Affective Touch Gesture Recognition for a Furry Zoomorphic Machine

Anna Flagg
Department of Computer Science
University of British Columbia
me@annaflagg.com

Karon MacLean
Department of Computer Science
University of British Columbia
maclean@cs.ubc.ca

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 through haptic emotional interaction, which requires sophisticated touch sensing and interpretation. We explore the development of touch 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. In this work, we build upon a previous design for a new type of fur-based touch sensor. Here, we integrate our fur sensor with a piezoresistive fabric location/pressure sensor, and adapt the combined design to cover a curved creature-like object. We then use this interface to collect synchronized time-series data from the two sensors, and perform machine learning analysis to recognize 9 key affective touch gestures. In a study of 16 participants, our model averages 94% recognition accuracy when trained on individuals, and 86% when applied to the combined set of all participants. The model can also recognize which participant is touching the prototype with 79% accuracy. These results promise a new generation of emotionally intelligent machines, enabled by affective touch gesture recognition.

Author Keywords
Haptics, Tangible interaction, Conductive fur, Smart fur, Affective touch, Piezoresistive pressure sensing, Gesture recognition.

MOTIVATION
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.

Figure 1. Complete affective touch-sensing zoomorphic prototype with pressure and conductive fur sensors (left), inner styrofoam and plastic skeleton (top right), skeleton with fabric pressure sensor (bottom right).

However, 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 range of potential uses in therapy, rehabilitation, education, treatment of cognitive disorders, and assistance for people with special needs.

THERAPEUTIC MACHINES: THE HAPTIC CREATURE
It has long been known that pets can have a positive effect on their owner's emotions. Could this phenomenon ever occur in interactions between people and artificial systems?

Exploring this direction, our group has developed the Haptic Creature [27], a furry lap-sized social robot that perceives the world through touch, and expresses itself through ear stiffness modulations, 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 by sensing and responding to human emotion. In a recent study in which participants were set up with wearable biometric sensors to indicate stress level, our group found that the Creature was capable of reducing anxiety markers in individuals who experience its active breathing [20]. This important result motivates further development of the Creature's therapeutic capabilities. In particular, the current method of sensing emotional state through wearable biometric sensors is too intrusive to be acceptable in the long run. This raises the question: how can we sense emotion less intrusively?

Posture, speech, touch, voice prosody, and physiological measures have been explored as possible paths to emotion sensing [16]. 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 emotional 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 [30]. This result justifies further work investigating the Creature as a platform for sensing the emotional content of touch.

CONTRIBUTIONS

Our goal is to sense human emotion through analysis of touch interaction with the Creature. This endeavour involves sensing and recognizing touch gestures, and then inferring emotional state from those touch gestures. In this paper, we investigate the former. We present the combination of physical design work and artificial intelligence methods that enable our new low-cost, low-tech touch sensing system to achieve recognition rates competitive with existing gesture recognition systems that use costlier sensing technologies.

Our physical prototype is a major expansion of our previous design for a new type of fur-based touch sensor [5]. We adapted this sensor to cover a curved surface, and integrated an additional piezoresistive fabric sensor to measure pressure, resulting in a small (∼20cm x 15cm) furry animal-like prototype that outputs synchronized time-series touch data (Figure 1). We then used the prototype as an interface to collect human touch data for 9 key emotional gestures, and classified them using machine learning analysis. The resulting realtime model predicts gesture type with an average accuracy of 94% for a given individual, and 86% when generalized across the combined set of all 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 for modelling gesture.

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.

Haptic Affective Robots

Huggable, PARO, Aibo and Probo are all touch-sensitive affective social robots relevant to our work [24, 21, 7, 9]. Of these, Huggable [24], 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. PARO [21] 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 [7] 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 [9] is an elephant-like social robot equipped with a large variety of sensors. It expresses 7 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.

Touch Sensing

Most affective systems focus on sensing touch through force, such as with Force Sensitive Resistors (FSRs) [8] or Quantum Tunnelling Composites (QTCs) [15]. PARO and Aibo use FSRs alone to identify touch, and hitherto the Haptic Creature has as well [21, 7, 28]. Huggable also uses FSRs, in conjunction with temperature sensors and capacitive sensors [24].

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 difficult nonlinearities, and they are also not readily 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.

Hand Motion Sensing

Current touch sensing technologies rely largely on force alone, which handicaps the system's sensitivity to light touches, and differentiation of gestures with 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. We have therefore developed, in previous work, a sensor aimed at describing *above-surface hand motion* information [5].

Now, we take the initial sensor presented in [5], improve it, adapt it to cover a curved surface, combine it with a pressure sensing design, and integrate the two into a creature-like prototype for extensive touch analysis (details below).

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, hardware complexity, degradation of performance on curved surfaces, and non-continuous sensor coverage. Piezoresistive fabrics, (which change resistivity when a pressure is applied, and can thus be used as pressure sensors), address many of these problems. In particular, these fabrics are inexpensive, flexible, and can be easily wrapped and sewn around irregular three-dimensional shapes, making them well-suited for use in the Haptic Creature.

As a first step, we selected one of the simpler piezoresistive fabric pressure sensor designs from the many proposed by Schmeder and Freed [19], adapted it to our three-dimensional robot body shape, and evaluated it for use in affective gesture recognition.

Affective Touch Gesture Recognition

The use of machine learning for touch gesture recognition in affective systems is in early stages, and formal accuracy rates are rare in the current literature. Of the works we have mentioned, only Huggable and the Haptic Creature projects involve 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 [24]. 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 [3], and reports 77% recognition accuracy for a set of four gestures.

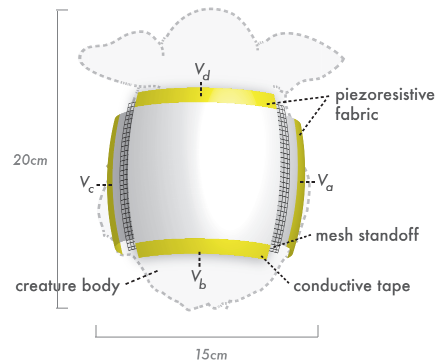

Figure 2. Adaptation of Schmeder and Freed's piezoresistive fabric pressure and position sensor.

We also took a feature-based approach, but rather than defining features in terms of the relationships among multiple sensors read instantaneously, we extracted standard sequence statistics as features from our synchronized time-series sensor curves [2]. In this way we incorporated the time-dependent nature of gestures into our model, giving it a "memory," which we have not seen in similar previous work, and which we hypothesize is key to performance.

APPROACH AND APPARATUS

This section is a step-by-step description of the design of our low-cost touch-sensitive zoomorphic prototype.

Conductive Fur Sensor

Our conductive fur sensor is a new type of touch sensor made up of conductive threads embedded in animal-like fur [5]. Inspired by Perner-Wilson and Satomi's conductive thread stroke sensor [12], our conductive fur touch and gesture sensor is in essence a big resistor: when we run a weak current through the fur, the touch of a hand will disturb the configuration of the conductive threads inside. This disturbance alters the number of electrical connections between threads and thus changes the overall sensor resistance, which is reflected in voltage changes across the fur. These voltage changes can be collected and analyzed in realtime. See [5] for a full discussion of the conductive fur sensor.

Piezoresistive Fabric Pressure Sensor

Schmeder and Freed [19] 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 2, which must each be settable as input or output, and the surface resistivity of the fabric must be small in comparison to the material's through-resistance. While at rest, the two pieces of piezoresistive material are physically separated by the mesh standoff layer, and thus a supplied voltage can cause current to flow across the surfaces of the two separate fabrics, but not 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 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. 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 the x position. We can similarly measure the y position. Further, if we apply a voltage to V_b and ground V_a, V_c will give a reading inversely proportional to pressure. This is because as current flows from the top layer to the bottom layer, more pressure decreases the resistance of the fabric, read at V_c. 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 [19] for more details.

Sensor Construction, Adaptation, and Fusion

We constructed a small rounded semi-spherical zoomorphic shape as a base for our prototype. Head and body are sufficiently defined to suggest an animal form, but do not represent any particular species. This 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 (Figure 2). Segmenting the body into top, bottom, left and right hemispheres allowed us to measure position and pressure on the curved surface using the method described above. 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. Then the conductive fur sensor is mounted on top, in similar garment-like fashion.

Constructing this three-dimensional shape involves wrapping the fabric layers fairly tightly, mounting the hardware securely, and accommodating the (albeit small) weight of the fur. A side effect of this process is 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 (rather than just one). Using three layers increases the stiffness of the vertical structure holding the top 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 them allows us to offset the weight of the hardware on top without needing to actually decrease its weight. Input/output leads from the conductive fur sensor and the fabric pressure sensor are then wired to the analog inputs of a Teensy 2.0 microcontroller [14]. 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. 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. A sampling rate of 50 Hz satisfies our speed requirements for measuring gestures, as 5-10 Hz has been established as the maximum bandwidth for comfortable human voluntary motion (as opposed to sensation) [22].

Construction Costs

One of the goals of this research is to develop *lower cost* sensing technologies that still perform well enough (see *Discussion*) for our purposes in affective touch recognition. Specifically, in our context is it 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 do not provide continuous coverage, and/or involve expensive high-performance hardware that is perhaps overkill in this context [29, 24, 21, 7, 9]. Low-tech methods may be sufficient because hand gestures occur at relatively low speeds [22], span large (hand-sized) areas, and do not involve highly precise pressure gradations. Further, the use of analytical methods may make up for some of what lower-tech hardware lacks. Our estimated total cost of manufacturing one of these zoomorphic prototypes is $73 (and would likely decrease in bulk production), far less than existing high-performance continuous touch-sensing technologies. For example, patches of similar size manufactured by Meka Robotics [11] or Roboskin [17] perform at much higher resolution and accuracy, but start in the thousands of dollars.

GESTURE EVALUATION AND ANALYSIS

We chose 9 key emotional gestures from Yohanan's 30-item human-animal touch dictionary [30] 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: ***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.

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 gesture definitions, participants held prototype on their laps, and performed 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. We chose a 2-second window of observation through informal experimentation: in pilot data collected from two

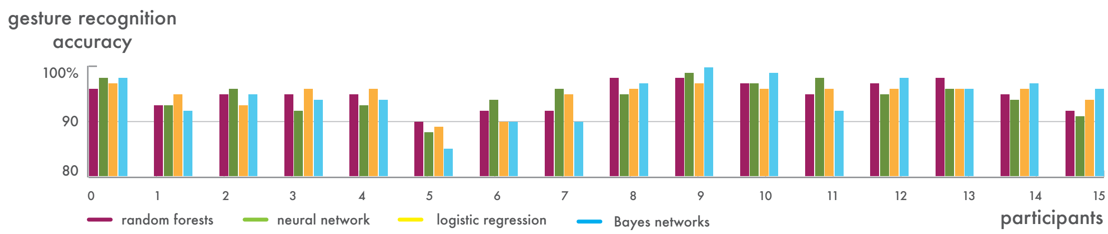

Figure 3. Gesture recognition results for random forests, neural networks, logistic regression, and Bayesian networks. Models were trained and tested on each of 16 study participants individually. (A graph range of 80-100% is used to maximize detail.)

participants, a window shorter than 2 seconds resulted in degraded model performance, while increasing the window did not improve performance. These results suggested that a 2-second window might be a reasonable choice for the main experiment (though further work will be needed to formally optimize window length for real-world applications). 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, making learning effect unlikely. And indeed, we did not observe a learning effect: however on reflection, randomizing the gesture types might have been a better approach. Participants maybe have become more comfortable as the study progressed, possibly making execution of later gestures truer and more natural. Further, we did not attempt to randomly intersperse gesture examples due to the added time involved in transitioning 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 diversity that lack of randomization might have lost. However, future evaluations should include randomization.

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. 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 [5], 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 models considered highly effective for classification [1, 10] using Weka [26], an open source framework supporting practical application of machine learning algorithms.

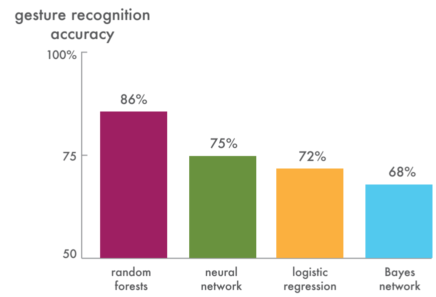

Figure 4. Gesture recognition results for random forests, neural networks, logistic regression, and Bayesian networks. Models were trained and tested on a combined set of all 16 participants. (50-100% displayed graph range).

Results

We evaluated our recognition models, in all cases using 100-fold cross-validation. Cross-validation is a standard approach to evaluating performance 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 [10]. Accuracy was defined as the percentage of data cases in the subset that were labeled correctly by the model. The model assumed that each case was a true instance of one of the gesture classes, and always returned one of these labels as an answer.

Figure 3 summarizes classification performance on *individual participants* for the following standard models: random forests, neural networks, logistic regression, and Bayesian networks. Figure 4 summarizes classification performance for the same models on a *combined* data set of all participants. The Hinton diagram in Figure 5 is a visualization of the confusion matrix for classification of the combined data set by the highest-performing model, random forests, which

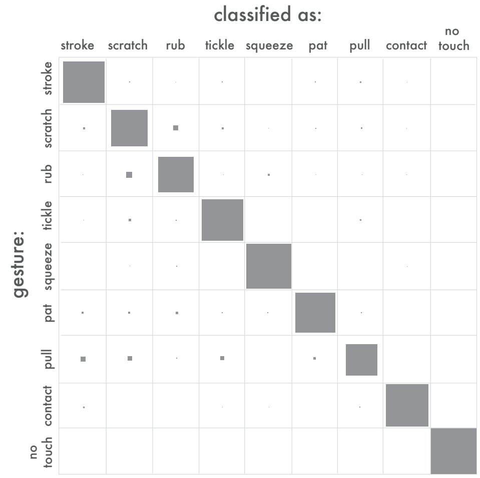

Figure 5. Hinton diagram of the confusion matrix corresponding to the random forests gesture recognition model, evaluated on the combined data set of 16 study participants. Size of grey squares represents classification: for example, the first row shows how many examples of "stroke" are recognized as *stroke*, how many are mislabeled as *scratch*, and so on for all gesture types.

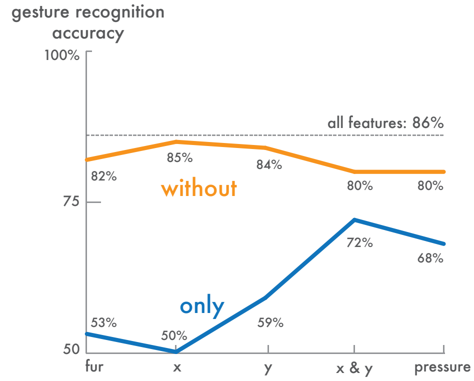

Figure 6. Contributions of individual sensor channels to gesture classification performance of random forests model on combined 16-participant data set. Upper curve shows performance when using features from all data channels *except* conductive fur, then *except* x position, *except* y position, *except* 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).

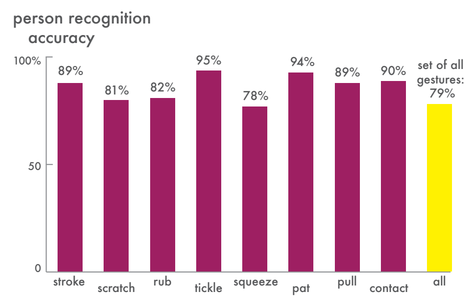

Figure 7. Results of evaluating a random forests model's ability to recognize people from their touch data. Model was trained and tested on *stroke* data for all participants, *scratch* data for all participants, etc for all gestures. Model was also evaluated on combined set of participant data for all gestures.

is 86% accurate. 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. In Figure 6 we report the relative drop in performance when *leaving out* any given channel, and when relying on *only* a given channel. Finally, Figure 7 reports the results of training a model to recognize specific people from their touch gesture data.

DISCUSSION

We evaluated our prototype's 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.

As shown in Figure 3, a model trained on a single person's touch data can generally achieve very high recognition performance. Accuracy is almost always above 90%, on average higher than 95%, and quite consistent across classifiers.

As might be expected, modelling all participants together decreases gesture recognition accuracy (Figure 4). The effect of model choice also becomes more evident, with performance among classifiers ranging from 68-86%. This result suggests that touch behaviours are by no means universal. 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 5, 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*.

We also attempt to characterize the value of each data channel to prediction of our set of gestures. As can be seen in Figure 6, 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 has the second weakest performance, only stronger than x position. However, performance drops 4% when the fur is removed, more than the performance drops corresponding to either x position or y position (1% and 2%). This result suggests that the fur sensor contributes a relatively orthogonal channel of information, valuable to recognition. Removing pressure information also has a big negative impact on performance. However, Schmeder and Freed [19] 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 does not interfere with the relative evident individual contributions of x position data and y position data.

Finally, we examine our model's ability to recognize an individual person from her touch interaction. As shown in Figure 7, 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 executed the *tickle* gesture.

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 above, 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.

In terms of practical value, 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 means, choosing a performance goal 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 and dislikes, which might not always cater strictly to audience needs and expectations. Here, in addition to examining standard recognition rates, we would also evaluate how easy it is for the audience to communicate effectively with the system, and how valuable the resulting interactions are. 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. We believe this level of performance would enable the system to: a) perceive and respond in a large number of complex, subtle ways, b) capture a significant majority of audience expectations, and c) express a "personality" and sense of unpredictability that might make for a compelling interaction in both the short and long term.

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 material properties. We will also compare 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 will 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, including time-series-specific classification routines. 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. Future systems can incorporate confidence estimates, so that the model is not forced to classify all data in terms of known gesture types, even if probability is 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 transition could involve buffering in 2-second periods, and perhaps running several buffers of different lengths simultaneously to help prevent shorter gestures from getting lost in too-long windows, as suggested in Chan, et. al. [3].

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 [4, 25]. In unsupervised learning, the model is shown lots of data and searches for hidden structure, for instance using clustering or dimensionality reduction [4]. In reinforcement learning, the system would learn by interacting with its environment, and observing the resulting feedback, such as with a Markov decision process [25]. This type of approach would be more complex than a 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.

In conclusion, based on the original objectives of this research, these early positive 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". These capabilities could enable an affective touch-based system to better engage with its audience on an emotional level, informing comforting and therapeutic interactions. We hope this work 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.

ACKNOWLEDGEMENTS

We gratefully acknowledge the GRAND NSERC Network Center for Excellence, which provided partial support for this work.

REFERENCES

1. C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
2. P. Brockwell and R. Davis. Introduction to time series and forecasting. In *Springer-Verlag, New York*, 2002.
3. J. Chang, K. E. MacLean, and S. Yohanan. Gesture recognition in the haptic creature. In *EuroHaptics 2010*, 2010.
4. 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.
5. A. Flagg, D. Tam, K. MacLean, and R. Flagg. Conductive fur sensing for a gesture-aware furry robot. In *IEEE Haptics Symposium*, 2012.
6. 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.
7. 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.
8. Interlink Forse Sensing Resistors, 2012. `http://www.interlinkelectronics.com/products.php`.
9. K. Goris, J. Saldien, I. Vanderniepen, and D. Lefeber. The huggable robot probo, a multi-disciplinary research platform. In *Eurobot '08*, 2008.
10. T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, second edition, 2009.
11. Meka. Dextrous Mobile Manipulators and Humanoids, 2012. `http://mekabot.com/?ref=Hizook-spotlight1`.
12. H. Perner-Wilson. Stroke sensor. Plusea, www.plusea.at, May 2011.
13. R. Picard. *Affective Computing*. MIT Media Laboratory, 1997.
14. PJRC. Teensy USB Development Board, 2012. `http://www.pjrc.com/teensy/`.
15. Peratech Ltd. Quantum Tunnelling Composites, 2012. `http://www.peratech.com/qtcmaterial.php`.
16. C. Rafael and S. D'Melle. Affect detection: An interdisciplinary review of models, methods, and their applications. In *IEEE Transactions on Affective Computing*, 2010.
17. Roboskin. Skin-based Technologies and Capabilities for Safe, Autonomous and Interactive Robots, 2012. `http://www.roboskin.eu/index.php/home`.
18. 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.
19. 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.
20. 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.
21. T. Shibata, K. Inoue, and R. Irie. Emotional robot for intelligent system: Artificial emotional creature project. In *Proceedings of IIZUKA*, pages 43–48, 2006.
22. 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.
23. R. Singh, T. Jaakkola, and A. Mohammad. *Machine Learning*. MIT OpenCourseWare, 2006.
24. 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.
25. R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
26. I. H. Witten and E. Frank. *Data mining: Practical machine learning tools and techniques*. Morgan Kaufmann Publishers, second edition, 2005.
27. 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.
28. 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.
29. 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.
30. 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.