Touched by ChatGPT: Using an LLM to Drive Affective Tactile Interaction

Qiaoqiao Ren
Faculty of Engineering and Architecture
IDLab-AIRO, Ghent University – imec
Qiaoqiao.Ren@ugent.be

Tony Belpaeme
Faculty of Engineering and Architecture
IDLab-AIRO, Ghent University – imec
Tony.Belpaeme@ugent.be

arXiv:2501.07224v1 [cs.RO] 13 Jan 2025

***Abstract*—Touch is a fundamental aspect of emotion-rich communication, playing a vital role in human interaction and offering significant potential in human-robot interaction. Previous research has demonstrated that a sparse representation of human touch can effectively convey social tactile signals. However, advances in human-robot tactile interaction remain limited, as many humanoid robots possess simplistic capabilities, such as only opening and closing their hands, restricting nuanced tactile expressions. In this study, we explore how a robot can use sparse representations of tactile vibrations to convey emotions to a person. To achieve this, we developed a wearable sleeve integrated with a 5x5 grid of vibration motors, enabling the robot to communicate diverse tactile emotions and gestures. Using chain prompts within a Large Language Model (LLM), we generated distinct 10-second vibration patterns corresponding to 10 emotions (e.g., happiness, sadness, fear) and 6 touch gestures (e.g., pat, rub, tap). Participants (N = 32) then rated each vibration stimulus based on perceived valence and arousal. People are accurate at recognising intended emotions, a result which aligns with earlier findings. These results highlight the LLM's ability to generate emotional haptic data and effectively convey emotions through tactile signals. By translating complex emotional and tactile expressions into vibratory patterns, this research demonstrates how LLMs can enhance physical interaction between humans and robots.**

***Index Terms*—Tactile interaction, affective computing, emotion classification, gesture classification, large language model.**

I. Introduction

Emotion is central to human communication, influencing interactions, decision-making, and relationships [1]. One promising approach is using haptic feedback, particularly vibration, to convey emotional states. Vibration is an intuitive and non-verbal channel that can convey a range of emotional cues [2], making it a valuable tool for enriching communication in digital systems where the use of traditional sensory modalities, such as vision and sound, may be limited [3].

Haptic devices have long been explored for their role in augmenting emotional communication. Early studies demonstrated the ability of vibrations to convey basic emotions, such as happiness or sadness, through simple intensity modulation [4]. Researchers developed wearable haptic devices to convey emotional states, finding moderate success in user decoding, particularly for high-arousal emotions [5]. Similarly, Salminen [6] used force feedback and skin stretch to evoke emotional responses, demonstrating that tactile communication could enhance immersion and empathy in digital environments. However, these approaches often rely on manually designed patterns, which can limit their scalability and effectiveness in representing complex emotions.

Russell's circumplex model of emotions [7], which organizes emotions along valence (pleasantness) and arousal (intensity) dimensions, provides a useful framework for designing haptic feedback systems. Studies have used this model to map emotions to haptic parameters such as vibration intensity, rhythm, and duration, enabling more structured approaches to haptic communication [8]. However, these efforts often focus on a limited set of emotions and fail to account for the intricacies of emotional expressions or interpersonal variability in emotion decoding.

The application of AI in emotion communication has been explored in visual and auditory modalities [9], but its integration into haptic systems remains relatively underdeveloped [10]. Traditional approaches for generating tactile patterns often rely on non-AI methods, such as heuristic-based systems [11], [12]. These methods have primarily focused on understanding how features –such as vibration frequency, amplitude, and duration– influence participants' perceptions of arousal and valence [13]. While insightful, these approaches face significant limitations, including hardware constraints, differences in the contact areas of vibration motors, variability in motor properties, and challenges in standardizing tactile stimuli across participants [14]. Therefore, it is still a significant challenge to reflect complex emotional expressions accurately using tactile communication. Previous studies have shown that even in human-human communication, decoding emotional signals can be inherently difficult. For example, the accuracy of emotion decoding between individuals in close relationships, such as couples, often falls below 50% for a limited set of emotions [15]. These findings underscore the need for more expressive and reliable haptic systems. Moreover, previous studies have predominantly focused on collecting affective data directly from participants, relying on manually designed patterns that may not capture the full range of emotional nuances. Recent developments in LLMs have introduced new possibilities in this area. LLMs, with their ability to analyse large datasets, offer a powerful tool for understanding the features of emotional expressions and translating them into tactile feedback [16]. For instance, Yu *et al.* [16] demonstrated the potential of tactile-language models

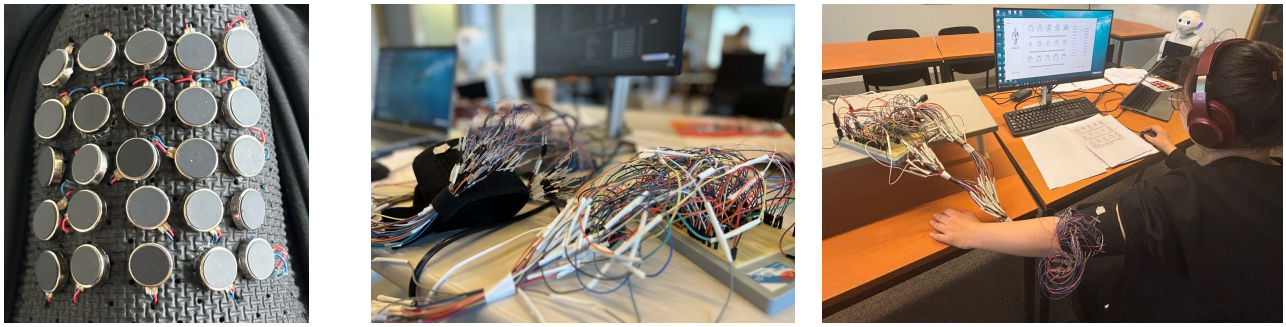

(a) Vibration motors distribution. (b) Vibration sleeve. (c) Convey emotions to the participants.

Fig. 1: Experimental setup.

in object property reasoning, bridging the gap between tactile perception and linguistic representation.

Our study uses a 5x5 grid of vibration motors embedded in a wearable device, designed to convey affective touch. We use an LLM to control the device, a notable and unique contribution. The LLM is first prompted to analyse the tactile features of emotions such as happiness, sadness, fear, and anger, as well as touch gestures like patting and rubbing. Next, the LLM generates corresponding vibration patterns based on the analysis. These patterns are mapped onto the 5x5 grid motors to produce varying intensities, rhythms, and spatial vibration distributions. By integrating LLM-driven analysis into the design process, we aim to create more expressive and dynamic haptic feedback systems capable of conveying various emotions. Our primary research question explores whether a LLM can generate tactile data that effectively conveys emotions to humans.

II. Materials & Methods

To validate the ability of an LLM to generate tactile signals, we collect data on how people interpret these stimuli. Participants are tasked with decoding 10 emotions (anger, fear, disgust, happiness, surprise, sadness, confusion, comfort, calm, and attention) and 6 gestures (hold, pat, tickle, rub, tap, and poke). Before data collection, participants were provided with definitions of all emotions (to ensure consistent interpretations for non-native English-speaking participants) and touch gestures to ensure they understood and agreed with the given definitions. Participants are asked to rate the arousal and valence on a scale from 1 to 10 for each stimulus generated by the LLM. Subjective ratings of arousal and valence are given on a 10-point scale, with lower scores indicating lower arousal and less pleasant stimuli, and higher scores reflecting higher arousal and more pleasant stimuli. Afterwards, they are asked to classify the stimulus as expressing a specific emotion or gesture. The following sections provide details on the equipment, the study design, and the data acquisition.

A. Equipment

1) Vibration sleeve: The hardware for the vibration device consists of a 5x5 grid of vibration motors embedded in an upper arm sleeve, powered by a 3.3V power supply. Each motor is controlled by a single Raspberry Pi using BC557B PNP transistors, which act as switches to regulate current flow to the motors, as shown in Fig. 1a and Fig. 1b. The system uses pulse-width modulation (PWM) for precise control over the intensity of each motor's vibration, which allows for modulation of the vibration intensity by rapidly switching the transistors on and off at varying duty cycles. By adjusting the duty cycle of the PWM signal, the perceived intensity of each motor's vibration can be finely controlled, allowing for a dynamic range of vibration patterns.

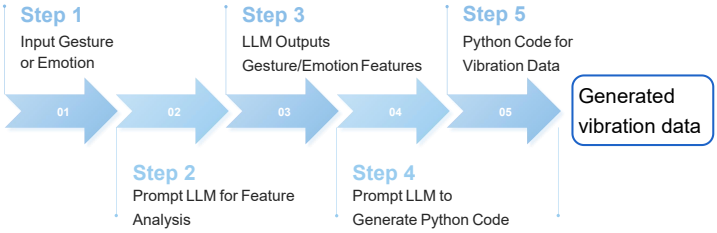

Fig. 2: Flow chart of LLM generated vibration data.

2) Prompts: This approach uses a chain of prompts to guide an LLM (*gpt-4o*) to generate vibration data for a 5x5 grid of vibration motors. Inspired by previous work where 10-second data was collected from humans and transformed into vibration patterns to convey emotions [17], we asked the LLM to generate 10 seconds of vibration data. The vibration data generation flow chart is shown in Fig. 2. The process begins by prompting the LLM to analyze the features of a specific gesture or emotion (e.g., happiness, sadness, rub, tap). The LLM is asked how such an emotion or gesture would manifest in terms of pressure on the grid, focusing on how the vibration magnitude changes over time and across the spatial arrangement of vibration motors. The prompt also directs the model to account for smooth transitions in pressure values, reflecting the natural progression of touch or emotional cues. This ensures a natural onset and decay of the motors.

Once the analysis is complete, a second prompt uses the insights from the first to generate Python code. This code simulates the vibration data by creating a data frame, capturing the evolving vibration patterns over time. The Python code

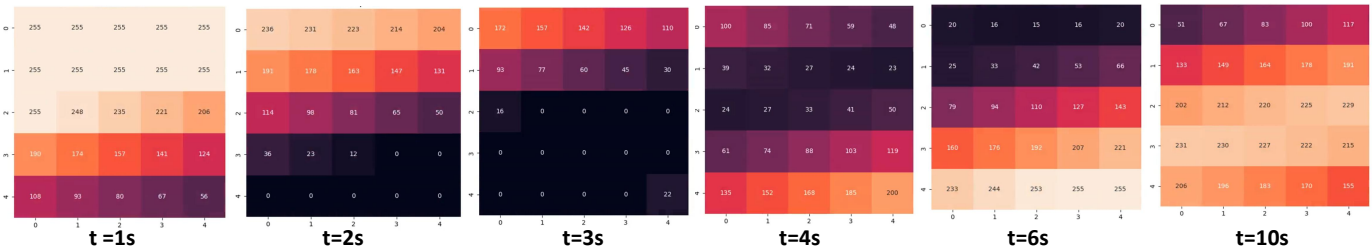

Fig. 3: Illustration of frames of LLM generated *rub* gesture.

ensures that the data is in a format that can be used directly to control the 5x5 grid of vibration motors. Key requirements like smooth pressure transitions and sustained changes over multiple time steps are emphasized to ensure the data is suitable for producing realistic haptic feedback. The output of this process is a CSV file for 10 seconds of vibration data. The motor distributions and visualization of LLM-generated data are shown in Fig. 1a and Fig. 3. The detailed prompts, code and visualization can be found in our GitHub repository[1].

III. Experiment Setup

To explore emotion decoding via touch-based expression and reveal if and how people can decode the emotions and gestures generated by the LLM, we set up a data collection experiment (Fig. 1c). We recruited 32 participants on a university campus, 17 identified as female, 15 as male, average age 27.5 ± 3.6 years. The data collection and study complied with the ethics regulations of *ANONYMOUS* and participants gave informed consent before data collection. The collection was split into two sessions: in the first session they *decode emotions conveyed by the robot*, in the second session they *decode tactile gestures conveyed from the robot*.

1) *Pre-session:* Participants were fitted with the vibration device and their perception threshold was set.
2) *Decoding Emotions from the Robot:* Participants were asked to interpret emotions from 10 stimuli delivered by the robot through the wearable device. For each stimulus, they rated arousal and valence and identified the specific emotion. Each stimulus lasted 10 seconds, with time provided between stimuli for participants to provide their responses. The next stimulus was offered once participants indicated they were ready.
3) *Decoding Tactile Gestures from the Robot:* In this session, participants decoded touch gestures from 6 stimuli delivered by the robot through the wearable device. For each stimulus, they identified the corresponding gesture and could replay it as needed until they were sure of their answer. Each stimulus lasted 10 seconds, with preparation time given between stimuli. As before, the next stimulus was offered when participants indicated they were ready.

[1]https://github.com/TouchedByLLM/LLM_generate_data

IV. Results and Discussion

A. Emotion decoding

We averaged participants' arousal and valence ratings for different emotional stimuli, which are shown in Table. I. Figure. 4 illustrates the distribution of the 10 emotions within the emotion complex model. Participants on average placed *happiness* in the high-arousal, positive-valence quadrant. *Disgust*, *fear*, *surprise* and *anger* were placed in the high-arousal, negative-valence quadrant. *Confusion* and *sadness* were assigned to the low-arousal, negative-valence quadrant, while *comfort* and *calming* were grouped in the low-arousal, positive-valence quadrant. This classification is consistent with the emotion categorisation described in Russell's complex model and previous researches. [18], [19].

The classification results in Fig. 5a showed that 22 out of 32 participants were correctly classified *anger*, making it the most accurately decoded emotion. This was followed by *comfort*, which 13 participants decoded correctly, and *calming*, decoded correctly by 12 participants. The overall decoding accuracy across the 10 emotions is 30.3%. Previous research [20] compared participant responses to the probability of randomly guessing correctly (the "chance level"). In our experiment, the chance level for emotions is 10%. Previous research [15], [21] explored emotion decoding among couples and strangers, and the results we obtained in our research are similar to the results they presented for strangers (37.5%). If we set the 37.5% as a baseline, we used a one-sample t-test to determine whether the participant's decoding accuracy is significantly higher than chance level; the results showed that there is

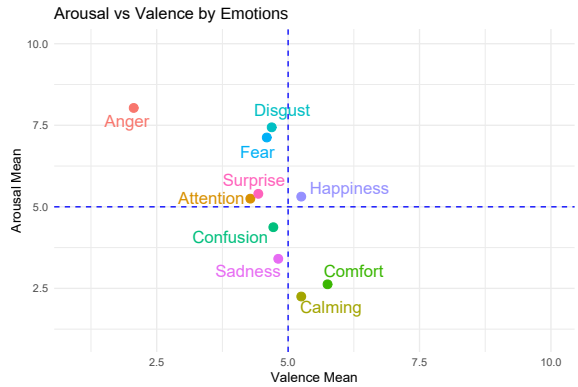

Fig. 4: Arousal and valence rating distribution of emotions.

TABLE I: Arousal and valence of different emotions and decoding accuracy (%)

Emotions	Happiness	Surprise	Fear	Disgust	Anger	Comfort	Attention	Calming	Confusion	Sadness
Arousal	5.3±2.0	5.3 ± 2.2	7.1 ± 1.9	7.4±1.7	8.0 ± 1.5	2.6±1.9	5.3 ± 1.9	2.3 ± 1.9	4.4 ± 2.1	3.4±2.1
Valence	5.3±2.3	4.4 ± 1.9	4.6 ± 2.8	4.7±2.5	2.1 ± 1.4	5.8±1.5	4.3 ± 1.8	5.3 ± 2.0	4.7 ± 2.3	4.8±2.2
Accuracy	21.9	**25.0**	**31.2**	21.9	**68.8**	**25.0**	**31.2**	**37.5**	12.5	28.1
Sig.(p)	0.06	0.03	< 0.01	0.06	< 0.01	0.03	< 0.01	< 0.01	0.42	0.28

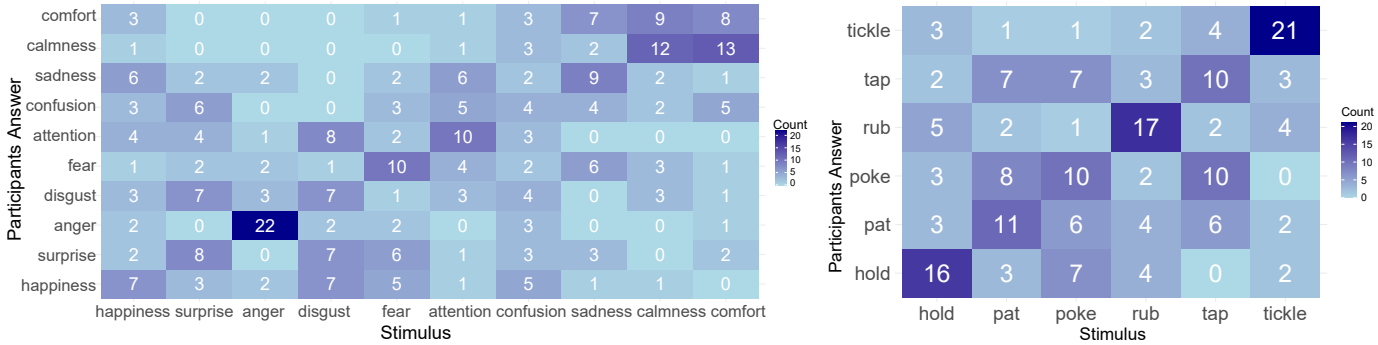

(a) Confusion matrix for emotions

(b) Confusion matrix for touch gestures

Fig. 5: Confusion matrix.

no significant difference between human-to-human emotion decoding among strangers in previous research and our human-robot decoding ($t(31) = -2.99, p = 0.99$). And Table. I presents the decoding accuracy for each emotion, along with its statistical significance level, in comparison to the chance level. In addition, compared to a chance level of 10%, the participants' decoding results (30.3%) are significantly above chance ($t(31) = 7.89, p < 0.001$). Therefore, we can conclude that the participants succeeded in decoding the emotional data generated by the LLM, and the LLM has the ability to generate affective touch.

TABLE II: Decoding accuracy(%)

Gestures	Hold	Pat	Poke	Rub	Tap	Tickle
Accuracy	53.1	34.4	31.2	53.1	31.2	65.6
Sig.(p)	< 0.01	0.02	0.045	0.01	0.045	< 0.01

B. Gesture decoding

As shown in the Table. II and Fig. 5b, the gesture that participants found easiest to decode was *tickle*, followed by *rub*. In contrast, *poke*, *pat* and *tap* were more frequently misclassified and often confused with each other. One possible explanation for this confusion is the limited size of the designed sleeve's touch-sensitive area. Although the 25-motor setup aims to facilitate relative touch perception, the small touch area (6cm by 6cm) on the upper arm may hinder participants from distinguishing significant differences between certain gestures. For instance, while *poke* involves a smaller contact area compared to *pat*, their rhythmic patterns can be similar, as is the case with *tap*. This overlap in tactile cues may lead to confusion and reduce the fidelity of gesture recognition when conveying gestures to participants. The participants' touch gesture decoding accuracy is significantly higher than the chance level of 16.7% ($t(31) = 7.67, p < 0.001$). Overall, participants decoded all touch gestures significantly higher than the chance level, while they could only decode six emotions significantly higher than the chance level.

C. Conclusion and future work

This study highlights the potential of using a wearable sleeve with a 5x5 grid of vibration motors, combined with LLM-generated vibration patterns, to effectively convey emotions and touch gestures. Participants successfully decoded 10 emotions and 6 gestures with accuracy levels that matched prior research in human-human tactile interaction and were significantly higher than chance. Notably, the results demonstrate that LLMs, despite never having been trained on tactile interaction experiences, can generate affective touch data purely based on textual information and descriptions. However, some emotions, such as *confusion*, did not perform significantly higher than the chance level. This is reasonable, as certain emotions are inherently more challenging to convey purely through touch; most participants also mentioned that they feel confusion and disgust are more challenging to decode than other emotions. These emotions often rely on nuanced contextual or introspective cues that are difficult to translate into tactile patterns alone. In addition, some gestures, such as *poke*, *pat*, and *tap*, posed challenges due to the limited tactile area of the sleeve, which may have hindered participants' ability to distinguish between subtle variations in the contact area and rhythm. This suggests opportunities for enhancing the design to improve gesture differentiation. The findings validate the system's capacity to translate complex emotional and tactile expressions by using LLM into vibratory signals, significantly advancing physical interaction between humans and robots. Future work will aim to refine the tactile interface, enhance gesture distinctiveness, and expand the repertoire of emotional and haptic patterns.

REFERENCES

[1] G. Assunção, B. Patrão, M. Castelo-Branco, and P. Menezes, "An overview of emotion in artificial intelligence," *IEEE Transactions on Artificial Intelligence*, vol. 3, no. 6, pp. 867–886, 2022.

[2] A. Mazzoni and N. Bryan-Kinns, "How does it feel like? an exploratory study of a prototype system to convey emotion through haptic wearable devices," in *2015 7th International Conference on Intelligent Technologies for Interactive Entertainment (INTETAIN)*, pp. 64–68, IEEE, 2015.

[3] J. M. Loomis and R. Klatzky, "Functional equivalence of spatial representations from vision, touch, and hearing: Relevance for sensory substitution," in *Blindness and brain plasticity in navigation and object perception*, pp. 167–196, Psychology Press, 2007.

[4] M. A. Eid and H. Al Osman, "Affective haptics: Current research and future directions," *IEEE Access*, vol. 4, pp. 26–40, 2015.

[5] T. Olugbade, L. He, P. Maiolino, D. Heylen, and N. Bianchi-Berthouze, "Touch technology in affective human–, robot–, and virtual–human interactions: A survey," *Proceedings of the IEEE*, vol. 111, no. 10, pp. 1333–1354, 2023.

[6] K. Salminen, *Emotional Responses to Friction-based, Vibrotactile, and Thermal Stimuli*. PhD thesis, School of Information Sciences, University of Tampere, 2015.

[7] J. A. Russell, "A circumplex model of affect.," *Journal of personality and social psychology*, vol. 39, no. 6, p. 1161, 1980.

[8] M. Teyssier, G. Bailly, C. Pelachaud, and E. Lecolinet, "Conveying emotions through device-initiated touch," *IEEE Transactions on Affective Computing*, vol. 13, no. 3, pp. 1477–1488, 2020.

[9] N. J. Shoumy, L.-M. Ang, K. P. Seng, D. M. Rahaman, and T. Zia, "Multimodal big data affective analytics: A comprehensive survey using text, audio, visual and physiological signals," *Journal of Network and Computer Applications*, vol. 149, p. 102447, 2020.

[10] G. S. Giri, Y. Maddahi, and K. Zareinia, "An application-based review of haptics technology," *Robotics*, vol. 10, no. 1, p. 29, 2021.

[11] B. Saket, C. Prasojo, Y. Huang, and S. Zhao, "Designing an effective vibration-based notification interface for mobile phones," in *Proceedings of the 2013 conference on Computer supported cooperative work*, pp. 149–1504, 2013.

[12] D. Mathew, "vsmileys: Imaging emotions through vibration patterns," *Alternative Access: Feeling and Games*, vol. 2005, pp. 75–80, 2005.

[13] H. Seifi and K. E. Maclean, "A first look at individuals' affective ratings of vibrations," in *2013 World Haptics Conference (WHC)*, pp. 605–610, IEEE, 2013.

[14] O. Schneider, K. MacLean, C. Swindells, and K. Booth, "Haptic experience design: What hapticians do and where they need help," *International Journal of Human-Computer Studies*, vol. 107, pp. 5–21, 2017.

[15] E. H. Thompson and J. A. Hampton, "The effect of relationship status on communicating emotions through touch," *Cognition and Emotion*, vol. 25, no. 2, pp. 295–306, 2011.

[16] S. Yu, K. Lin, A. Xiao, J. Duan, and H. Soh, "Octopi: Object property reasoning with large tactile-language models," *arXiv preprint arXiv:2405.02794*, 2024.

[17] Y. Ju, D. Zheng, D. Hynds, G. Chernyshov, K. Kunze, and K. Minamizawa, "Haptic empathy: Conveying emotional meaning through vibrotactile feedback," in *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–7, 2021.

[18] M. K. Noordewier and S. M. Breugelmans, "On the valence of surprise," *Cognition & emotion*, vol. 27, no. 7, pp. 1326–1334, 2013.

[19] M. M. Bradley, M. Codispoti, B. N. Cuthbert, and P. J. Lang, "Emotion and motivation i: defensive and appetitive reactions in picture processing.," *Emotion*, vol. 1, no. 3, p. 276, 2001.

[20] L. Nummenmaa, E. Glerean, R. Hari, and J. K. Hietanen, "Bodily maps of emotions," *Proceedings of the National Academy of Sciences*, vol. 111, no. 2, pp. 646–651, 2014.

[21] M. J. Hertenstein, D. Keltner, B. App, B. A. Bulleit, and A. R. Jaskolka, "Touch communicates distinct emotions.," *Emotion*, vol. 6, no. 3, p. 528, 2006.