Generative AI for Affective Vibration: Human-Centered Evaluation of LLM and VAE Models

Xiang Li
University of the Arts London
London, United Kingdom
x.li0220231@arts.ac.uk

Nick Bryan-Kinns
University of the Arts London
London, United Kingdom
n.bryankinns@arts.ac.uk

Abstract

Haptic feedback plays a significant role in expressing emotions, however, there is a lack of research on haptics compared with visual and audio channels. In this paper we investigate AI and machine learning methodologies for generating affective vibrotactile feedback. Two generative AI (GenAI) approaches to vibration generation were examined using a custom dataset of vibration-emotion pairings: a Variational Autoencoder (VAE) approach and a fine-tuned large language model (LLM) approach. A quantitative user study involving 15 people validated the GenAIs' capabilities to generate vibrations conveying a range of levels of emotion valence or arousal. Subjective interviews were conducted afterwards which provided valuable insights for multimodal interaction design and future research topics of affective haptics.

CCS Concepts

• **Human-centered computing** → **Haptic devices**; **User studies**;
• **Computing methodologies** → *Artificial intelligence.*

Keywords

Affective Haptics, Large Language Model, Variational Autoencoder, Generative AI, Human-Computer Interaction

ACM Reference Format:
Xiang Li and Nick Bryan-Kinns. 2025. Generative AI for Affective Vibration: Human-Centered Evaluation of LLM and VAE Models. In *Creativity and Cognition (C&C '25), June 23–25, 2025, Virtual, United Kingdom.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3698061.3734410

1 Introduction

Haptic feedback plays a significant role in communicating emotions [8, 19, 39] and yet whilst generative AI has been widely applied to generating affective visual and audio feedback [6, 24, 34], AI-generated vibrations remain unexplored. One reason for this may be a lack of clear methodologies. Our primary goal in this paper is to explore and evaluate different approaches for generating emotional vibrations with AI from a human-centered perspective, thus fostering creativity in haptic feedback and shifting the focus on emotional communication to this sensory modality. This includes three major components: i) construction of a vibration-emotion dataset; ii) exploration of generative AI methods for vibration; iii) a human-centered study to provide insights for designing computational affective haptic feedback.

C&C '25, Virtual, United Kingdom

ACM ISBN 979-8-4007-1289-0/25/06
https://doi.org/10.1145/3698061.3734410

2 Related Work

Haptic feedback encompasses tactile inputs across the body, including kinesthetic, thermal, force, and vibrotactile feedback [19, 36] which can be passive or active forms [10, 35]. In this paper we focus on active vibrotactile feedback to convey emotion, also referred to as affective vibration.

2.1 Design and Technologies for Vibration Rendering

Researchers have been exploring the role of haptic feedback in conveying emotions for decades [13]. In the Human-Computer Interaction (HCI) community, research has explored the potential of vibrotactile patterns to convey emotions [2, 11, 17]. Indeed, affective haptics have been increasingly utilized in various domains, including storytelling [15], mental therapy [18], film [21], virtual reality [33, 35], and enhancing visual-audio experiences [1, 12, 29, 32]. However, despite the diversity of applications of haptic feedback, it is still difficult to design vibrations to convey emotions.

Technologies for rendering vibrations have also advanced over the past decades. For example, by coding Pulse Width Modulation (PWM) outputs, hardware such as vibration motors can convey various vibrations [21] and are frequently integrated into handheld devices such as mice [11, 36] and game controllers [9]. Some other vibration data libraries [15, 31] have been published for standardization. Specifically, Hasti Seifi's team has developed several open-sourced haptic datasets [27, 28], employing the Waveform Audio File Format (WAVE) to present vibrations. Their work enables broader applications of vibrations across computational devices. On mobile devices, vibration rendering is further simplified by hardware producers. For example, Apple's Haptic and Audio Pattern (AHAP) [14] format allows developers to customize vibration patterns with text editors. This JSON-like text format opens possibilities for leveraging large language models (LLMs), offering a simpler and more efficient solution.

2.2 Emotion Quantification for Affective Computing

Designing systems to convey emotion through vibrations relies on being able to identify human emotional responses to these stimuli either through physiological signals or subjective scales. Commonly used physiological signals indicating human emotional response include Electroencephalography (EEG), Electrocardiogram (ECG), Galvanic Skin Response (GSR), and facial activity [5, 22, 25, 30]. As for subjective scales, the dimensions of valence and arousal [26] are most commonly applied in affective computing systems[6, 16, 20, 38]. Subjective scales such as the Self-Assessment Manikin (SAM) [3], utilize these valence-arousal dimensions to measure emotion, and are considered reliable and effective. Compared to monitoring emotion using physiological signals, using the SAM scale costs little and requires no specialized equipment, so is regarded as an affordable option.

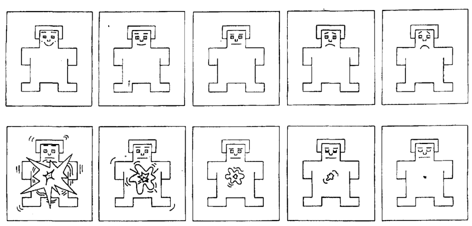

Figure 1: The Self-Assessment Manikin (SAM) used to rate the affective dimensions of valence (top panel) and arousal (bottom panel). Originally from [3]

3 Methodology

To evaluate GenAI methods for generating affective vibrations, research was undertaken split into three parts:

(1) Vibration-emotion dataset construction: section 3.1
(2) Model training using two different GenAI approaches for comparison: section 3.2
(3) Human-centred evaluation to compare the effectiveness of the two GenAI approaches to affective vibration generations: section 4

3.1 Vibration-Emotion Dataset Construction

Datasets of vibration data are limited. In Hasti Seifi's VibViz dataset [28], a large collection of vibrotactile effects was categorized with emotion labels. However, due to differences in rendering devices, as well as the limited number of annotators involved in the original dataset, it is necessary to conduct additional data annotation. The data annotation process in this paper follows methods from the VibViz project as well as research on building emotional music datasets [34].

The original VibViz dataset contains 120 vibration samples in WAVE audio format. Before annotation in our study, these vibrations were converted into the Apple Haptic and Audio Pattern (AHAP) format for rendering on mobile devices. Fifteen participants were then invited (M_{age} = 23.73; 9 males, 6 females) through online recruitment to annotate vibration patterns using a mobile app developed for the study. Participants were required to use a mobile phone preinstalled with the app and follow instructions provided by the contact person either in person or online.

On the data annotation app vibrations were presented one by one in random order. Five seconds after experiencing a single vibration, participants were asked to rate the valence and arousal of the emotion they perceived on a 9-point scale. There was a five-second break between each annotation to reduce potential fatigue. The user interfaces of this annotation flow can be seen in Figure 2. The mean and standard deviations of participants' ratings for each vibration were calculated. Vibrations with clearly high/low valence-arousal mean ratings or with low variability were retained for AI training.

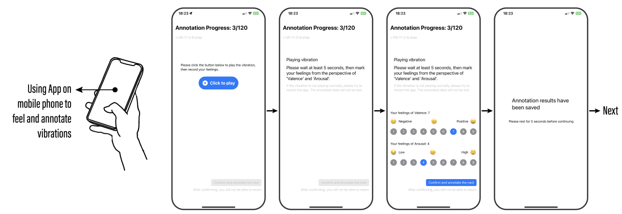

Figure 2: Interfaces of Data Annotation App and Annotation Procedure.

3.2 AI Models

Two of the most common forms of GenAI were selected to generate affective vibrations: audio-based generation and text-based generation as discussed in this section. Both approaches were trained on the dataset constructed in section 3.1.

3.2.1 VAE Approach: Audio-based Vibration Generation. Extensive research has been conducted on audio generation using emotion labels [6, 34, 37, 38]. Leveraging Variational Autoencoders (VAEs) [7] offers an approach to mapping annotated data into a multi-dimensional latent space [4] which is followed in this study for audio-based vibration generation by training the model on the vibration-emotion dataset from section 3.1.

In Figure 3 each spectrogram represents a vibration pattern generated by a pair of emotion-valence values. The horizontal axis represents time (in seconds), the vertical axis represents frequency (in Hz), and the color intensity represents the amplitude (in dB). These spectrograms visually demonstrate the variety and complexity of the generated vibration patterns, demonstrating that valence-arousal inputs have influenced the generated vibration outputs of the VAE.

3.2.2 LLM Approach: Text-based Vibration Generation. For our second approach we chose the popular ChatGPT API [23] to fine-tune the large language model (LLM) ChatGPT-3.5-turbo. The dataset from section 3.1 was converted into AHAP format in advance, ensuring compatibility with the fine-tuning requirements of the ChatGPT API. Compared to the original LLM, the fine-tuned LLM demonstrated a significant improvement in generating AHAP patterns tailored to emotional inputs.

In Figure 4, each diagram represents a vibration pattern generated by the fine-tuned LLM. The horizontal axis represents time (in seconds), and the vertical axis represents intensity ranging from 0.0 (lowest) to 1.0 (highest). These patterns demonstrate the diversity of the fine-tuned LLM's generated vibration outputs.

4 Evaluation

A user study was undertaken to evaluate the quality and emotion-expression accuracy of the two GenAI approaches as described in this section.

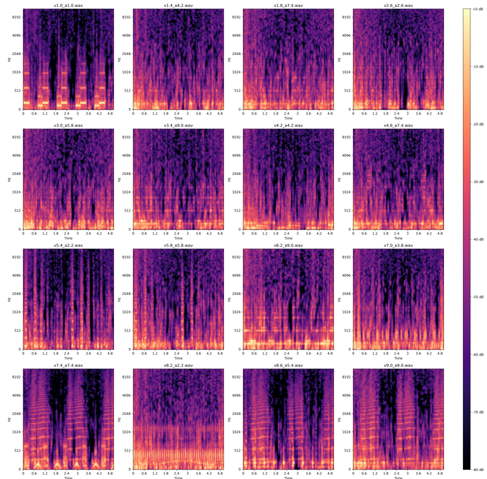

Figure 3: Spectrograms of VAE-generated vibration samples

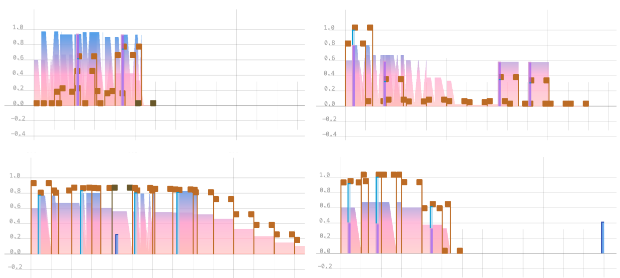

Figure 4: Vibration patterns of sampled generation outputs of the fine-tuned LLM

4.1 Participants, Materials, and Procedure

Fifteen participants were invited to the user study ($M_{\text{age}} = 24.0$ years; 7 males, 8 females). All participants were healthy adults with experience using mobile phones. The user study material consisted of two components: a mobile phone with a customized app preinstalled for displaying generated vibrations, and an online form on a laptop for listing tasks and collecting data (Figure 5).

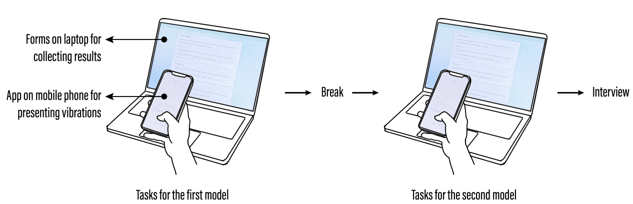

Figure 5: User Study Procedure

The key user interfaces of the app are shown in Figure 6, the app uses a grid system for haptic pattern selection. The haptic patterns are assigned into four quadrants, each corresponding to one of four emotions: happy (high-valence, high-arousal), angry

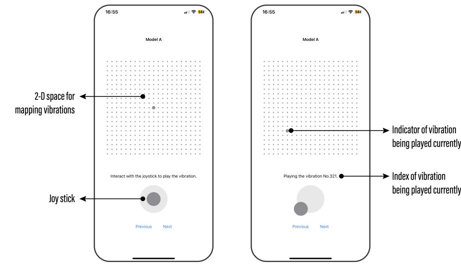

Figure 6: User interfaces of the user test app on the mobile phone

(low-valence, high-arousal), sad (low-valence, low-arousal), and relaxed (high-valence, low-arousal) but for the user study there is no visual indication of this arrangement in the user interface. The size of the grid is 21x21, which represents 441 generated vibrations mapped onto the 2-dimensional valence-arousal emotion space. The size and number are chosen to balance the accuracy of emotion expression and the cost of model generation. A virtual joystick below the emotion space enables continuous navigation across this grid, allowing participants to explore the 2-D space and sample the vibrations with real-time feedback.

The user's tasks were split into two groups in random order: VAE and LLM. In each group, participants were required to explore vibrations on the app using one hand and select vibrations that best represented one of the four emotions. For each emotion, two selections were required. During the study, the details of the dimension representation and models were not disclosed to the participants. To minimize fatigue, participants were given a 5-minute break between the two models' tasks.

A semi-structured interview was conducted afterward to qualitatively gather their feedback. The main interview questions were:

- What are the differences between two groups of vibrations based on your feelings? And which group of vibrations do you prefer?
- Were there any specific emotions that were difficult to identify using the vibrations? And what criteria did you use to identify emotions represented by the vibrations?
- Have you had any prior experience with applications or interactive designs involving vibrations?

4.2 Quantitive Analysis

All 15 participants completed the tasks. Their responses were analyzed to determine whether the emotion they perceived matched the quadrant to which the vibration belonged. By aggregating the results across all participants, the emotion quadrant matching accuracy was calculated as the ratio of the total number of vibrations correctly matched to the total number of annotated vibrations.

Figure 7 illustrates that both models performed better than random generation but not exceptionally well. The LLM was more effective in expressing high-arousal emotions (happy and angry), while the VAE was more effective in expressing low-valence emotions (sad and relaxed). Notably, the VAE performs significantly well in conveying relaxation.

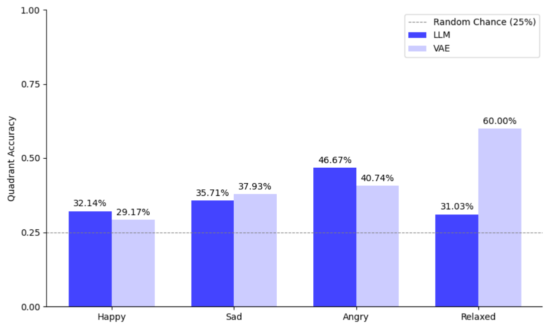

Figure 7: Emotion quadrant matching accuracies of LLM and VAE

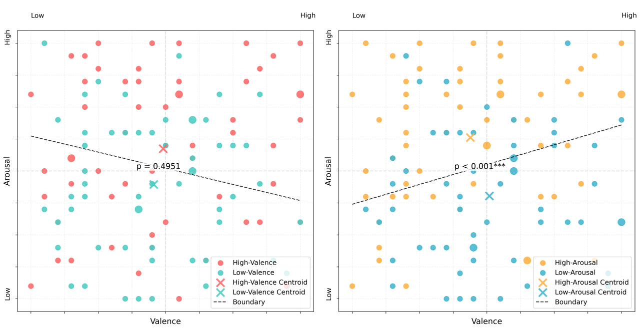

Figure 8: Distribution of LLM-generated vibrations selected by participants.

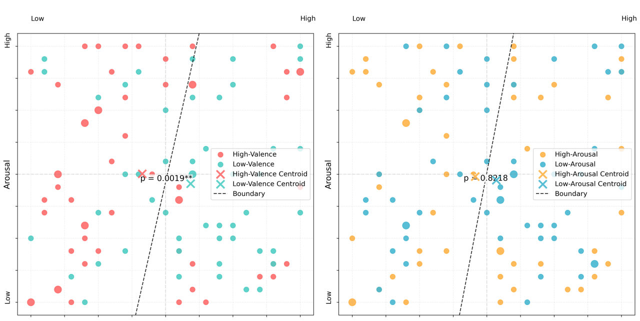

Figure 9: Distribution of VAE-generated vibrations selected by participants.

The labeled vibrations were further mapped onto the valence-arousal grid according to the value that generates them. Using centroid distribution and T-test, the deviations between the input emotional parameters and the participants' perceived emotions were revealed. For the LLM (Figure 8), the distribution showed difficulty in accurately reflecting valence. However, significant differences ($p < .001$) were observed in arousal between the low-arousal and high-arousal groups. For the VAE (Figure 9), the generated vibrations fail to distinguish arousal. Significant differences ($p = .0019$) were observed in valence groups but not in the correct distribution.

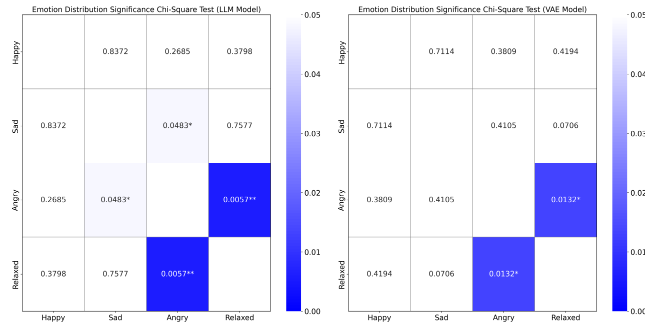

Figure 10: Pairwise significance differences between four emotions in both models

To investigate differences across the four emotion groups, CHI-squared tests were conducted. For the general differences among the four emotions, neither the LLM ($p = .1019$) nor the VAE ($p = .0995$) model shows significant differences. However, when comparing specific emotion pairs, significant differences were partially found (Figure 10). For the VAE, angry (low-valence, high-arousal) differed significantly from relaxed (high-valence, low-arousal) ($p = .0132$). For the LLM, significant differences were found between angry and relaxed ($p = .0057$) as well as between angry and sad (low-valence, low-arousal) ($p = .0483$).

Overall, both the LLM and VAE models demonstrate the potential for generating affective vibrations and exhibit strengths in specific aspects. The LLM was more effective in emotion expression especially for high-arousal, whereas the VAE performed better at conveying relaxation. However, both models have notable weaknesses, as neither of them was able to generate affective vibrations with high accuracy. In terms of overall performance, the LLM performed better. These insights were further validated during the interview sessions.

4.3 Qualitative Feedback Summarization

All 15 participants responded to our interviews. Using thematic analysis, we identified key insights by categorizing their text responses and extracting frequently mentioned keywords. Each participant's response has been labeled and referenced in the analysis.

Eight participants (P1, P4, P6, P7, P10, P11, P13, P14) described vibrations generated by the LLM as intense, clear, or easy to interpret, while vibrations from the VAE were considered as smooth and soft. Key participant feedback includes:

"Model B (LLM) is more intense, while A (VAE) is more gentle and subtle." (P1)

"Model A (LLM) is more direct, and Model B (VAE) is more smooth." (P6)

As shown in Figure 11, more than half of the participants preferred LLM-generated vibrations. Only one participant preferred VAE vibrations, and others expressed no clear preference.

In addition to evaluation between the VAE and LLM, features of affective vibrations were explored to identify design principles. Participants highlighted that certain vibration features like rhythm and intensity could effectively evoke specific emotions (P2, P8, P9, P11, P12). Strong, continuous vibrations were often associated with

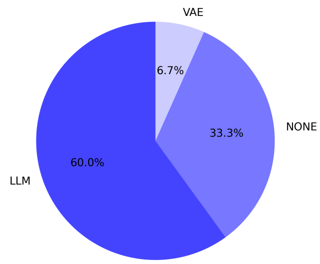

Figure 11: Participants' preference towards LLM and VAE

emotions like anger (high-arousal, low-valence). Rhythmic vibrations were linked to high-valence emotions (happy and relaxed). Participants found it easier to identify emotions like angry due to its high intensity, while relax and sad were more difficult to identify. Several participants noted that they associated vibrations with human behaviors, such as sobbing for sadness (P2, P3, P8). Table 1 summarizes the vibration features associated with four emotions.

Arousal-Valence	Intensity	Rhythm
High Arousal, Low Valence	High	Continuous
High Arousal, High Valence	High	Rhythmic
Low Arousal, High Valence	Low	Rhythmic
Low Arousal, Low Valence	Low	Like sobbing/crying

Table 1: Features of Vibrations for Four Emotions

For potential applications, most participants referred to vibrations provided by game controllers equipped with haptic engines. However, these are predominantly used for physical simulation rather than affective feedback. For example, "opening treasure chests" (P6) or "diving into ink in 'Splatoon'" (P4). Other vibration-related scenarios associated with affective feedback mentioned include music games (P15) and music performances (P9).

Despite the diverse perspectives on interacting with haptic feedback, we suggest the following guidance for designers of affective vibrations:

- Rhythms are essential for expressing emotional valence, while continuity and intensity are essential for expressing arousal as summarized in Table 1.
- Vibrations commonly convey negative emotions on mobile phones in daily life, while the opposite conclusion is found in entertaining scenarios such as gaming and music performances.
- Many users interpret vibrations through synesthesia and metaphor, such as associating certain vibrations with heartbeats, sobbing, or alarms e.g. "In judging vibrations, I need to rely on associations to understand emotions" (P14).
- Based on keyword counts of participants' responses, vibrations are more effective in expressing emotions with high arousal or low valence.
- Participants with professional musical backgrounds highlighted the similarities of the emotional expression among music rhythms, drum patterns, and vibrations. "Those with normal rhythms feel positive, those with abnormal rhythms do not" (P11).

When choosing between VAE and LLM from the perspective of designers and users, we suggest that for most scenarios utilizing vibrotactile feedback, fine-tuned LLM performs better because of its diverse outputs. VAE would be an alternative in specific scenarios such as relaxation. Moreover, in scenarios requiring large-scale AI-generated content, using LLM may incur substantial costs in terms of time, energy, and financial resources.

5 Limitations and Future Work

Our study on the AHAP format limits non-iOS devices from using the AI-generated vibrations, future iterations will extend compatibility across platforms. Since AHAP describes vibrations through structured, human-readable text, it provides the potential for systematic analysis of AI-generated vibrations and platform-independent vibrotactile design guidelines.

Compared to the 5-10 participants in relevant projects [28, 34], our project engaged a larger sample (15), but remains insufficient to fully eliminate potential biases and randomness. However, starting from an open-sourced dataset of 120 WAVE-format clips, we expanded it to 441 samples in both WAVE and AHAP formats. Through user studies, these vibrations were assessed and filtered, and those with clear emotion patterns were selected.

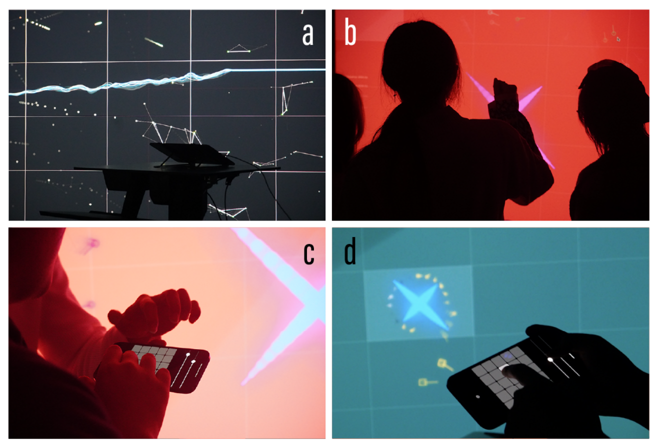

Figure 12: Records of the music performance, taken by the author. (a) An overview of the dynamic background of music visualization and the main DJ console. (b) The audience members engaging at the music performance. (c) Two audience members gathering to interact with the remote-control app while experiencing the rendered vibrations. (d) A single audience member interacting with the music via the remote-control app while experiencing vibrations.

As an example application, these selected vibrations were used in an interactive music performance (Figure 12), where each audience member could act as a DJ using a mobile phone remote-control app. By rendering vibrations with a valence-arousal space, the app establishes connections between music filters and vibrations. Audience feedback suggested that the vibrations offered a novel sensory channel, enriching the emotional atmosphere and enhancing the sense of engagement. This application illustrates the potential of AI-generated vibrotactile feedback in future creative practices.

These selected vibrations may also contribute to future research of computational haptic design. Our future work will concentrate on two directions: i) iterations to models and datasets, aiming to develop a generative system capable of accurately mapping vibrations onto the valence-arousal space; ii) quantitative evaluations on the insights and phenomena identified through user studies. These could evolve into focused research topics, ultimately contributing to design principles that advance human-computer interaction.

6 Conclusions

This project contributes to the field of affective haptics from both technical and design perspectives. With our vibration-emotion dataset, two AI-based approaches for vibration generation were investigated: an audio-based method using Variational Autoencoder (VAE) and a text-based method using fine-tuned large language model (LLM). Through a human-centered study, the VAE's ability to express certain emotion and the LLM's ability to express arousal was examined. From the interview, five insights for designers and researchers were provided to further explore the design of vibrations. Despite limitations such as dataset and sample size, the findings demonstrate the potential of these methods for advancing the computational generation and application of affective vibrations.

References

[1] Akshita, Harini Alagarai Sampath, Bipin Indurkhya, Eunhwa Lee, and Yudong Bae. 2015. Towards Multimodal Affective Feedback: Interaction between Visual and Haptic Modalities *(CHI '15)*. Association for Computing Machinery, 2043–2052. doi:10.1145/2702123.2702288

[2] Ahmad Hoirul Basori, Daut Daman, Mohd Shahrizal Sunar, and Abdullah Bade. 2008. The Potential of Human Haptic Emotion as Technique for Virtual Human Characters Movement to Augment Interactivity in Virtual Reality Game. *IJVR* 7 (2008), 27–32. Retrieved November 24, 2024 from https://www.researchgate.net/publication/220222188_The_Potential_of_Human_Haptic_Emotion_as_Technique_for_Virtual_Human_Characters_Movement_to_Augment_Interactivity_in_Virtual_Reality_Game.

[3] Margaret M. Bradley and Peter J. Lang. 1994. Measuring emotion: The self-assessment manikin and the semantic differential. *Journal of Behavior Therapy and Experimental Psychiatry* 25, 1 (1994), 49–59. doi:10.1016/0005-7916(94)90063-9

[4] Nick Bryan-Kinns, Bingyuan Zhang, Songyan Zhao, and Berker Banar. 2024. Exploring Variational Auto-encoder Architectures, Configurations, and Datasets for Generative Music Explainable AI. *Machine Intelligence Research* 21 (2024), 29–45. doi:10.1007/s11633-023-1457-1

[5] Aaron Frederick Bulagang, Ng Giap Weng, James Mountstephens, and Jason Teo. 2020. A review of recent approaches for emotion classification using electrocardiography and electrodermography signals. *Informatics in Medicine Unlocked* 20 (2020), 100363. doi:10.1016/j.imu.2020.100363

[6] Adyasha Dash and Kathleen Agres. 2024. AI-Based Affective Music Generation Systems: A Review of Methods and Challenges. *ACM Comput. Surv.* 56, 11, Article 287 (July 2024), 34 pages. doi:10.1145/3672554

[7] Carl Doersch. 2021. Tutorial on Variational Autoencoders. doi:10.48550/arXiv.1606.05908 arXiv:1606.05908

[8] Mohamad A. Eid and Hussein Al Osman. 2016. Affective Haptics: Current Research and Future Directions. *IEEE Access* 4 (2016), 26–40. doi:10.1109/ACCESS.2015.2497316 Conference Name: IEEE Access.

[9] Sony Interactive Entertainment. [n. d.]. *DualSense wireless controller.* Retrieved 25 November, 2024 from https://www.playstation.com/en-us/accessories/dualsense-wireless-controller/.

[10] Uwe Gruenefeld, Alexander Geilen, Jonathan Liebers, Nick Wittig, Marion Koelle, and Stefan Schneegass. 2022. ARm Haptics: 3D-Printed Wearable Haptics for Mobile Augmented Reality. *Proceedings of the ACM on Human-Computer Interaction* 6 (2022), 193:1–193:18. Issue MHCI. doi:10.1145/3546728

[11] Tampereen Yliopisto Hallikas, Grigori Evreinov, Leena Vesterinen, Antti Nyman, Jalo Kääminen, Jarno Jokinen, and Deepa Vasara. 2005. vSmileys: Imaging Emotions through Vibration Patterns. (2005). Retrieved November 29, 2024 from https://www.researchgate.net/publication/277297050_vSmileys_Imaging_Emotions_through_Vibration_Patterns.

[12] Lauren Hayes. 2015. Skin Music (2012): an Audio-Haptic Composition for Ears and Body. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition* (Glasgow, United Kingdom) *(C&C '15)*. Association for Computing Machinery, New York, NY, USA, 359–360. doi:10.1145/2757226.2757370

[13] Matthew J. Hertenstein, Dacher Keltner, Betsy App, Brittany A. Bulleit, and Ariane R. Jaskolka. 2006. Touch communicates distinct emotions. *Emotion (Washington, D.C.)* 6, 3 (2006), 528–533. doi:10.1037/1528-3542.6.3.528

[14] Apple Inc. no date. *Representing haptic patterns in AHAP files.* Retrieved 25 November, 2024 from https://developer.apple.com/documentation/corehaptics/representing-haptic-patterns-in-ahap-files.

[15] Ali Israr, Siyan Zhao, Kaitlyn Schwalje, Roberta Klatzky, and Jill Lehman. 2014. Feel Effects: Enriching Storytelling with Haptic Feedback. *ACM Trans. Appl. Percept.* 11, 3 (2014), 11:1–11:17. doi:10.1145/2641570

[16] Yang Jiao and Yingqing Xu. 2020. Affective Haptics and Multimodal Experiments Research, Masaaki Kurosu (Ed.). Springer International Publishing, 380–391. doi:10.1007/978-3-030-49062-1_26

[17] Hae youn Joung and Ellen Yi-Luen Do. 2011. Study of affective communication method in tactile hand gesture feedback. In *Proceedings of the 8th ACM Conference on Creativity and Cognition* (Atlanta, Georgia, USA) *(C&C '11)*. Association for Computing Machinery, New York, NY, USA, 351–352. doi:10.1145/2069618.2069691

[18] Chelsea Kelling, Daniella Pitaro, and Jussi Rantala. 2016. Good vibes: the impact of haptic patterns on stress levels *(AcademicMindtrek '16)*. Association for Computing Machinery, 130–136. doi:10.1145/2994310.2994368

[19] Karon E MacLean. 2022. Designing affective haptic experience for wellness and social communication: where designers need affective neuroscience and psychology. *Current Opinion in Behavioral Sciences* 45 (2022), 101113. doi:10.1016/j.cobeha.2022.101113

[20] Sara Marques-Villarroya, Juan Jose Gamboa-Montero, Cristina Jumela-Yedra, Jose Carlos Castillo, and Miguel Angel Salichs. 2022. Affect Display Recognition Through Tactile and Visual Stimuli in a Social Robot, Filippo Cavallo, John-John Cabibihan, Laura Fiorini, Alessandra Sorrentino, Hongsheng He, Xiaorui Liu, Yoshio Matsumoto, and Shuzhi Sam Ge (Eds.). Springer Nature Switzerland, 130–140. doi:10.1007/978-3-031-24667-8_12

[21] Antonella Mazzoni and Nick Bryan-Kinns. 2016. Mood Glove: A haptic wearable prototype system to enhance mood music in film. *Entertainment Computing* 17 (2016), 9–17. doi:10.1016/j.entcom.2016.06.002

[22] Ali Momennezhad. 2018. EEG-based emotion recognition utilizing wavelet coefficients. *Multimedia Tools and Applications* 77 (2018), 1–18. doi:10.1007/s11042-018-5906-8

[23] OpenAI. no date. *OpenAI Platform.* Retrieved 25 November, 2024 from https://platform.openai.com.

[24] Jia Qing Poh, John See, Neamat El Gayar, and Lai-Kuan Wong. 2024. Are You Paying Attention? Multimodal Linear Attention Transformers for Affect Prediction in Video Conversations. In *Proceedings of the 2nd International Workshop on Multimodal and Responsible Affective Computing* (Melbourne VIC, Australia) *(MRAC '24)*. Association for Computing Machinery, New York, NY, USA, 15–23. doi:10.1145/3689092.3689409

[25] Aasim Raheel, Muhammad Majid, and Syed Muhammad Anwar. 2021. DEAR-MULSEMEDIA: Dataset for emotion analysis and recognition in response to multiple sensorial. *Information Fusion* 65 (2021), 37–49. doi:10.1016/j.inffus.2020.08.007

[26] James A. Russell and Lisa Feldman Barrett. 1999. Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology* 76, 5 (1999), 805–819. doi:10.1037/0022-3514.76.5.805 Place: US Publisher: American Psychological Association.

[27] Hasti Seifi, Farimah Fazlollahi, Michael Oppermann, John Andrew Sastrillo, Jessica Ip, Ashutosh Agrawal, Gunhyuk Park, Katherine J. Kuchenbecker, and Karon E. MacLean. 2019. Haptipedia: Accelerating Haptic Device Discovery to Support Interaction & Engineering Design *(CHI '19)*. Association for Computing Machinery, 1–12. doi:10.1145/3290605.3300788

[28] Hasti Seifi, Kailun Zhang, and Karon E. MacLean. 2015. VibViz: Organizing, visualizing and navigating vibration libraries. 254–259. doi:10.1109/WHC.2015.7177722

[29] Sichao Song and Seiji Yamada. 2017. Expressing Emotions through Color, Sound, and Vibration with an Appearance-Constrained Social Robot *(HRI '17)*. Association for Computing Machinery, 2–11. doi:10.1145/2909824.3020239

[30] Ramanathan Subramanian, Julia Wache, Mojtaba Khomami Abadi, Radu L. Vieriu, Stefan Winkler, and Nicu Sebe. 2018. ASCERTAIN: Emotion and Personality Recognition Using Commercial Sensors. *IEEE Transactions on Affective Computing* 9, 2 (2018), 147–160. doi:10.1109/TAFFC.2016.2625250 Conference Name: IEEE Transactions on Affective Computing.

[31] David Ternes and Karon E. MacLean. 2008. Designing Large Sets of Haptic Icons with Rhythm, Manuel Ferre (Ed.). Springer, 199–208. doi:10.1007/978-3-540-69057-3_24

[32] Yacine Tsalamlal, Michel-Ange Amorim, Jean-Claude Martin, and Mehdi Ammi. 2018. Modeling Emotional Valence Integration From Voice and Touch. *Frontiers in Psychology* 9 (2018), 1966.

[33] Ryoko Ueoka, Ali AlMutawa, and Hikaru Katsuki. 2016. Emotion hacking VR (EH-VR): amplifying scary VR experience by accelerating real heart rate using false vibrotactile biofeedback *(SA '16)*. Association for Computing Machinery, 1–2. doi:10.1145/2988240.2988247

[34] Yajie Wang, Mulin Chen, and Xuelong Li. 2024. Continuous Emotion-Based Image-to-Music Generation. *IEEE Transactions on Multimedia* 26 (2024), 5670–5679. doi:10.1109/TMM.2023.3338089 Conference Name: IEEE Transactions on Multimedia.

[35] Zhizhuo Yang and Dongdong Weng. 2016. Passive haptics based MR system for geography teaching *(VRCAI '16)*. Association for Computing Machinery, 23–29. doi:10.1145/3013971.3013995

[36] Gyeore Yun, Minjae Mun, Jungeun Lee, Dong-Geun Kim, Hong Z Tan, and Seungmoon Choi. 2023. Generating Real-Time, Selective, and Multimodal Haptic Effects from Sound for Gaming Experience Enhancement *(CHI '23)*. Association for Computing Machinery, 1–17. doi:10.1145/3544548.3580787

[37] Jincheng Zhang, György Fazekas, and Charalampos Saitis. 2023. Fast Diffusion GAN Model for Symbolic Music Generation Controlled by Emotions. doi:10.48550/arXiv.2310.14040 arXiv:2310.14040 [cs, eess]

[38] Sicheng Zhao, Yaxian Li, Xingxu Yao, Weizhi Nie, Pengfei Xu, Jufeng Yang, and Kurt Keutzer. 2020. Emotion-Based End-to-End Matching Between Image and Music in Valence-Arousal Space *(MM '20)*. Association for Computing Machinery, 2945–2954.

[39] Ran Zhou, Zachary Schwemler, Akshay Baweja, Harpreet Sareen, Casey Lee Hunt, and Daniel Leithinger. 2023. TactorBots: A Haptic Design Toolkit for Out-of-lab Exploration of Emotional Robotic Touch. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 370, 19 pages. doi:10.1145/3544548.3580799