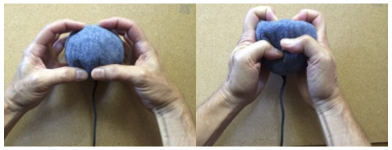

(a) Squeezing

(b) Stroking

(c) Tapping, Punching

Sidebar 1: Sample touch input

Translating Affective Touch into Text

Daniel Shapiro, Zeping Zhan, Peter Cottrell, Katherine Isbister
UCSC
Santa Cruz, CA, USA
(dgs,pcottrel,kisbiste)@ucsc.edu,zzhan@nyu.edu

ABSTRACT

This paper presents a game-like experience that translates tactile input into text, which captures the emotional qualities of that touch. We describe the experience and the system that generates it: a plush toy instrumented with pressure sensors, a machine learning method that acquires a mapping from touch data into a feature vector of affect values, and a mechanism that transcribes that feature vector into text. We conclude by discussing the range of novel interactions that such a nuanced tactile interface can support.

KEYWORDS

Affective touch, affect display, haptic input, gestural input, fidgets, translation, synesthesia

ACM Reference Format:
Daniel Shapiro, Zeping Zhan, Peter Cottrell, Katherine Isbister. 2019. Translating Affective Touch into Text. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI'19 Extended Abstracts), May 4–9, 2019, Glasgow, Scotland UK*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3290607.3313015

INTRODUCTION

Fidget objects have become massively popular in recent years. In May of 2017 every one of the top 10 selling toys on Amazon was a form of spinner [8], while popular sites recommended stress

CHI'19 Extended Abstracts, May 4–9, 2019, Glasgow, Scotland UK

ACM ISBN 978-1-4503-5971-9/19/05.
https://doi.org/10.1145/3290607.3313015

Physical action	Output lyric
Holding + shaking	Control through fear A reign of terror is here Control through fear The terrors powers made clear
Slow stroking	Baby, if you want to, be my lover You better take me home Cause it's a long long way to paradise And I'm still on my own.
Light squeeze	Clouds of fear Permeate All I see Heightening Nervousness Threatens me

Sidebar 2: Sample I/O of the touch-to-text prototype

Sidebar 3: The Fidget Ball

balls and worry stones to children and adults as a means of managing stress and improving focus [14]. Manipulating hand-held objects with the right kinesthetic and tactile affordances seems to satisfy an important need in day-to-day life. This paper explores the potential of such sensor-enabled, computationally enhanced fidget objects to impact digital interactions. We focus on two questions:

- Can a computational system discern affect from touch traces upon fidget objects?
- Can we illustrate affective tactile input via a game?

We know intuitively that touch carries affect, but we have limited insight into the component problems of extracting affect from tactile data. We cannot enumerate the features of tactile data that carry affect, or directly interpret the affective content in tactile data. However, recent advances in deep learning hold promise for establishing connections between touch traces and affect.

This paper presents first results from a system of this kind. The performance component accepts input from a soft fidget object instrumented with pressure sensors, and outputs a song lyric that expresses the emotional content of that touch. The underlying learning system inputs a body of touch data with associated textual descriptions obtained from people, processes the text into an affect vector via sentiment analysis tools, and acquires a mapping from touch to affect via deep learning. The system outputs a lyric from a database whose affect most closely matches the user's touch.

RELATED WORK

A fundamental appreciation that touch carries affect dates back to work in the 1970s [2]. More recent research examines the relation between touch and affect as mediated by tactile input devices, affective gestures, and gesture recognition systems. Work on devices has examined the design space of smart materials for capturing affective touch [1, 12], and of fidgets that afford affective touch [13]. Research focused on affective gestures includes the itemization of touch types common in human-human and human-animal interactions [15] (including analysis of the role of those gestures in human-robot communication), and work that links hand-held fidget use with emotional self-regulation [9].

A number of authors have examined the problem of recognizing affect from touch traces, largely in the context of keyboard strokes and touchpad patterns with mobile devices [3, 6, 7, 10, 11]. Some of this work has recognized specific gestures using plush toys instrumented with sensors [4], which is closer to the setting we explore. However, to the best of our knowledge, no one has yet built an end-to-end system that accepts touch traces as input, identifies affective characteristics in that touch, and uses that analysis to generate output.

THE EXPERIENCE

Interaction with the touch-to-text system is simple. The user picks up the fidget ball, which is tethered to a laptop, and manipulates it for 10 seconds (an interval chosen somewhat arbitrarily to span a single intent). The system responds by printing out a song lyric that reflects the affective content of

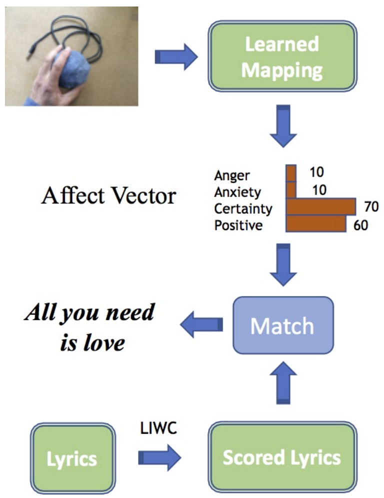

Sidebar 4: The Performance System

Sidebar 5: Dimensions of Affect Recognized by the Touch To Text System

- anger
- anxiety
- certainty
- positive emotion
- sadness
- tentativeness

that touch. The physical interaction can vary from fairly gentle to forceful, slow to rapid, and can repeat or change continuously. The user's motions are unconstrained. Sidebar 1 illustrates several.

The system's output follows the character of its input. For example, scratching, twisting, squeezing, tapping, punching, holding, drumming, massaging, and rubbing all generate distinct lyrics. Sidebar 2 presents several examples. Different executions of the same action produce distinct output, while combinations and interpolations between gestures produce text containing both emotional notes.

THE PHYSICAL SYSTEM

Sidebar 3 illustrates the fidget ball. Its skeletal structure consists of a 9mm in diameter silicon ball that can be deformed with minimal pressure. The ball is covered in a double layer of fabric, with six 18.5 mm circular capacitive force-sensitive resistors inserted between the layers and distributed symmetrically around the ball's surface. The sensors are connected to an Ardunio based Adafruit Feather M0, with a RTC board attached to provide time-stamping abilities. The ball's surface is only disturbed by a tethering USB cord at the "bottom", which enables data collection and provides power.

The fidget ball outputs a vector of 6 numerical values from 0-1023, which correlate to the amount of voltage passing through the corresponding sensor; 0 is an untouched sensor and 1023 is a heavily compressed sensor. The top layer of fabric is a navy colored, soft felted cotton typically used in baby blankets. This color and texture were chosen to minimize visual and tactile stimulation that might bias user manipulations.

TRANSFORMING TOUCH TO TEXT

Our mechanism for translating touch into text has two components: a performance system that generates text from tactile input, and a learning system that acquires the required mappings from training data. We describe each in turn.

The Performance System

Sidebar 4 illustrates the architecture of the performance system, which contains two steps:

(1) Collect ~10 seconds of raw touch data from the user's manipulation of the fidget ball and transform it into a feature vector of affect data.
(2) Realize that affect vector as text by matching it onto a database of {lyric, affect vector} pairs.

The first step transforms user input into an affect vector via a mapping acquired by the learning system, described below. The second step realizes that affect vector as text via a simple retrieval strategy. It applies the Linguistic Inquiry and Word Count package (LIWC) to process ~30,000 individual stanzas from a database of song lyrics into affect vectors, V. It extracts the 6 fields identified in Sidebar 5 (from the 55 output by LIWC) and retrieves the lyric that most closely matches the affect in the user's touch sample. That match is calculated with even weights on each 0-99 valued field:

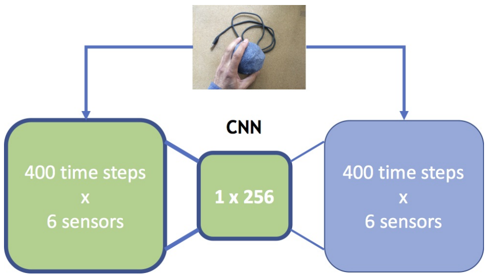

Sidebar 6: The learning system employs an auto-encoder to extract a compressed representation of touch.

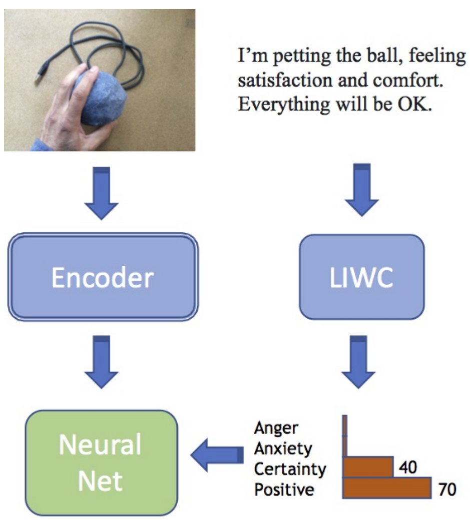

Sidebar 7: Stage 2 of the learning system trains a neural net to map compressed touch data to an affect vector.

$$argmin_{l \in \text{lyrics}} \sum_{i in 1..6} (V_i(l) - V_i(usertouch))^2 \tag{1}$$

The choice of fields is the result of a engineering compromise; they must simultaneously represent affect that the fidget ball can plausibly detect, and affect present in the textual descriptions of touch employed as supervised data by the learning system. The specific fields are less important than the net behavior, since our primary goal is to demonstrate the feasibility of transforming touch into text.

The Learning System

The learning component of the touch to text system acquires two mappings: a compact embedding of the user's touch signal, and a mapping from that representation to a vector of affect values. Learning occurs sequentially, across these two stages.

Sidebar 6 illustrates the first stage, which acquires the touch embedding. It employs a common auto-encoder design that seeks to predict the input touch signal, while forcing the information to pass through a ~10:1 compression. We employed a CNN with three layers of convolution and max-pooling pairs in the encoder, and three layers of up-sampling and convolution pairs in the decoder. The encoder outputs a 1x256 vector representing 10 seconds of user touch.

Sidebar 7 illustrates the second stage architecture, which acquires the touch to affect mapping. It inputs supervised data in the form of {touch, text} pairs, where the text describes the affective quality of the touch. It passes the text through LIWC to produce a 1x6 affect vector, and compresses the touch data into a 1x256 vector via the encoder described above. Then, given the low dimensionality of this data, we train a Neural Net with 3 densely connected layers to predict affect from embedded touch.

Training

We trained the learning system with data collected from a user study of 16 participants, who performed a touch on the fidget ball while describing their touch in words. We solicited {affective touch, affective text} pairs via prompts encoding gestural tasks, emotional communication tasks, and situation interpretation tasks (as illustrated by category in Sidebar 8).

We trained the 1st stage autoencoder with 30,000 synthetic touch samples of punching, stroking, and squeezing of varying amplitudes, augmented with 20,000 samples of unconstrained experimenter generated touch, extracted from 15 minutes of touch partitioned via rolling windows with 1/25th second offset. We added 110,000 samples obtained from 80 minutes of touch data from study participants partitioned in the same way. Of the resulting 160,000 sample corpus, we employed 75% for training the 1st stage CNN, and 25% for testing. We trained the 2nd stage neural net on 76,000 supervised {touch, text} pairs obtained from the user study. Within this set, we employed 90% of the data for training, and 10% for testing the second stage model during learning. We evaluated the two learned

Prompt type	Description
Gestural	Duplicate demonstrated gesture: a light squeeze, pressing fidget between fingers of both hands.
Task	Your pet obeyed a command: Praise it
Interpretation	Translate this lyric into touch: Good day sunshine Good day sunshine Good day sunshine

Sidebar 8: Sample prompts from the user study.

	Light Squeeze	Hold and Shake	Slow Stroke
Posemotion	0	0	10.85
Anxiety	34.25	26.25	0
Anger	0	0	0.63
Sad	0	0	0
Tentative	0	0	3.42
Certain	12.89	4.05	0

Sidebar 9: Probing the touch to affect map

models in combination (autoencoder and embedded touch to affect map) on a held-back validation set of ~8,000 {touch, text} pairs. The mean squared error after training was 7.76; an average difference between prediction and ground truth of 2.78 on each dimension of affect (a 0-99 scale).

RESULTS AND DISCUSSION

This paper provides a feasibility demonstration of a novel capability to transform affective touch into text. As illustrated in sidebars 1 and 2, the prototype accepts a wide range of touch inputs and responds to multiple dimensions of affect in touch. Sidebar 9 shows that the learned mapping captures degrees of expression in that touch. The underlying deep learning mechanism obviates the need to engineer features for analyzing touch, and automates its interpretation as affect. Our approach of employing an affect vector as an intermediate representation is relevant wherever touch and textual descriptions of desired output occur. For example, future learning applications might transform affective touch into music or visual art, map affective touch to diagnostic and therapeutic responses, or employ affective touch to parameterize game interactions.

That said, the current prototype has several limitations. First, the user experience of translating touch into lyrics is bounded. This I/O is consistent with a demonstration of principle, but a more engaging interaction that responds to affective touch vs mirrors it is easily possible. The palette of 6 affect features is also limiting, although it can be expanded with more nuanced textual analysis tools and alternate input devices. Accelerometers would let the system respond better to the tossing and shaking we observed, while suggestive form factors could elicit more specialized touch, as with Skweezees [5] that invite deformation, or a tribble with piezoelectric fur that invites petting [4]. Third, our use of deep learning requires a significant amount of supervised touch data that is difficult to obtain. The common approach of adapting data from analogous problems is less available here as our device is custom made, although it can be replicated. Finally, while this work acquires and evaluates a mapping from touch to text in machine learning terms, it lacks end-to-end validation showing that the inferred affect is in fact what the user intended. This can be addressed via further user study.

More broadly, this work raises interesting questions about the vocabulary of touch. We have shown that some affective expressions hold across individuals, but we also observed differences in the magnitude of participants' touch and their selection of gestures, especially in response to task prompts. This suggests an opportunity for learning systems to acquire and utilize personalized vocabularies of touch, as well as a challenge to separate the functional and affective components of touch signals.

In summary, this work demonstrates the feasibility of utilizing affective touch in a computational system in order to make interaction increasingly human centered, engaging, and visceral. From that perspective, translating touch to text is one step of a wider exploration to employ the nuances of touch for communication, interaction, and control.

REFERENCES

[1] Heber Cruz Zurian Seyed Reza Atefi Erik Billing Fernando Seoane Martinez Paul Lukowicz Bo Zhou, Carlos Andres Velez Altamirano. 2017. Textile Pressure Mapping Sensor for Emotional Touch Detection in Human-Robot Interaction. *Sensors* 17, 11 (2017).

[2] Manfred Clynes. 1978. *Sentics: The Touch of the Emotions.* Anchor Press/Doubleday.

[3] Clayton Epp, Michael Lippold, and Regan L. Mandryk. 2011. Identifying Emotional States Using Keystroke Dynamics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 715–724. https://doi.org/10.1145/1978942.1979046

[4] Anna Flagg and Karon MacLean. 2013. Affective Touch Gesture Recognition for a Furry Zoomorphic Machine. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction (TEI '13)*. ACM, New York, NY, USA, 25–32. https://doi.org/10.1145/2460625.2460629

[5] Luc Geurts, Jolien Deville, Vero Vanden Abeele, Jelle Saldien, and Karen Vanderloock. 2013. Skweezee Studio: Turn Your Own Plush Toys into Interactive Squeezable Objects. In *Proceedings of the 8th International Conference on Tangible, Embedded and Embodied Interaction (TEI '14)*. ACM, New York, NY, USA, 377–380. https://doi.org/10.1145/2540930.2567896

[6] Surjya Ghosh, Niloy Ganguly, Bivas Mitra, and Pradipta De. 2017. TapSense: combining self-report patterns and typing characteristics for smartphone based emotion detection. 1–12.

[7] Javier Hernandez, Pablo Paredes, Asta Roseway, and Mary Czerwinski. 2014. Under Pressure: Sensing Stress of Computer Users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 51–60. https://doi.org/10.1145/2556288.2557165

[8] Katherine Isbister. 2017. Fidget Toys Aren't Just Hype. Retrieved August 7, 2018 from https://www.scientificamerican.com/article/fidget-toys-arent-just-hype/

[9] Michael Karlesky and Katherine Isbister. 2016. Understanding Fidget Widgets: Exploring the Design Space of Embodied Self-Regulation. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. https://doi.org/10.1145/2971485.2971557

[10] Hsiao T-C Lee P-M, Tsui W-H. 2015. The Influence of Emotion on Keyboard Typing: An Experimental Study Using Auditory Stimuli. *PLoS ONE* 10, 6 (e0129056 2015). https://doi.org/10.1371/journal.pone.0129056

[11] A.F.M. Nazmul Haque Nahin, Jawad Mohammad Alam, Hasan Mahmud, and Kamrul Hasan. 2014. Identifying emotion by keystroke dynamics and text pattern analysis. *Behaviour & Information Technology* 33, 9 (2014), 987–996. https://doi.org/10.1080/0144929X.2014.907343 arXiv:https://doi.org/10.1080/0144929X.2014.907343

[12] Katherine Isbister Peter Cottrell, April Grow. 2018. Soft-bodied Fidget Toys: A Materials Exploration. In *Proceedings of the Twelfth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, New York, NY, 42–48. https://doi.org/10.1145/3173225.3173266

[13] Katherine Isbister Suzanne B. da Câmara, Rakshit Agrawal. 2018. Identifying Children's Fidget Object Preferences: Toward Exploring the Impacts of Fidgeting and Fidget-Friendly Tangibles. In *DIS '18,*. 301–311. DOI:https://doi.org/10.1145/3196709.3196790

[14] Lexi Walters Wright. [n. d.]. 6 Types of Fun Fidgets for Kids With ADHD. Retrieved August 7, 2018 from https://www.understood.org/en/learning-attention-issues/child-learning-disabilities/add-adhd/6-types-of-fun-fidgets-for-kids-with-adhd?view=slideview

[15] MacLean K.E. Yohanan, S. 2012. The Role of Affective Touch in Human-Robot Interaction: Human Intent and Expectations in Touching the Haptic Creature. *International Journal of Social Robotics* 4 (2012), 163–180. https://doi.org/10.1007/s12369-011-0126-7