study, and subsequent machine learning analysis to classify gesture, evaluating above system as an approach to affective gesture recognition in a Creature-like context

6) computationally inexpensive (realtime) model, which predicts gesture type with an average accuracy of 94% for a given individual, and 86% when generalized across the entire set of participants, and recognizes who out of the 16 participants is touching the prototype with 79% accuracy.

We believe this work will provide the Haptic Creature with a fundamental basis model for gesture recognition. We hope our contributions can pave the way for a more emotionally sophisticated system, capable of providing therapy to children with anxiety disorder, elderly adults with dementia, and others in need.

1.3 Thesis Organization

The next chapter summarizes relevant previous works in touch sensing, affective robotics, and machine learning. The focus is an overview of the status quo in affective touch recognition, and the specific problems and challenges in current methods that have motivated the directions this research has taken (Chapter 2).

Among these directions is the invention of our conductive “smart fur” sensor, a new type of touch and gesture sensor designed to measure hand motion information unavailable to conventional pressure sensors (the norm in touch systems). Chapter 3 covers the smart fur sensor, its origins in creative thread-based circuit designs by Perner-Wilson and Satomi [13], our iterative design and construction process, and the physical and analytical advances we made to create a sensor capable of contributing to touch gesture recognition. In this chapter we also describe our informal evaluation, done early in the larger research process to establish the potential of the smart fur sensor, and justify its further development as a complementary technology to the current touch sensing methods in the Haptic Creature.

Chapter 4 deals with the process of combining the smart fur sensor with pres-

sure sensing, using an adaptation of Schmeder and Freed's piezoresistive fabric pressure sensor [20]. This chapter covers the operating principle and construction details of this design, and how we adapted it to three dimensions to cover a furry lap-pet object, integrated to produce synchronized touch pressure, position, and fur sensor data streams. This chapter also describes our 16-participant study, during which we collected data for 9 gestures key to emotional communication, and the gesture recognition model resulting from analyzing this data using machine learning methods.

In the final chapter, we summarize the contributions of this research, and what our results might mean for the future of the Haptic Creature project, and to future haptic affective systems.

Chapter 2

Related Work

This project is at the intersection of touch sensing, affective robotics, and machine learning. We discuss selected previous works relevant to our goals, and how our work differs and builds upon them.

2.1 Haptic Affective Robots

Huggable, PARO, Aibo and Probo are all touch-sensitive affective social robots relevant to our work [8, 10, 22, 25]. Of these, Huggable [25], a furry robotic bear, has the most advanced touch sensing: its initial recognition model identifies 9 touch gestures with data from the robot's full-body sensitive skin, which includes a wide range of sensors (Figure 2.1). PARO [22] is the famous interactive robotic seal that recognizes patterns in its environment, including common verbal phrases, and it has a long-term memory of owner touch behaviour. Specifically, it differentiates between being stroked and hit, and tries to amend its own behaviour accordingly, repeating actions that have been rewarded with stroking, and avoiding actions that have resulted in hitting. Robot dog Aibo [8] grows gradually from a puppy personality to a mature dog over time, and is able to connect with people and its environment in many ways, including recognizing its owner, learning tricks, and locating its charging station. Aibo has touch sensors on its head, chin and back, and responds to touch interaction based on location of touch. Probo [10] is an elephant-like social robot equipped with a large variety of sensors. It expresses 7




Figure 2.1: MIT Media Lab's Huggable [25].

emotions by changing its facial expression, and is used to ease anxiety in hospitalized children. Probo focuses its touch sensing on recognizing whether it is being hugged, scratched, or hurt.

Some of these projects share goals similar to those of the Haptic Creature project - namely, recognition of human emotion, and appropriate responses to provide a therapeutic effect. But none has yet solved the complex problem of accurate emotion recognition, nor do any of them go beyond the most rudimentary processing in terms of sensing emotional touch. It is our goal to contribute to touch gesture recognition in the Haptic Creature, in the hopes of enabling this system to recognize emotion from human touch behaviours.

2.2 Touch Sensing

Most affective systems focus on sensing touch through force, such as with Force Sensitive Resistors (FSRs) [9] or Quantum Tunnelling Composites (QTCs) [16]. PARO and Aibo use FSRs alone to identify touch, and hitherto the Haptic Creature

has as well [8, 22, 29]. Huggable also uses FSRs, in conjunction with temperature sensors and capacitive sensors [25].

These are promising directions, but they are still in early stages of gesture recognition, and none is likely to individually have the needed sensing scope. Huggable contains the most advanced recognition engine, but it uses over 1500 high-tech sensors, relies partly on location of touch to define gesture, and does not have complete recognition capabilities. The capacitive sensors are also quite expensive, and may be vulnerable to interference. 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 light touches, but our group and others have found that they suffer from intractable nonlinearities, and they are also not easily available at this time.

To attempt to ease these problems of expense, hardware complexity, degraded performance on curved surfaces, lack of continuous sensor coverage, and insensitivity to light touches, we propose the combination of two low-cost sensor types: a conductive fur sensor, and a piezoresistive fabric pressure sensor. The next two sections address works related to these two sensors.

2.3 Hand Motion Sensing

Current touch sensing technologies rely largely on force detection alone, which handicaps the system's sensitivity to light touches, and differentiation of gestures that are distinct but involve similar hand pressure. For instance, a firm stroke and a scratch could involve similar pressure, and it is the subtleties of hand position and motion over time that defines each. A firm stroke involves the flat of the hand moving smoothly and repeatedly along the skin to exert force, usually in one direction. In a scratch, fingernails ruffle the fur back and forth with high-pressure contact against its length, especially at the roots.

Sensing force alone would likely not provide the best differentiation between a firm stroke and a scratch, and there are many such examples. Thus, tasked with improving the Haptic Creature's sensing capabilities using inexpensive, computationally-



Figure 2.2: Perner-Wilson and Satomi’s conductive thread stroke sensor [13].

efficient and hardware-light methods, we explore above-surface hand motion information as an input to touch sensing.

Our physical design is inspired by Perner-Wilson and Satomi’s concept for a low-tech stroke sensor [13]: a circuit made up of conductive threads placed vertically like fur, that senses when a long stroking motion is performed (Figure 2.2).

We applied the idea in Perner-Wilson and Satomi’s conductive thread stroke sensor to a new, considerably more sensitive conductive fur touch and gesture sensor. This new sensor includes physical advances as well as an added analytical component, designed to recognize and differentiate gesture types by measuring hand motion data. Chapter 3 contains full details of our design process beginning with Perner-Wilson and Satomi’s seed idea, and our subsequent development, analysis and evaluation.




Figure 2.3: Variants of Schmeder and Freed's piezoresistive fabric pressure sensor [20].

2.4 Pressure Sensing with Piezoresistive Fabric

As mentioned above, at least low-resolution pressure information is important for gesture recognition, but there are many drawbacks to existing pressure sensing technologies in affective robots. These include expense (accompanying higher-than-needed sensing quality), hardware complexity, maintaining performance on curved surfaces, and providing continuous sensor coverage. Piezoresistive fabrics address many of these problems.

Piezoresistive materials are special types of semiconductors that change their electrical resistivity when a pressure is applied, and can thus be used as pressure sensors in artificial systems. In particular, many low-cost pressure sensing designs have been developed involving flexible piezoresistive fabrics, as discussed in detail in Freed's work [7, 19]. Since these fabrics are inexpensive, flexible, and can be easily wrapped and sewn around irregular three-dimensional shapes, this mate-

rial is a promising direction for pressure sensing in the Haptic Creature and many related applications.

Previous piezoresistive fabric pressure sensing designs have been used largely in medical sensing, robotics, and, more recently, electronic musical instrument design [7, 19]. But we have not yet encountered their use for affective touch gesture recognition. Future work will be required to optimize functionality and performance in these designs, both absolute and relative to each other, for touch gesture recognition. As a first step, in this work we selected one of the simpler existing designs created by Schmeder and Freed [20] (Figure 2.3), adapted it to our three-dimensional robot body shape, and evaluated it for use in affective gesture recognition.

Chapter 4 details the construction, adaptation and integration process, as well as the resulting gesture classification results.

2.5 Affective Touch Gesture Recognition

The use of machine learning for touch gesture recognition in affective systems is in early stages. Of the works we have mentioned, only Huggable and the Haptic Creature projects involve modelling more than a few basic gesture types.

The Huggable team has experimented with supervised neural networks using feature-based sensor data, and reports a 61.6% true positive recognition rate, and 96.5% true negative recognition rate, averaging to 79% accuracy [25]. This report includes 8 of the original 9 gesture types, due to some believed technical difficulties with the remaining gesture type (which had resulted in very poor accuracy). The Haptic Creature group has made use of features with an eventual probabilistic structure in mind [5], and reports 77% recognition accuracy for a set of four gestures.

We also took a feature-based approach, but rather than defining features in terms of the relationships among multiple sensors read instantaneously, we extracted features from time-series sensor curves [4]. Reading synchronous data from the fur sensor and the fabric pressure sensor, we extracted several standard sequence statistics. In this way we incorporated the time-dependent nature of gestures into our model, giving it a “memory,” which we hypothesize is key to

performance. These sequence features were fed into Weka [27], an open source framework supporting practical application of machine learning algorithms. Weka allows us to quickly compare performance of many different algorithms on our data.

Making use of standard sequence features is one natural first approach to classifying gesture [4]. Chapter 4 discusses our results using this approach, which we believe significant both because formal touch gesture recognition rates are rare in the current literature, and because our results are promising relative to previous reports.

However, these results should also be considered in light of other possibly significant factors, such as gesture selection. There is considerable overlap between our gesture set and the gesture sets used in previous works, but they are not identical. It is likely that certain gestures are easier to recognize than others, and that gesture sets that are more diverse are easier to differentiate than gesture sets containing many very similar gesture types. It is also likely that performance will degrade considerably on real-world tests, outside of a controlled lab environment, as people relax into more natural, in-home behaviour. For instance, gestures may be more difficult to recognize when people are holding the prototype in varying ways, carrying it around, placing it in different locations, etc.

These issues and more are discussed in Chapter 4, where we will also touch on what our model says about the data, what should be done to improve performance, and how our results inform design as we move towards a more sophisticated recognition engine in future.

Chapter 3

The Smart Fur Sensor¹

In Chapter 2 we discussed some of the problems involved in relying solely on pressure information to recognize touch. In an attempt to ease these problems, we propose augmenting pressure sensing with information from a new type of conductive "smart fur" touch sensor, designed to quantify hand motion (Figure 3.1).

Beginning with its origins in creative thread-based circuitry by Perner-Wilson and Satomi, this chapter details the design of the conductive fur sensor, and our construction and development process, including physical improvements and an added analytical model. Following a small user evaluation, the chapter ends with our conclusions about the sensor's validity as a contributing hardware for gesture recognition in the Haptic Creature, an important milestone that informs the further directions of this research.

3.1 Origins

Perner-Wilson and Satomi proposed several low-tech wearable circuit designs built with conductive threads and fabrics. One of these designs is the conductive thread stroke sensor [13] mentioned in Section 2.3: a circuit made up of conductive threads sewn vertically like fur into insulating fabric, and connected to an LED for output (Figure 2.2). When stroked, the threads brush against each other, closing the circuit and providing power to turn on the light. When the hand moves

¹A version of Chapter 3 has been published: Flagg, A., Tam, D., MacLean, K., and Flagg, R. *Conductive Fur Sensing for a Gesture-Aware Furry Robot*, in IEEE Haptics Symposium (2012).



Figure 3.1: Our conductive “smart fur” sensor.

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. However this is not an inherent limitation of the concept, as we show in our extension. Perner-Wilson and Satomi’s sensor is also responsive to one specific 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 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. Furthermore, when only binary open/closed information is used, sensor output is not rich enough to differentiate the subtleties of different types of gestures; such analysis was indeed not part of the original presentation.

We made several physical and analytical advances to combat these challenges, discussed in Section 3.3.1. First, however, we cover the objectives and basic operating principle behind the conductive fur sensor design.

3.2 Objectives

The conductive fur sensor was 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 lap-pet. 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 low-cost, 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, 80%+ is a good initial target (see Section 4.4 for the motivation behind this choice).

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.

These objectives shaped our approach in designing the conductive fur sensor.

3.3 Approach

Our design 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 these physical changes in the fur for visibility into the gesture space.

To this end, we sewed the vertical conductive threads from Perner-Wilson and




Figure 3.2: Circuit for our design. Touches change the fur configuration and consequently the net fur resistance. The resulting fluctuating voltage is sampled at 144 Hz.

Satomi’s stroke sensor design [13] into thick animal-like fur, creating a weak circuit. When someone touches the fur, the hand motion disturbs the configuration of the conductive threads inside, causing different numbers of threads to connect and disconnect. These changing electrical connections change the overall resistance in the circuit, which is reflected in voltage level changes across the fur (Figure 3.2). These voltage changes can be analyzed over time to classify gesture.

3.3.1 Advances in Physical Design and Data Collection

To create our new conductive fur sensor, we made the following changes to Perner-Wilson and Satomi’s original stroke sensor idea (Figure 3.3):

Thick, animal-like fur: Embedding the conductive threads into a sample of the thick fur used in the Haptic Creature (rather than Perner-Wilson and Satomi’s individual sensor threads) maintains realism and visual, tactile attractiveness.

Sampling voltage over time: Rather than sampling a single *stroke* or *no stroke* state, we sample voltage over time. Voltage fluctuates according to connections



Figure 3.3: Perner-Wilson and Satomi’s conductive thread stroke sensor [13] (left), our conductive fur touch and gesture sensor (right).

between the threads, i.e., changes in the fur’s physical state that occur during a touch.

Densely placed conductive threads: Using changing voltage levels over time allowed us to position the conductive 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 iterated on thread patterns to increase sensitivity without saturating the circuit, including combining threads with low resistance (which we refer to as *conductive threads*) with threads with high resistance (*resistive threads*). Adding resistive threads helps prevent a saturated voltage signal, which 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.

We now discuss details of the sensor architecture and construction.




Figure 3.4: 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).

3.3.2 Architecture

A LilyPad Arduino microprocessor [1] samples voltage 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 a resistor to the LilyPad's ground port. The fur's second strip of conductive fabric is wired up to 5V power (Figure 3.4).

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 3.4).

Table 3.1: Conductive fur sensor materials

<i>Material</i>	<i>Specifications</i>
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: roughly 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)

3.3.3 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.

We prototyped many thread layouts for the sensor. Here we describe the iteration with the best sensitivity, long-term practicality and tactile attractiveness. Table 4.1 lists the materials in this iteration, which is made up of a 10cm square patch of the Creature’s fur sewn onto an insulating piece of fabric, 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 3.5). 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 3 threads). The two thread lengths are used to help differentiate gestures that occur at different positions in the fur. A *stroke* mostly disrupts the tips of the fur and




Figure 3.5: 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 3.

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 voltage trajectory.

To reproduce our results, it is not necessary to replicate the precise thread positions in Figure 3.5, 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, which will not affect time series analysis. For example, we have found through experimentation that the following are effective alterations to the pattern in the Figure 3.5: 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.

3.4 Analysis

In evaluating the conductive fur sensor, our goal was to measure its ability to augment existing sensing technologies in a useful way: specifically, are there classes of gesture that the fur can model more accurately than traditional pressure sensors?

As a proof of feasibility, we focused on identifying a few key gestures from Yohanan's touch dictionary [31], leaving inference of the emotional content to later work.

We selected *stroke*, *scratch* and *tickle* on the basis of crucial affective content, inadequate differentiation by existing sensor technology, and a potentially good match to the fur-based sensor.

These gestures are defined as follows by Yohanan [31]:

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

scratch: rubbing the fur with one's fingernails,

tickle: touching the fur with light finger movements

Our subsequent analysis is made up of two parts: 1) an informal experimentation period where we use experimenter data to suggest potential classification techniques for the above 3 gestures (Figure 3.6), and 2) a small study involving 7 participants, carried out to evaluate our model. We briefly describe both in the next sections.

3.4.1 First Iteration Fur Sensor Data

Our analysis began with a data set of 30 2-second samples each of *stroke*, *scratch* and *tickle* gestures, performed by one of the experimenters (Figure 3.6). We used this data to identify analytical techniques that might work to differentiate gesture, and later evaluate more objectively (Section 3.4.3).

The resulting time series voltage 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

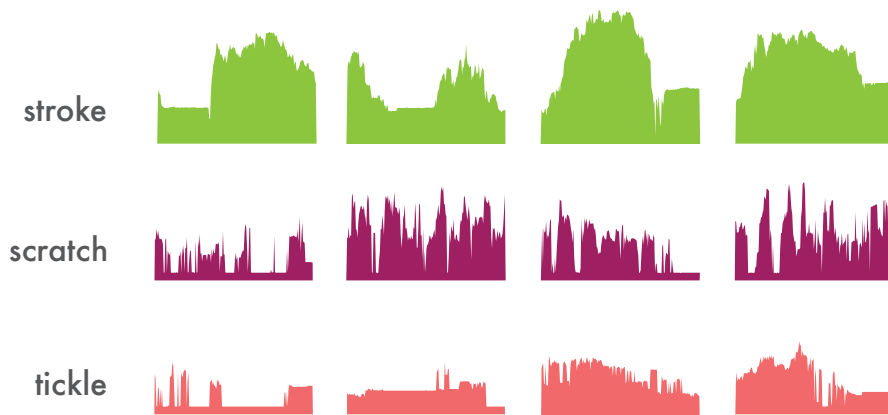


Figure 3.6: 2-second voltage samples by experimenter for *stroke*, *scratch* and *tickle*. Scale is constant; axis values omitted to focus on curve shapes.

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 *tickle* is gentler and smaller in scope, causing small rifts in the flow (Figure 3.6).

3.4.2 Gesture Recognition

The data curves in Figure 3.6 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 *tickle*, 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 [11]. For instance, consider frequency as a feature in our initial training set. Since it is higher for *scratch* and *tickle* than for *stroke*, a




Figure 3.7: Distribution of features over the 30 2-second gesture samples. Scale is consistent for all data.

touch with high frequency is likely to be either a *scratch* or a *tickle*. Then to further distinguish between *scratch* and *tickle*, 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 compared model performance for several standard statistics. We also plotted the distribution of feature values over the training data to visualize why the model performs better with some than others. We experimented with the following statistics: minimum, maximum, mean, first quartile, total variation, area, median, third quartile, interquartile range, variance, skewness, and kurtosis (Figure 3.7).

From the graphs in Figure 3.7, we see that maximum is able to distinguish all 3 gestures quite well, with relatively little overlap. Total variation is also powerful, because the ranges of *tickle* and *scratch* are well separated. Inclusion of these features does indeed improve model performance.

Skewness seems to have little predictive value, with considerable overlap among the 3 gestures. Kurtosis ranges are similarly overlapping. First and third quartiles contain significant overlap for *scratch* and *tickle*, and although both somewhat useful for distinguishing *stroke* from the other two gestures, neither of them does this job as well as maximum, median, or area, so they are not necessary. And mean and area are similar to each other, so we need include only one of them.

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 feature data to Weka [27], 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 Bayes network [3], neural network, and logistic regression [11]. We found that a logistic regression model accurately classifies 98% of this initial experimenter-supplied set, miscalculating in only one case. See Appendix A for detailed model parameters.

3.4.3 Informal Evaluation

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 *tickle* gestures, 10 of each. Prior to performing the gestures, the participants were shown written definitions of each gesture from Yohanan’s touch dictionary [31].

We trained a logistic regression classifier on each of our seven independent data sets, and also trained an eighth combined model on all the data. This was done to test the classifier’s ability to personalize to an individual’s gestures, as well as its potential for modelling variable touch behaviours across a group of people. Leave-one-out cross-validation was again used to evaluate the resulting recognition accuracy. Appendix A contains model parameter settings.




Figure 3.8: Logistic regression gesture recognition accuracy for all participants, and for the combined model based on all participant data.

3.4.4 Results

The resulting individual models are able to distinguish gesture type for each of the six participants with at least 86% accuracy, and 67% for one participant. The combined model trained on all participant data achieves 82.4% accuracy (chance=33% accuracy) (Figure 3.8).

Table 3.2: Classification confusion matrix for the combined model trained on *stroke*, *scratch* and *tickle* gesture data (n=7).

	Classified as <i>stroke</i>	Classified as <i>scratch</i>	Classified as <i>tickle</i>
<i>Stroke</i>	61	2	7
<i>Scratch</i>	1	60	9
<i>Tickle</i>	7	11	52

The combined model’s corresponding confusion matrix (Table 3.2) suggests that *tickle* is the hardest of the 3 gestures to model accurately, confused primarily with *scratch*. *Stroke* and *scratch* can also be misclassified as *tickle*, but are well distinguished from each other.

3.5 Discussion

Using our conductive 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 result indicates that further effort to utilize the fur sensor to aid in recognition of more gesture types 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 high performance 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 had several limitations. It sampled seven individuals, too little to capture the variation among large numbers of people 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. were not taken into account.

Perhaps most importantly, we considered only 3 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




Figure 3.9: Demo prototype responding to a gentle stroke (top left), a playful scratch (top right), and a breath (bottom).

gesture types. Having not attempted recognition of other gestures, we do not know our sensor’s limits even in its present prototype form.

3.6 Demo Prototype²

Based on the above results, we created a simple actuated prototype to demonstrate the sensing and recognition capabilities of our combined fur sensor and model system (Figure 3.9). Built as an addition to the version of this chapter that was published as a paper, the goal of this prototype was to demonstrate to a conference audience that the fur sensor can tell the difference between two types of gestures: a

²Section 3.6 describes unpublished additional work done for conference demonstration purposes.

soft, gentle stroke, and a more energetic, playful scratch. Additionally, we wished to convey that this recognition can take place in real time, with a high degree of sensitivity and accuracy, without being training on strangers.

The prototype features a hand-sized rectangular patch of conductive fur connected to a servo motor, and controlled by an Arduino Uno microcontroller [2]. Using a simplified version of the gesture recognition model presented above, the device expresses itself in two ways: when a gentle touch is sensed, the servo motor moves the fur calmly and slowly up and down, similar to a restful breathing. When the more playful touch is sensed, the servo moves the fur quickly and eagerly in smaller bursts of excitement. A short demonstration video is included with this thesis in the supplementary materials.

With the help of a zoomed-in streaming webcam projection, a successful live demo of this prototype was given at the 2012 Haptics Symposium to an audience of 300.

3.7 Next Steps

This chapter 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 physical changes in conductive fur configurations. We described our approach to recognizing gestures from this sensor using machine learning techniques on time-series circuit data. Finally, 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.

Given these promising results, we now turn our attention to carrying out a more formal evaluation of the sensor for use in gesture recognition, as well as its further development and integration into a more complete system. Specifically, this process will involve combining the conductive fur sensor with pressure sensing, and integrating both into a three-dimensional creature-like object, a truer representation of our eventual Haptic Creature application goal. We will also increase the number of participants from whom we collect data, and expand our affective gesture set. These steps and more are covered in the next chapter.

Chapter 4

A Creature Object: Combining Smart Fur and Fabric Pressure Sensing¹

In the previous chapter, we presented a design for a new type of touch sensor built with conductive fur, and our initial evaluation of its potential for use in affective gesture recognition. Now, based on promising first results, we are ready to take it to the next level in an integrated creature-like prototype, with the goal of inferring several key emotional touch gestures, using a combination of synchronized touch sensing hardware.

This chapter covers the design of this $\sim 15 \times 20$ cm lap-pet object (Figure 4.1), including construction and integration of two low-tech, low-cost (Section 4.2.5) sensors: the conductive fur sensor, and an adaptation of an existing piezoresistive fabric pressure and position sensor. After a 16-participant study consisting of 9 key affective gestures, we present a gesture recognition model that averages 94% accuracy when evaluated on one individual at a time, and a combined accuracy of 86% when evaluated on the entire set of all participant data. We also attempt to recognize the identities of the 16 participants from their touch data, averaging 79% accuracy. These results inform recommendations for the future of gesture recognition in the Haptic Creature, and future emotional touch systems.

¹A version of Chapter 4 will be submitted for publication: Flagg, A., and MacLean, K. *Affective Touch Recognition for a Furry Zoomorphic Machine*. Manuscript in preparation.



Figure 4.1: Hand stroking the touch-sensing creature prototype. Prototype is equipped with conductive fur and piezoresistive fabric pressure sensing, and sized $\sim 15 \times 20$ cm.

4.1 Objectives and Approach

This chapter describes the integration of the conductive fur sensor with pressure sensing into a furry lap-animal prototype, and a quantitative evaluation of this prototype for use in recognition of some key emotional gestures. The goal of this evaluation is to establish the validity of our approach in the context of real pet capabilities, including: interpreting basic human touch patterns, learning personalized behaviours of a given individual, and recognizing and responding to important people. Evaluating potential in this scope will establish whether further development of our approach is justified.

As covered in detail in Chapter 3, further requirements for our system include hardware that is small and lightweight, easy to clean, inexpensive to construct, and involve relatively little wiring. We also require that the system have the potential to sample data and evaluate the classification model in real time. Additionally, we

set ourselves the initial application-derived goal of 80%+ recognition performance (see Section 4.4 for the reasoning behind this choice).

4.2 Integrated Sensor Construction and Electronics

Chapter 3 discusses the development of the conductive fur sensor in detail. Now, after some minor changes to the fur sensor, we move on to our process of adapting Schmeder and Freed’s piezoresistive fabric position and pressure sensor [20] to an arbitrary non-flat surface, done to fit our robot-body context. This fabric pressure sensor and the fur sensor are then connected to a single microprocessor to provide synchronized real-time data output, and mounted on a small creature body prototype.

The next sections detail our approach to each of these steps.

4.2.1 Fur Sensor Improvements

We made a few small changes to the fur sensor hardware to make it lighter, and tougher. First, we replaced the insulating layer of neoprene with a much thinner rubber layer, which reduces the weight of the sensor. However, this insulating layer is no longer thick enough to hold the threads in place, making it necessary to secure the threads by tying them into knots into the insulating layer, rather than simply looping them. This process increases manufacture time, but makes the sensor much more robust to wear and tear. Threads are secure in the fur now, very difficult to pull out.

4.2.2 Piezoresistive Fabric Pressure Sensor

Schmeder and Freed [20] introduced a pressure and position sensor comprised of a rectangular standoff layer of plastic mesh sandwiched between two layers of piezoresistive fabric. The sensor is wired at its four edges V_a , V_b , V_c , V_d as shown in Figure 4.2. Each of these nodes must be settable as input or output: that is, reading in a voltage measurement, or supplying an output voltage; and the surface resistivity of the fabric must be small in comparison to the material’s through-resistance.

Here, we briefly discuss the basic operating principle behind this piezoresistive




Figure 4.2: Adaptation of Schmeder and Freed’s piezoresistive fabric pressure and position sensor [20].

fabric pressure sensor design; full details can be found in Schmeder and Freed [20].

When the sensor is not being touched, the two pieces of piezoresistive material are physically separated by the mesh standoff layer. In this case, a supplied voltage can cause current to flow across the surfaces of the two separate fabrics, but no current will travel between them. A touch, however, puts them in physical contact, allowing current to flow both across the fabric surfaces, and between them, depending on where a voltage is applied. Using this fact, Schmeder and Freed demonstrate how it is possible to manipulate different input/output combinations of the nodes to measure different variables.

Referring to Figure 4.2: setting V_b to supply a high voltage and grounding V_d allows us to measure the x position of a touch with V_a . This is because a touch brings the two piezoresistive pieces in contact, creating a voltage read at V_a that is proportional to the resistance in the path along the surface of the lower fabric layer between the touch and the V_a node. Since this resistance is proportional to the physical distance between the touch and V_a , it provides an estimate of x posi-




Figure 4.3: Construction summary of the creature prototype hardware. Top down from left to right: (a) creature body consisting of styrofoam and plastic “skeleton” and “skin”, (b) first layer of piezoresistive fabric, (c) 3 mesh standoff layers, (d) second layer of piezoresistive fabric, and (e) touch-sensing conductive fur on top.

tion. We can similarly measure y position by setting V_a to supply a high voltage, grounding V_c , and reading V_b .

Finally, if we apply a voltage to V_b and ground V_a , both V_c and V_d will give a reading inversely proportional to pressure. This is because a touch connects the two pieces of piezoresistive fabric, allowing current to flow from the top fabric layer to the bottom fabric layer. Either of the two nodes V_c or V_d can then provide a pressure measurement, because as more pressure is applied, the resistance of the fabric decreases. Position of the touch does not interfere with the pressure reading, because the through-resistance of the fabric is large compared to its surface resistance. See Schmeder and Freed [20] for more details.

4.2.3 Construction and Adaptation

Next, we constructed a small rounded semi-spherical zoomorphic shape. Similar to the Haptic Creature, the head and body are sufficiently defined to suggest an animal form, but do not represent any particular species. The body's inner "skeleton" is carved out of styrofoam, on top of which is attached a soft, thin layer of plastic foam material to give the impression of skin-like elasticity.

We adapted the fabric pressure sensing design to this curved three-dimensional shape. By segmenting the body prototype into top, bottom, left and right hemispheres as in Figure 4.2, we can use the method discussed in Section 4.2.2 of alternating voltage input/output combinations to measure x and y position and pressure across the curved surface. The layers of the fabric pressure sensor are sewn to fit garment-like to the body of the prototype, and attached at the edges using conductive tape, which also provides connecting points for the wired nodes. Finally, the conductive fur sensor is sewn around the body on top of the pressure sensing fabric, in similar garment-like fashion (Figure 4.3).

4.2.4 Sensor Fusion

Applying the two sensors onto our three-dimensional shape involves wrapping the piezoresistive fabric fairly tightly, mounting it securely, and accommodating the (albeit small) weight of the fur on top. Altogether this applies a bias pressure on the pressure sensor. To avoid saturating the pressure reading, we use three plastic mesh standoff layers between the two piezoresistive pieces in the pressure sensor (rather than just one). Using three standoff layers increases the stiffness of the vertical structure between the top and bottom layers, adding more physical support holding the upper piezoresistive layer up and away from the bottom layer. Since the sensor is activated not by pressure directly, but by the pressure between the two piezoresistive pieces, increasing the structure between these pieces allows us to offset the weight of the hardware on top without needing to actually decrease its weight. Thus as the weight of the hardware pushes the top piezoresistive layer down towards the bottom piezoresistive layer, the middle standoff layer pushes it back upwards. In effect, the added structure partially desensitizes the sensor, meaning it is not activated by the upper hardware, and requires additional pressure

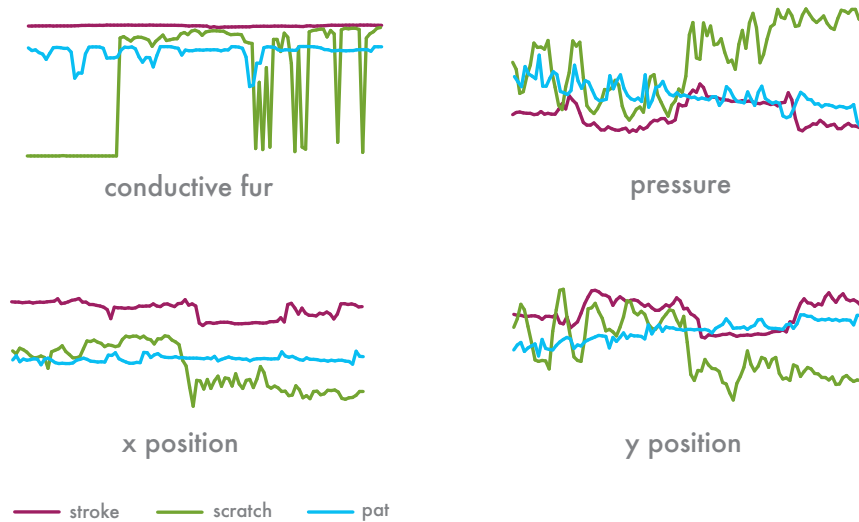


Figure 4.4: Sample fur sensor, pressure, x, and y data curves for stroke, scratch and pat gestures. (Axis scales excluded for clarity.)

(e.g. from a hand) to register a touch. Experimentation proved this arrangement allows secure construction as well as maximal pressure sensing range.

Input/output leads from the conductive fur sensor and the fabric pressure sensor are wired to the analog inputs of a Teensy 2.0 microcontroller [15]. We sample at 50 Hz, and at each iteration measure pressure, x position, y position, and fur sensor data by exciting and then reading the corresponding analog ports. This results in a synchronized stream of 4 time-series data curves, which can then be analyzed to classify gesture patterns (see Figure 4.4).

At the moment, recognition processing is done offline, however the models we have chosen can be evaluated very quickly, and thus are capable of performing at interactive rates when eventually integrated into the Creature system. By “interactive rates” we mean response time will be fast enough to keep up with the quickest gesture speed, missing no input data and causing no perceptible delay. A sampling rate of 50 Hz satisfies this requirement, as 5-10 Hz has been established as the maximum bandwidth for comfortable human voluntary motion [23].

Table 4.1: Estimated manufacture costs for a 20cm x 15cm gesture sensing prototype

<i>Part</i>	<i>Cost</i>
piezoresistive fabric sensor	\$30
Teensy microcontroller	\$16
conductive fur sensor	\$15
styrophoam and plastic foam skin layer	\$5
insulating rubber	\$5
wiring	\$2
Total	\$73

4.2.5 Construction Costs

One of the goals of this research is to develop *lower-cost* sensing technologies that still perform well enough (see Section 4.4) for our purposes in affective touch recognition. Specifically, in our context it is very important to have full continuous touch-sensitive coverage in order to hope to capture every touch; extremely high-resolution, high-accuracy sensing is less necessary. Many existing projects, including the current sensing component of the Haptic Creature [30], do not provide continuous coverage, and/or involve very expensive high-performance sensing technologies that are perhaps overkill solutions in this context [8, 10, 22, 25]. Low-tech methods may be sufficient for our purposes because affective touch gestures occur at relatively low speeds (<10 Hz) [23], span large (hand, finger-sized) areas, and do not involve highly precise pressure gradations. Further, the use of more sophisticated analytical methods may make up for some of what lower-tech hardware lacks.

Table 4.1 gives an estimated cost breakdown of manufacturing one of our 20cm x 15cm zoomorphic gesture sensing prototypes. Production in bulk would likely decrease costs considerably. As can be seen in the table, estimated total cost is less than \$75, far less than existing high-performance continuous touch-sensing technologies. For example, patches of similar size manufactured by Meka Robotics [12] or Roboskin [18] perform at much higher resolution and accuracy, but start in the thousands of dollars.

4.3 Gesture Evaluation and Analysis

Our analysis begins with a set of key emotional gestures. From Yohanan’s 30-item human-animal touch dictionary [31], we chose 9 key gestures that are a) appropriate within our lap-pet context, b) important for emotional communication, and c) feasible to perform on our prototype.

Our final set consists of the following gestures, as defined in Yohanan [31]:

stroke: moving hand gently over the prototype body, often repeatedly

scratch: rubbing the prototype with the fingernails

tickle: touching the prototype fur with light finger movements

squeeze: firmly pressing the prototype body between the fingers or both hands

pat: gently and quickly touching the body with the flat of the hand

rub: moving the hand repeatedly to and fro with firm pressure

pull: gently and randomly pulling at hairs in the fur

contact without movement: any undefined touch without motion

no touch: prototype left untouched

4.3.1 Gesture Data Collection

We collected gesture data from 16 participants (9 female), with cultural backgrounds in Canada, the United States, the Middle East, China, and Southern Asia. After viewing a list of the above gestures along with their definitions, participants were asked to hold the prototype on their laps, and perform each gesture continuously for several seconds. From each participant, we collected 25 2-second examples of each of the 9 gestures, sampling at 50 hertz, measuring x position, y position, pressure, and fur sensor data at each iteration. This process results in 4 time-series curves for each example collected. We chose a 2-second window of observation through informal experimentation: in pilot data collected from two participants, we found that a window shorter than 2 seconds (for instance, 1.5 seconds) resulted in degraded model performance, while increasing the window to larger than 2 seconds did not improve performance. In addition, while our participants took 1 second or less to perform even the longest gestures, we wanted each example to capture more than one continuous instance of each gesture, so as not to give our model the benefit of clearly delineated start and stop times. In this way

we better mimicked our eventual application, which will read in data continuously, and thus have no way of knowing the span of a gesture.

The gestures were collected from each participant in the same order. Randomizing gesture categories seemed unnecessary at the time, since beyond the supplied general gesture definitions, interpretation of each gesture type was left up to the participants, who were not told what were “correct” or “incorrect” touch behaviours. Thus, since gestures were not “tasks,” we did not expect learning effect to be a problem. Indeed, we did not observe a learning effect: however on reflection, randomizing the gesture types might have been a better approach. It is possible that participants became more comfortable with the prototype and the procedure itself as the study progressed, which might have made execution of later gestures truer and more natural.

Further, we did not attempt to randomize the gesture sampling by interspersing gesture examples. This choice was made due to the desire to collect a large amount of data, and the added time involved in asking participants to transition between gesture types for each of the 2 second samples. This was a limitation of our study, because the repetition involved could have encouraged the participants to settle into comfortable patterns, whereas randomization might have sparked more diversity in touch behaviours within a gesture.

In our analysis, we make sure to measure recognition accuracy across participants (as well as within), which we believe incorporates some of the gesture diversity that lack of randomization might have lost. However, future evaluations should include randomization.

4.3.2 Machine Learning Analysis

In machine learning analysis, recognition depends on the combination of measurable properties or “features” to help differentiate between data categories. Crucial to classification performance, a strong feature can distinguish between two or more gestures; i.e., its typical range is distinct for two or more gestures (see Section 3.4.2).

We extracted several standard sequence statistics from our time-series data to be used as features for training a classifier. Based on our previous work modelling

gesture data with the fur sensor (Section 3.4.2), we hypothesized that the following sequence features aid in prediction: maximum, minimum, mean, median, area under the curve, variance, and total variation. These features were calculated for each of the 4 time-series curves, resulting in a set of 28 features.

We then evaluated several standard classification models [11] using Weka [27], an open source framework supporting practical application of machine learning algorithms. Weka allowed us to quickly test many feature-based techniques that are considered highly effective for classification, including logistic regression, Bayesian networks, neural networks, and random forests [3, 11]. Model parameter settings can be found in Appendix A.

4.3.3 Results

We examined gesture recognition accuracy for models trained on each of our 16 participants individually, as well as for models trained on the combined group of all participants. In both cases we used 100-fold cross validation: the data is split randomly into 100 equally-sized subsets, and the model was tested on each of these subsets, after having been trained on the remaining data. Cross-validation is a standard approach to evaluating machine learning algorithms in a way that a) takes advantage of all available data, and b) avoids variation due to random partitioning that might throw off our estimate of the true predictive value of the model [11].

Accuracy was defined as the percentage of data cases in the subset that were labeled correctly by the model, averaged across the 100 subsets. The model assumed that each case is a true instance of one of the gesture classes, and always returned one of these labels as an answer. In future, we could incorporate a threshold confidence value, where a data case is labeled as “none of the above” if confidence for all known gesture types is low.

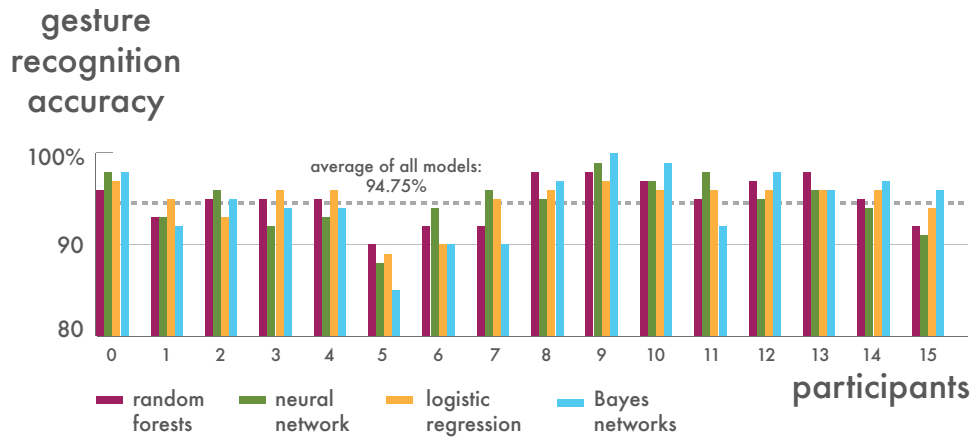


Figure 4.5: Gesture recognition results for random forests, neural networks, logistic regression, and Bayesian networks. Models were trained and tested on each of 16 study participants individually, using 100-fold cross-validation. (A graph range of 80-100% is used to maximize detail.)

Figure 4.5 summarizes classification performance on individual participants for some common feature-based models: random forests, neural networks, logistic regression, and Bayesian networks. Performance ranges from 85-100% accuracy. The neural network, logistic regression and Bayesian models average 95% accuracy, and the random forests model averages 94%. In particular, the random forests model classifies each participant with 90% accuracy or higher.

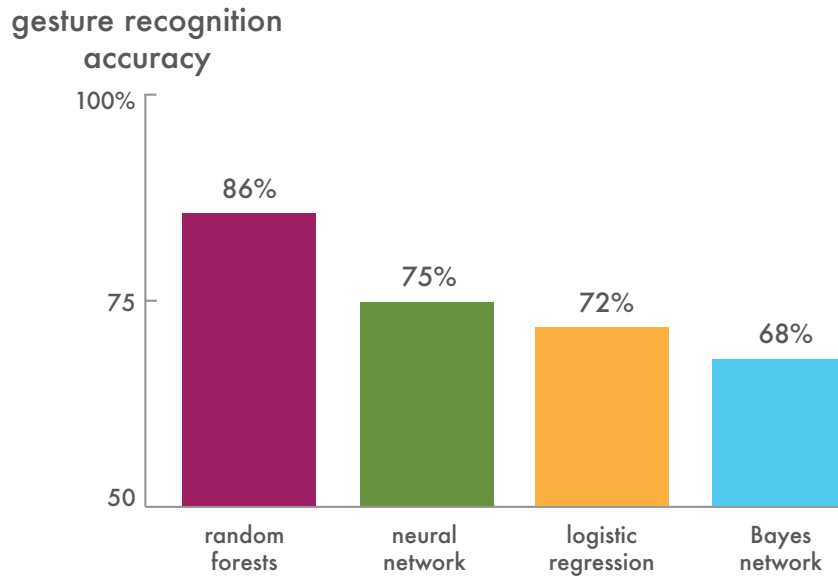


Figure 4.6: Gesture recognition results for random forests, neural networks, logistic regression, and Bayesian networks. Models were trained and tested on a combined set of all 16 participants, using 100-fold cross-validation. (50-100% displayed graph range).

Figure 4.6 summarizes classification performance on a combined data set of all participants, for the same models as above, evaluated using 100-fold cross-validation.




Figure 4.7: Hinton diagram of the confusion matrix corresponding to the random forests gesture recognition model, evaluated using 100-fold cross-validation on the combined data set of 16 study participants. Size of grey squares represents classification: for example, the first row shows how many examples of “stroke” are recognized as *stroke*, how many are mislabeled as *scratch*, and so on for all gesture types.

The Hinton diagram in Figure 4.7 is a visualization of the confusion matrix for the highest-performing model, random forests, which classifies the set with 86% accuracy. The random forests model is selected as the best of the reported models, because while all perform similarly well on individual data, the random forests model is significantly better on the combined set.

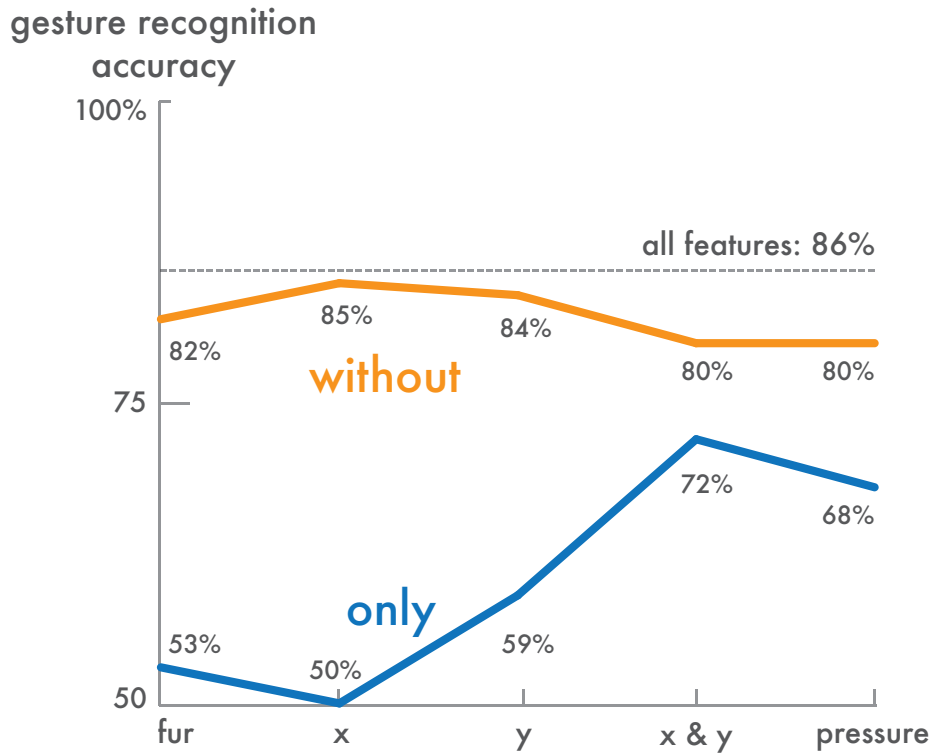


Figure 4.8: Contributions of individual sensor channels to gesture classification performance of random forests model on combined 16-participant data set, evaluated using 100-fold cross validation. Upper curve shows performance when using features from all data channels *except* conductive fur, then *except* x position, *except* y position, *except* combined x and y position, and *except* pressure. Lower curve shows accuracy when using features from *only* conductive fur, then *only* x position, *only* y position, *only* combined x and y position, and *only* pressure. (50-100% displayed graph range).

In Figure 4.8 we report the relative drop in performance when the fur sensor data is left out, when position data is left out, and when pressure data is left out (upper curve). We also report accuracy when the model is given only conductive fur data, only position data, and only pressure data (lower curve). Combined x and y position is the strongest channel for this set of gestures: the model achieves 72% accuracy using position alone, and suffers a 6% drop when this data is removed.

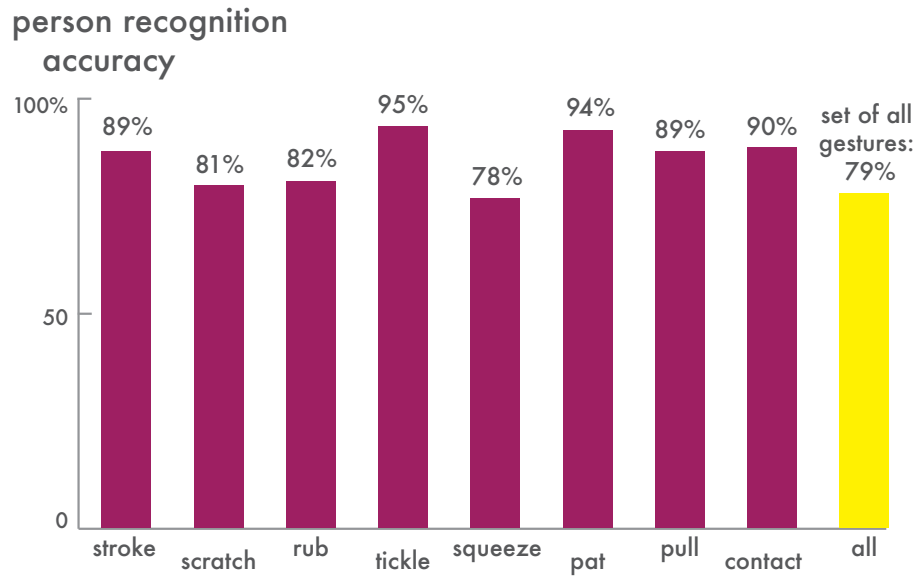


Figure 4.9: Results of evaluating a random forests model’s ability to recognize people from their touch data. Model was trained and tested on *stroke* data for all participants, *scratch* data for all participants, etc for all gestures. Model was also evaluated on combined set of participant data for all gestures. All tests used 100-fold cross-validation.

Finally, Figure 4.9 reports the results of training a model to recognize specific people from their touch gesture data. In the combined set of data from all gestures, the model can recognize who out of the 16 participants is touching the prototype with an accuracy of 79%.

4.4 Discussion

We evaluate our prototype on its ability to differentiate a single person’s gestures, to recognize gestures across participants, and to recognize individual people from their touch interaction.

All three results are of interest to us. Individual gesture recognition shows the system’s capacity for personalization, much like how a pet develops a relationship with an owner, and learns to interpret the owner’s specific behaviours. Combined

gesture results show the system’s potential for generalizability across a wide audience, representing its understanding of the basic rules of human touch behaviour. Person recognition results indicate how the system could be trained to recognize and respond to important people in the environment, such as the owner.

High performance for individuals: As shown in Figure 4.5, a model trained on a single person’s touch data can generally achieve very high recognition performance. In our tests, accuracy is almost always above 90%, on average higher than 95%, and quite consistent across classifiers.

Performance decreases in group model but remains high: As might be expected, modelling all participants together decreases gesture recognition accuracy, as shown in Figure 4.6. The effect of model choice also becomes more evident, with performance among classifiers ranging from 68-86%.

This results suggests that touch behaviours are by no means universal - in fact, quite the opposite. It also appears that while people vary widely in their individual interpretations of a given gesture, they are likely to stick to their respective interpretations, and perform each gesture in a consistent way.

The random forests model has highest performance, classifying the set with 86% accuracy. From the confusion matrix visualization in Figure 4.7, it appears that the *rub* gesture is easy to confuse with *scratch*, and vice-versa, as might be expected. *Pull* is quite difficult to classify, often confused with *stroke*, *scratch*, or *tickle*.

Position most essential, pressure & fur contribute unique data: We also attempt to characterize the value of each data channel to prediction of our set of gestures. As can be seen in Figure 4.8, each channel contributes to classification performance. Combined x and y position is the best individual classifier, and when left out, performance suffers a 6% drop, as it does when pressure is left out. As an individual classifier, the fur sensor is only stronger than x position, and performs significantly worse than pressure alone, y position alone, or combined x and y position. However, when the conductive fur data is removed from the model, performance drops 4%, more than the performance drops corresponding to either x position or y position (1% and 2%). This suggests that the conductive fur sensor contributes a relatively orthogonal channel of information, valuable to recognition. Removing pressure information also has a big negative impact on performance,

more so than removing x data or removing y data.

However, Schmeder and Freed [20] described some possible noisy dependencies in the fabric pressure sensor between pressure and position reads, suggesting our pressure/position data is not completely accurate at present. Further work will be needed to ensure this issue is not interfering with the relative evident individual contributions of x position data and y position data.

Model can sometimes recognize people from their touch: Finally, we examine our model’s ability to recognize an individual person from her touch interaction. As shown in Figure 4.9, the 16 participants can each be recognized from the way they perform a given gesture with an accuracy of 78% or higher. In the combined set of all gesture types, we find that a person is recognizable by the model with 79% accuracy. This accuracy could potentially improve over time, as the model is exposed to more and more data from a given person.

Some gestures are more telling than others: for instance, people appear to be most recognizable from the way they perform the *tickle* gesture, in which the model is 95% accurate. This result makes sense given our observations of the widely varying ways participants execute the *tickle* gesture.

Performance competitive with previous affective touch projects: Despite the drop in gesture recognition performance when generalizing to a group, we are still able to model our selected gestures quite successfully relative to previous literature in affective gesture recognition. As discussed in Section 4.2.5, expensive high-performance sensors are not necessary in this context; our low-tech approach performs well enough for competitive touch recognition results, despite the low one-off cost of ~\$73 for a 20cm x 15cm prototype.

When comparing to previous results, however, it is important to remember that these accuracy rates are related to gesture selection. At present, we have focused on a set of gestures key to emotional communication, and while there is overlap between our gesture set and previous sets, they are not identical. It is likely that different gesture sets will result in different recognition rates both in terms of recognition and differentiation, and a more valid comparison to previous work will require evaluation of identical gesture sets, a topic for future work.

Hypothesizing practical use goals: In terms of practical usefulness, classification accuracy rates are usually evaluated in terms of the specific problem domain.

Since limited work has been done in this area to establish what “good” performance is for practical applications of affective touch gesture recognition, choosing a performance goal for ourselves is nontrivial. We hypothesize that in the realm of an emotional companion-like machine, strict correctness is perhaps less important than in other computing applications, since a certain amount of unpredictability might be expected. Further, the idea of “correctness” can be defined as either a) how well the system recognizes a person’s *intended* gesture, regardless of how it was actually carried out, or b) how well the system itself *defines* a gesture to support practical use. In the first case, we would evaluate our system by its ability to meet audience expectations, without giving it the benefit of any expectations of its own. In the second case, the system has its own model of expectations based on general norms of touch behaviour, and probably its own corresponding likes, dislikes and various other beliefs and reactions, which might not always cater strictly to audience needs and expectations. In this second case, in addition to examining standard recognition rates, we would also evaluate the system’s model by how easy it is for the audience to communicate effectively with it, and how valuable the resulting interactions are.

If we assume the second definition of correctness, given these factors we might initially hypothesize that a real-world (out of lab) recognition performance of 80%+ for ~10 or more gestures is a good preliminary goal for an affective touch recognition system. We believe such a system would be capable of a) perceiving and responding in a large number of complex, subtle ways, b) capturing a significant majority of audience expectations, and c) expressing a “personality” and sense of unpredictability that might make for a compelling interaction in both short and long term.

4.5 Future Work

Our next steps include both hardware and analytical improvements. Hardware-wise, we will attempt to improve the pressure and position sensing currently in our prototype by experimenting with different piezoresistive fabrics, with different surface and through-resistances, and different material specifications. We will also compare gesture recognition performance in alternate pressure sensor types,

possibly including multiplexed designs for higher position resolution. And we will compare performance in larger and smaller prototypes, including designs which combine multiple conductive fur and pressure sensing “patches” on different parts of the prototype anatomy. We will also experiment with additional sensor types, such as an accelerometer. Further, the system will need to be tested on a larger pool of participants, so as to better represent our audience population, and make for a better-informed model. Future data collection should also include more randomization of gesture types, the lack of which was a limitation of our study. Perhaps most importantly, for a true evaluation of our design and its implications, we need to get the system out of our lab environment and into the hands of real people. The way people perform gestures in the constrained setting of the lab chair may be very different from how they behave in the real world, when they are free to carry the device, put it in different places, hold it in different ways. It is difficult to predict how performance will be affected by these factors; experiments collecting and modelling this real-world data will tell us more.

On the analysis side, we will investigate more machine learning schemes to attempt to improve performance. As feature selection is often found to be more important than choice of classifier, we will also explore different feature types beyond the standard ones that we are using at present, especially those that are specific to time-series data. We will also compare time-series-specific classification routines, which we hope will result in higher performance than our current general approach. In particular, a Hidden-Markov Model [24] might be a natural choice to investigate, given the online nature of our ultimate application.

Future models should incorporate confidence estimates, so that the model is not forced to classify everything it sees in terms of its known gesture types, even if probability is very low. Also, before integrating any recognition system into the Haptic Creature, we will need to transition to continuous gesture recognition that works with ongoing sequences of data, rather than fixed-length windows. Given our current approach, this could involve buffering in 2-second periods. Since some gestures tend to take more time than others, perhaps running a few buffers of different length simultaneously could help prevent shorter gestures from getting lost in too-long windows, as suggested in Chan, et. al. [5].

Up until now, we have only discussed supervised learning, in which models are

trained on *labeled* data. However other interesting possibilities include the use of unsupervised learning and reinforcement learning, both of which attempt to mimic how a real living creature learns [6, 26]. In unsupervised learning, the model is shown lots of data and searches for hidden structure, for instance using clustering or dimensionality reduction [6]. In reinforcement learning, the system would learn by interacting with its environment, and observing the resulting feedback, such as with a Markov decision process [26]. This type of approach would be more complex than a simpler supervised model, but carries fascinating potential. Such a system would learn directly from its own unique environment, and thus develop its own “personality.” No two would be alike. It would be sophisticated and unpredictable, constantly learning and evolving based on observations, experiences and interactions.

To summarize: based on the original objectives of this research (Section 4.1), the reported results justify further development of our approach to affective gesture recognition. Specifically, the following are promising directions for future research: a) interpreting touch behaviours of a given person, b) generalizing across a population to understand basic patterns of human touch, c) learning to recognize and respond to important people, such as an owner, and d) designing models for learning and “personality”. We believe these types of capabilities will enable the system to engage with its audience on an emotional level, and ultimately provide a sophisticated level of therapeutic support.

Chapter 5

Conclusions and Future Work

This thesis began with an idea for harnessing a computer’s therapeutic power through sophisticated recognition of affective touch in a human/animal metaphor. In this research, we have investigated the possibility of enabling such a system through the combination of low-tech hardware with artificial intelligence methods. Low cost has been a particular priority, and as discussed in Section 4.2.5, our one-off manufacture costs for a 20cm x 15cm prototype was ~\$73.

We have presented the design, construction and evaluation of an early prototype for what affective gesture recognition might look like in a furry artificial lap creature system. This involved the invention of a new type of conductive “smart fur” touch sensor, designed to augment conventional touch sensing technologies by quantifying hand motion information. We described the physical design process behind this unique approach to touch and gesture recognition, as well as the original inspiration for the idea. Next, the integration of this conductive fur sensor with an adaptation of an existing piezoresistive pressure sensor was described, along with our work prototyping this hardware into a zoomorphic interface for collecting affective touch gesture data.

After running a 16-participant study evaluating the above combined sensing system, we have analyzed and modelled sensor data for 9 gestures key to emotional communication. Given our original design objectives, the reported results show that our system may be capable of recognizing this set of gestures, and possibly others, in a real world setting, accurately enough to inform a meaningful, nontrivial

emotional interaction (Section 4.4). Specifically, our approach shows promise as a low-cost means to enable the Creature system to understand and respond to basic rules of human touch, personalize to individual touch behaviours and preferences, and appreciate important people in its environment.

In the immediate future, our next steps include development and improvement of the physical design of the sensing system, further analysis work to improve our learning models, and evaluations of the system in real-world use (Section 4.5). Given positive results validating continued use of this approach, the sensing system can then be integrated into the Haptic Creature project. At that point, we can begin the fascinating process of attempting to infer human emotion through analysis of touch. Even partial success at this stage will be a momentous milestone for the Creature project, as the system will then have closed the loop: perceiving emotion, and expressing emotion.

In the mean time, we imagine many other applications of the fur sensor and the combined touch sensing system. These could include the design of haptic interactive children's toys, emotional teaching tools, smart carpets and furniture, wearables, and other conceivable low-tech flexible devices requiring touch sensing.

In the longer term, we hope that the capabilities partly aided by this research will have empowered the Creature as a platform for emotional communication and support. If so, other future haptic affective systems could similarly utilize and build upon this work for applications involving affective touch. Some potential areas of use include education, companionship, therapy, rehabilitation, treatment of cognitive disorders, and assistive technologies.

Given these exciting potential applications, we believe the development of touch-based emotional machine design is a valuable direction for future research. We hope the work presented in this thesis can be a small first step in this direction, ushering in a generation of smarter, more emotionally sophisticated machines that can help people feel better.

Bibliography

- [1] Arduino. LilyPad, 2012. <http://arduino.cc/en/Main/ArduinoBoardLilyPad>. → pages 19
- [2] Arduino. Uno, 2012. <http://arduino.cc/en/Main/ArduinoBoardUno>. → pages 29
- [3] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006. → pages 25, 40
- [4] P. Brockwell and R. Davis. Introduction to time series and forecasting. In *Springer-Verlag, New York*, 2002. → pages 12, 13
- [5] J. Chang, K. E. MacLean, and S. Yohanan. Gesture recognition in the haptic creature. In *EuroHaptics 2010*, 2010. → pages 12, 49
- [6] R. O. Duda, P. E. Hart, and D. G. Stork. *Unsupervised Learning and Clustering, Ch. 10 in Pattern classification*. John Wiley and Sons, Inc., 2001. → pages 50
- [7] A. Freed. Novel and forgotten current-steering techniques for resistive multitouch, duotouch, and polytouch position sensing with pressure. In *New Interfaces for Musical Expression*, 2009. → pages 11, 12
- [8] B. Friedman, P. H. K. Jr., and J. Hagman. Hardware companions?: what online aibo discussion forums reveal about the human-robotic relationship. In *CHI '03*, pages 273–280, 2003. → pages 7, 9, 37
- [9] FSRs. Interlink Forse Sensing Resistors, 2012. <http://www.interlinkelectronics.com/products.php>. → pages 8
- [10] K. Goris, J. Saldien, I. Vanderniepen, and D. Lefeber. The huggable robot probo, a multi-disciplinary research platform. In *Eurobot '08*, 2008. → pages 7, 37

- [11] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, second edition, 2009. → pages 23, 25, 40
- [12] Meka. Dextrous Mobile Manipulators and Humanoids, 2012. <http://mekabot.com/?ref=Hizook-spotlight1>. → pages 37
- [13] H. Perner-Wilson. Stroke sensor. Plusea, www.plusea.at, May 2011. → pages viii, 4, 5, 10, 14, 17, 18
- [14] R. Picard. *Affective Computing*. MIT Media Laboratory, 1997. → pages 1
- [15] PJRC. Teensy USB Development Board, 2012. <http://www.pjrc.com/teensy/>. → pages 36
- [16] QTCs. Peratech Ltd. Quantum Tunnelling Composites, 2012. <http://www.peratech.com/qtcmaterial.php>. → pages 8
- [17] C. Rafael and S. D’Melle. Affect detection: An interdisciplinary review of models, methods, and their applications. In *IEEE Transactions on Affective Computing*, 2010. → pages 3
- [18] Roboskin. Skin-based Technologies and Capabilities for Safe, Autonomous and Interactive Robots, 2012. <http://www.roboskin.eu/index.php/home>. → pages 37
- [19] J. Roh, A. Freed, Y. Mann, and D. Wessel. Robust and reliable fabric, piezoresistive multitouch sensing surfaces for musical controllers. In *New Interfaces for Musical Expression*, 2011. → pages 11, 12
- [20] A. Schmeder and A. Freed. Support vector machine learning for gesture signal estimation with a piezo resistive fabric touch surface. In *New Interfaces for Musical Expression*, 2010. → pages viii, ix, 4, 6, 11, 12, 32, 33, 34, 47
- [21] Y. Sefidgar. *TAMER: Touch-guided Anxiety Management via Engagement with a Robotic Pet, Efficacy Evaluation and the First Steps of Interaction Design*. M.Sc. Thesis, University of British Columbia, 2012. → pages 3
- [22] T. Shibata, K. Inoue, and R. Irie. Emotional robot for intelligent system: Artificial emotional creature project. In *Proceedings of IIZUKA*, pages 43–48, 2006. → pages 7, 9, 37

- [23] K. B. Shimoga. Finger force and touch feedback issues in dexterous telemanipulation. In *Proceedings of Fourth Annual Conference on Intelligent Robotic Systems for Space Exploration*, 2009. → pages 36, 37
- [24] R. Singh, T. Jaakkola, and A. Mohammad. *Machine Learning*. MIT OpenCourseWare, 2006. → pages 49
- [25] 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. → pages viii, 7, 8, 9, 12, 37
- [26] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998. → pages 50
- [27] I. H. Witten and E. Frank. *Data mining: Practical machine learning tools and techniques*. Morgan Kaufmann Publishers, second edition, 2005. → pages 13, 25, 40, 56
- [28] S. Yohanan. *The Haptic Creature: Social Human-Robot Interaction through Affective Touch*. Phd Thesis, University of British Columbia, 2012. → pages viii, 3
- [29] 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. → pages 9
- [30] S. Yohanan and K. E. MacLean. A tool to study affect touch: Goals & design of the haptic creature. In *Conference on Human Factors in Computing Systems (CHI '09)*, 2009. → pages 37
- [31] S. Yohanan and K. 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. → pages 4, 22, 25, 38
- [32] S. Yohanan, J. Hall, K. MacLean, E. Croft, M. V. der Loos, M. Baumann, J. Chang, D. Nielsen, and S. Zoghbi. Affect-driven emotional expression with the haptic creature. In *Proceedings of UIST, User Interface Software and Technology*, 2009. → pages 2

Appendix A

Model Parameters

The following is a summary of the parameter settings used in the analytical models presented in this thesis. To reproduce our results, upload the corresponding data sets to Weka [27] with these parameter values. Data will be posted online shortly, pending ethics approval.

Logistic Regression Parameters (Chapters 3 & 4)

ridge regularization parameter : 1.0 E-8

maximum iterations: not fixed, optimize weights fully (*option '-1' in Weka*)

Bayesian Network Parameters (Chapter 4)

conditional probabilities estimator algorithm: simple estimator

network search algorithm: K2

Neural Network Parameters (Chapter 4)

decay learning rate: false

hidden layers: 19 (*option 'a' in Weka*)

learning rate: 0.3

momentum: 0.2

normalize attributes: true

seed: 0

number of training epochs: 500

validation set size: no set, train for specified epochs (*option '0' in Weka*)

validation threshold: 20

Random Forests Parameters (Chapter 4)

maximum depth: unlimited tree depth (*option '0' in Weka*)

number of features to be used for random selection: 5 (*option '0' in Weka*)

number of trees: 10

seed: 1

Appendix B

Study Documents

- Study Consent Form
- Study Recruitment Form



RESEARCHER'S COPY CONSENT FORM

Department of Computer Science
2366 Main Mall
Vancouver, B.C. Canada V6T 1Z4
tel: [REDACTED]
fax: [REDACTED]

Project Title: The Haptic Creature Project – Affective Touch Influence
(UBC Ethics #H01-80470)

Principal Investigators: Dr. Karon MacLean; Associate Professor; Dept of Computer Science; [REDACTED]
Student Investigator: Anna Flagg; M.Sc. Candidate; Dept of Computer Science; [REDACTED]

The purpose of this study is to examine how people perform different emotion-related touch gestures. The gestures will be standard touch interactions such as a person might use with a dog or cat, i.e., “pat,” “stroke,” “scratch,” etc. In this study you will be asked to perform several of these standard touch gestures on a small patch of “smart” fur. We will collect sensor readings from the fur prototype, and this data will be used later on to train a computerized model to recognize gestures. If you are not sure about any instructions, do not hesitate to ask.

- REIMBURSEMENT: \$10
- TIME COMMITMENT: 1 × 30 minute session
- CONFIDENTIALITY: *You will not be identified by name in any study reports. Data gathered from this experiment will be stored in a secure Computer Science account accessible only to the experimenters.*

You understand that the experimenter will ANSWER ANY QUESTIONS you have about the instructions or the procedures of this study. After participating, the experimenter will answer any other questions you have about this study.

Your participation in this study is entirely voluntary and you may refuse to participate or withdraw from the study at any time without jeopardy. Your signature below indicates that you have received a copy of this consent form for your own records, and consent to participate in this study.

If you have any concerns about your treatment or rights as a research subject, you may contact the Research Subject Info Line in the UBC Office of Research Services at [REDACTED].

You hereby CONSENT to participate and acknowledge RECEIPT of a copy of the consent form:

Printed Name _____ Date _____ Signature _____

Smart Fur Project Touch Gesture Recognition

Thank you for your interest in participating in our study!
For instructions on how to sign up, please go to the bottom of the page.

Research Overview

I am a member of the SPIN research group in the Department of Computer Science at the University of British Columbia. I am recruiting participants in user studies as part of my M.Sc. Research under the supervision of Dr. Karon MacLean.

Our research is investigating the manner in which humans interact through the sense of touch. For this purpose, we are developing the Haptic Creature: a small, furry robot capable of expressing and recognizing emotion through the sense of touch.

Study Details

The purpose of this study is to examine how people interact through the sense of touch. In particular, this study examines how different people perform different emotion-related touch gestures, such as those one might use with a dog or cat, i.e. "pet", "stroke", "scratch", etc.

You will perform several of these standard gestures on a small patch of "smart" fur. We will collect sensor readings from the fur prototype during your touch interaction, and this data will be used later on to train a computerized model to recognize gestures.

General Information

- * The study will take approximately half an hour to complete.
- * You will be compensated \$10 for your participation.
- * The study will be conducted at the Vancouver campus of the University of British Columbia.

Restrictions on Participation

- * You must be between 19 and 50 years old.

This study has been approved by The University of British Columbia; Office of Research Services; Behavioural Research Ethics Board (#H01-80470).

If you have any questions, do not hesitate to email me:
Anna Flagg
<redacted>

Signup Instructions

<logistical details on how to sign up will be provided here.>